

Globally Injective ReLU Networks

Michael Puthawala

MAP19@RICE.EDU

*Department of Computational and Applied Mathematics
Rice University, MS-134
6100 Main St.
Houston, TX, 77005, USA*

Konik Kothari

KKOTHAR3@ILLINOIS.EDU

*Department of Computer Science
University of Illinois at Urbana-Champaign
Champaign, IL, 61820, USA*

Matti Lassas

MATTI.LASSAS@HELSINKI.FI

*Department of Mathematics and Statistics
University of Helsinki
P.O. Box 68, FI-00014 University of Helsinki, Finland*

Ivan Dokmanić

IVAN.DOKMANIC@UNIBAS.CH

*Department of Mathematics and Computer Science
University of Basel
4051 Basel, Switzerland*

Maarten de Hoop

MDEHOOP@RICE.EDU

*Department of Computational and Applied Mathematics
Rice University, MS-134
6100 Main St.
Houston, TX, 77005, USA*

Editor: Ruslan Salakhutdinov

Abstract

Injectivity plays an important role in generative models where it enables inference; in inverse problems and compressed sensing with generative priors it is a precursor to well posedness. We establish sharp characterizations of injectivity of fully-connected and convolutional ReLU layers and networks. First, through a layerwise analysis, we show that an expansivity factor of two is necessary and sufficient for injectivity by constructing appropriate weight matrices. We show that global injectivity with iid Gaussian matrices, a commonly used tractable model, requires larger expansivity between 3.4 and 10.5. We also characterize the stability of inverting an injective network via worst-case Lipschitz constants of the inverse. We then use arguments from differential topology to study injectivity of deep networks and prove that any Lipschitz map can be approximated by an injective ReLU network. Finally, using an argument based on random projections, we show that an end-to-end—rather than layerwise—doubling of the dimension suffices for injectivity. Our results establish a theoretical basis for the study of nonlinear inverse and inference problems using neural networks.

Keywords: bilipschitz networks, injective generative models, random neural networks, deep learning, universal approximation

1. Introduction

Many applications of deep neural networks require inverting them on their range. Given a neural network $N : \mathcal{Z} \rightarrow \mathcal{X}$, where \mathcal{X} is often the Euclidean space \mathbb{R}^m and \mathcal{Z} is a lower-dimensional space, the map $N^{-1} : N(\mathcal{Z}) \rightarrow \mathcal{Z}$ is only well-defined when N is injective. The issue of injectivity is particularly salient in two applications: generative models and (nonlinear) inverse problems.

Generative networks model a complicated distribution p_X over \mathcal{X} as a pushforward of a simple distribution p_Z through N . Given an x in the range of N , inference requires computing $p_Z(N^{-1}(x))$ which is well-posed only when N is injective. In the analysis of inverse problems Arridge et al. (2019), uniqueness of a solution is a key concern; it is tantamount to injectivity of the forward operator. Given a forward model that is known to yield uniqueness, a natural question is whether we can design a neural network that approximates it arbitrarily well while preserving uniqueness. Similarly, in compressed sensing with a generative prior N and a possibly nonlinear forward operator A injective on the range of N , we seek a latent code z such that $A(N(z))$ is close to some measured $y = A(x)$. This is again only well-posed when N can be inverted on its range (Balestriero et al., 2020). Beyond these motivations, injectivity is a fundamental mathematical property with numerous implications. We mention a notable example: certain injective generators can be trained with sample complexity that is polynomial in the image dimension (Bai et al., 2018).

1.1 Our Results

In this paper we study injectivity of neural networks with ReLU activations. Our contributions can be divided into layerwise results and multilayer results.

1.1.1 LAYERWISE RESULTS

For a ReLU layer $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$ we derive sufficient and necessary conditions for invertibility on the range. For the first time, we construct deterministic injective ReLU layers with minimal expansivity $m = 2n$. We then derive specialized results for convolutional layers which are given in terms of filter kernels instead of weight matrices. We also prove upper and lower bounds on minimal expansivity of globally injective layers with iid Gaussian weights. This generalizes certain existing pointwise results (Theorem 7 and Appendix A.2). We finally derive the worst-case inverse Lipschitz constant for an injective ReLU layer which yields stability estimates in applications to inverse problems.

1.1.2 MULTILAYER RESULTS

A natural question is whether injective models are sufficiently expressive. Using techniques from differential topology we prove that injective networks are *universal* in the following sense: if a neural network $N_1 : \mathcal{Z} \rightarrow \mathbb{R}^{2n+1}$ models the data, $\mathcal{Z} \subset \mathbb{R}^n$, then we can approximate N_1 by an injective neural network $N_2 : \mathcal{Z} \rightarrow \mathbb{R}^{2n+1}$. As N_2 is injective, the image set $N_2(\mathcal{Z})$ is a Lipschitz manifold. We then use an argument based on random projections to show that an *end-to-end* expansivity by a factor of ≈ 2 is enough for injectivity in ReLU networks, as opposed to layerwise 2-expansivity implied by the layerwise analysis.




Figure 1: An illustration of a ReLU layer $N : \mathbb{R}^2 \rightarrow \mathbb{R}^3$, $x = N(z)$, that is not globally injective. Differently colored regions in the z -space are mapped to regions of the same color in the x -space. While N is locally injective in the pink, blue and green wedges in z -space, the orange, brown, and violet wedges are mapped to coordinate axes. N is thus *not* injective on these wedges. This prevents construction of an inverse in the range of N .

We conclude with preliminary numerical experiments to show that imposing injectivity improves inference in GANs while preserving expressivity.

1.2 Why Global Injectivity?

The attribute “global” relates to global injectivity of the map $N : \mathcal{Z} \rightarrow \mathbb{R}^m$ on the low-dimensional latent space \mathcal{Z} , but it does not imply global invertibility over \mathbb{R}^m , only on the range $N(\mathcal{Z}) \subset \mathbb{R}^m$. If we train a GAN generator to map iid normal latent vectors to real images from a given distribution, we expect that any sampled latent vector generates a plausible image. We thus desire that any $N(z)$ be produced by a unique latent code $z \in \mathcal{Z}$. This is equivalent to global injectivity, or invertibility on the range. Our results relate to the growing literature on using neural generative models for compressed sensing (Bora et al., 2017; Daskalakis et al., 2020). They parallel the related guarantees for sparse recovery where the role of the low-dimensional latent space is played by the set of all k -sparse vectors. One then looks for matrices which map *all* k -sparse vectors to distinct measurements (Foucart and Rauhut, 2013). As an example, in the illustration in Figure 1 images corresponding to latent codes in orange, brown, and violet wedges cannot be compressively sensed. Finally, a neural network is often trained to directly reconstruct an image x from its (compressive) low-dimensional measurements $y = A(x)$ without introducing any generative models. In this case, whenever A is Lipschitz, it is immediate that the learned inverse must be injective.

1.3 Related Work

Closest to our work are the papers of Bruna et al. (2013), Hand et al. (2018) and Lei et al. (2019). Bruna et al. (2013) study injectivity of pooling motivated by the problem of signal recovery from feature representations. They focus on ℓ^p pooling layers; their Proposition 2.2 gives a criterion similar to the DSS (Definition 1) and bi-Lipschitz bounds for a ReLU layer

(similar to our Theorem 8). Unlike Theorems 2 and 8, their criterion and Lipschitz bound are in some cases not precisely aligned with injectivity; see Appendix E.1.

Compressed sensing with GAN priors requires inverting the generator on its range (Bora et al., 2017; Shah and Hegde, 2018; Wu et al., 2019; Mardani et al., 2018; Hand et al., 2018; Aberdam et al., 2020; Daskalakis et al., 2020). Lei et al. (2019); Daskalakis et al. (2020) replace the end-to-end inversion by the faster and more accurate layerwise inversion when each layer is injective. They show that with high probability a ReLU layer with an iid normal weight matrix can be inverted about a fixed point if the layer expands at least by a factor of 2.1, and that the inverse about a fixed point can be computed with the algorithm from Bora et al. (2017). This result is related to our Theorem 7 which gives conditions for global injectivity or layers with random matrices. Hand and Voroninski (2017) show that when the weights of a ReLU network obey a certain weighted distribution condition, the loss function for the inversion has a strict descent direction everywhere except in a small set. The condition is in particular satisfied by random matrices with expansivity $n_j = \Omega(n_{j-1} \log n_{j-1})$, where n_j is the output dimension of layer j .

A continuous analogy of our convolutional construction (Definition 11) was considered by Mallat et al. (2018). They show that ReLU acts as a phase filter and that the layer is bi-Lipschitz and hence injective when the filters have a diverse phase and form a frame. Discretizing their model gives a statement related to Corollary 6 and Theorem 13.

Injectivity is automatic in invertible neural networks such as normalizing flows (Grover et al., 2018; Kingma and Dhariwal, 2018; Grathwohl et al., 2018). Specialized architectures with simple Jacobians give easy access to the likelihood (Dinh et al., 2014, 2016; Gomez et al., 2017), which facilitates application to inverse problems (Ardizzone et al., 2018). Inference with GANs can be achieved by jointly training a generative network and its inverse (Donahue et al., 2016; Dumoulin et al., 2016), which is well-defined when the generator is injective. Relaxed injective probability flows resemble GAN generators but are trained via approximate maximum likelihood (Kumar et al., 2020). Injectivity is promoted by keeping the smallest singular value of the Jacobian away from zero at the training examples, a necessary but not sufficient condition. In general, Jacobian conditioning improves GAN performance (Heusel et al., 2017; Odena et al., 2018). Finally, lack of injectivity interferes with disentanglement (Chen et al., 2016; Lin et al., 2019). In this context, injectivity seems to be a natural heuristic to increase latent space capacity without increasing its dimension (Brock et al., 2016).

1.4 Notation

Given a matrix $W \in \mathbb{R}^{m \times n}$ we define the notation $w \in W$ to mean that $w \in \mathbb{R}^n$ is a row vector of W . For a matrix $W \in \mathbb{R}^{m \times n}$ with rows $\{w_j\}_{j=1}^m$ and $x \in \mathbb{R}^n$, we write

$$S(x, W) := \{j \in [[m]] : \langle w_j, x \rangle \geq 0\}$$

and $S^c(x, W)$ for its complement, with $[[m]] = \{1, \dots, m\}$. Further, given a matrix W and index set \mathcal{I} we define the notation $W|_{\mathcal{I}}$ to be a matrix whose i 'th row is the same as W if $i \in \mathcal{I}$, and a row of zeroes otherwise. We let $\mathcal{NN}(n, m, L, \mathbf{n})$ be the family of functions $N_{\theta} : \mathbb{R}^n \rightarrow \mathbb{R}^m$ of the form

$$N(z) = W_{L+1} \phi_L(W_L \cdots \phi_2(W_2 \phi_1(W_1 z + b_1) + b_2) \cdots + b_L) \quad (1)$$

Indices $\ell = 1, \dots, L$ index the network layers, $b_\ell \in \mathbb{R}^{n_\ell+1}$ are the bias vectors, $W_\ell \in \mathbb{R}^{n_{\ell+1} \times n_\ell}$ are the weight matrices with $n_1 = n$, $n_L = m$, and ϕ_ℓ are the nonlinear activation functions. We will denote $\text{ReLU}(x) = \max(x, 0)$. We write $\mathbf{n} = (n_1, n_2, \dots, n_{L-1})$ and $\theta = (W_1, b_1, \dots, W_L, b_L)$ for the parameters that determine the function N_θ . We also write $\mathcal{NN}(n, m, L) = \bigcup_{\mathbf{n} \in \mathbb{Z}^{L-1}} \mathcal{NN}(n, m, L, \mathbf{n})$ and $\mathcal{NN}(n, m) = \bigcup_{L \in \mathbb{Z}} \mathcal{NN}(n, m, L)$.

In Section 2.4 we use multi-index notation. Multi-indices are p -tuples denoted by the capitol letters, such as I, J, K, N and M . In the case that $p = 1$, the operations between multi-indices is equivalent to the equivalent operations for scalars. If I is an p -dimensional multi-index, then

$$I := (I_1, \dots, I_p) \in \mathbb{N}^p$$

where I_p refers to the p 'th index of I . We use the notation $c \in \mathbb{R}^N$ to refer to $c \in \mathbb{R}^{N_1} \times \dots \times \mathbb{R}^{N_p}$, and $\sum_{I=1}^N c_I = \sum_{I_1=1}^{N_1} \dots \sum_{I_p=1}^{N_p} c_{I_1, \dots, I_p}$, for $I = (I_1, \dots, I_p)$. The symbol 1 can refer to number 1 or a p -tuple $(1, \dots, 1)$, and is clear from context. We likewise use the notation $I \leq J$ to mean that $I_k \leq J_k$ for all $k = 1, \dots, p$ and likewise for $\geq, <$ and $>$. If $I \not\leq J$, then there is at least one k such that $I_k > J_k$.

We use the notation $\|\cdot\|_{p,m}$ and $\|\cdot\|_{p,n}$ to denote the standard l_p norms for vectors in \mathbb{R}^m and \mathbb{R}^n respectively.

2. Layerwise Injectivity of ReLU Networks

For a one-to-one activation function such as a leaky ReLU or a tanh, it is easy to see that injectivity of $x \mapsto W_i x$ implies the injectivity of the layer. We therefore focus on non-injective ReLU activations.

2.1 Directed Spanning Set

Unlike in the case of a one-to-one activation, when $\phi(x) = \text{ReLU}(x)$, $x \mapsto \phi(Wx)$ cannot be injective for all $x \in \mathbb{R}^n$ if W is a square matrix. Noninjectivity of $\text{ReLU}(Wx)$ occurs in some simple situations, for example if Wx is not full rank or if $Wx \leq 0$ element-wise for a nonzero x , but these two conditions on their own are not sufficient to characterize injectivity. Figure 2 gives an example of a weight matrix that yields a noninjective layer that has these two properties. In order to facilitate the full characterization of injectivity of ReLU layers, we define a useful theoretical device:

Definition 1 (Directed Spanning Set) *Let $W \in \mathbb{R}^{n \times m}$. We say that W has a directed spanning set (DSS) of $\Omega \subset \mathbb{R}^n$ with respect to a vector $x \in \mathbb{R}^n$ if there exists a \hat{W}_x such that each row vector of \hat{W}_x is a row vector of W ,*

$$\langle x, w_i \rangle \geq 0 \quad \text{for all } w_i \in \hat{W}_x,$$

and $\Omega \subset \text{span}(\hat{W}_x)$. When omitted, Ω is understood to be \mathbb{R}^n .

For a particular $x \in \mathbb{R}^n$, it is not hard to verify if a given $W = \{w_i\}_{i=1, \dots, m}$ has a DSS of Ω . One can simply let $\hat{W}_x = \{w_i \in \mathbb{R}^n : \langle x, w_i \rangle \geq 0\}$. Then, W has a DSS of Ω with respect to x if and only if \hat{W}_x spans Ω , which can be checked efficiently. Note that having full rank




Figure 2: Illustration of the DSS definition. *Left:* A configuration of 4 vectors in \mathbb{R}^2 that do not have a DSS w.r.t. all $x \in \mathbb{R}^2$. In this case, the vectors do not generate an injective layer. The set of labels indicate which w_j have positive inner product with vectors in the wedge; there are two wedges with only one such $\{w_j\}$; we have $\text{ReLU}(Wx_1) = \text{ReLU}(Wx_2)$. *Center:* A configuration where four vectors have a DSS for all $x \in \mathbb{R}^2$. These vectors correspond to a minimally-expansive injective layer; see Corollary 6. *Right:* A plot of $\|\text{ReLU}(Wx)\|_1$ where W is given as in the left figure. Note that $x \mapsto \text{ReLU}(Wx)$ is linear in every wedge.

is necessary for it to have a DSS of \mathbb{R}^n w.r.t. *any* $x \in \mathbb{R}^n$. It is however not sufficient. For example, $\text{Id}_n \in \mathbb{R}^{n \times n}$, the identity matrix in \mathbb{R}^n , is clearly full rank, but it doesn't have a DSS of \mathbb{R}^n w.r.t. $x = (-1, 0, \dots, 0) \in \mathbb{R}^n$. We also note that in this paper we focus on globally injective networks, and so Ω is always taken to be either \mathbb{R}^n or \mathbb{R}^p , where $p < n$. Further, this latter case is only discussed to eventually produce a statement showing that a matrix has a DSS on the full space \mathbb{R}^n . This only comes up in the context of our discussion on convolution layers in Section 2.4, specifically Theorem 13 and Lemma 14.

To check whether W has a DSS for all $x \in \mathbb{R}^n$, note that W partitions \mathbb{R}^n into open wedges S_k , $\mathbb{R}^n = \bigcup_k S_k$, with constant sign patterns. That is, for $x_1, x_2 \in S_k$, $\text{sign}(Wx_1) = \text{sign}(Wx_2)$. (See also the proof of Theorem 7 in Appendix A.2.2.) Checking whether W has a DSS for all x is equivalent to checking that for every wedge there are at least n vectors $W_k \subset W$ with positive sign, $\langle w, x \rangle > 0$ for $x \in S_k$, and that W_k spans \mathbb{R}^n . Since the number of wedges can be exponential in m and n (Winder, 1966) this suggests an exponential time algorithm. We also note the connection between DSS and spark in compressed sensing (Foucart and Rauhut, 2013), defined as the size of the smallest set of linearly dependent vectors in W . If every wedge has n or more positive signs, then full spark $n + 1$ is sufficient for W to have a DSS w.r.t all $x \in \mathbb{R}^n$. Computing spark is known to be NP-hard (Tillmann and Pfetsch, 2013); whether one can do better for DSS remains an open question.

2.2 Fully Connected Layer

The notion of a DSS immediately leads to our main result for fully connected layers.

Theorem 2 (Conditions for Injectivity of $\text{ReLU}(Wx)$) *Let $W \in \mathbb{R}^{m \times n}$ where $1 < n \leq m$ be a matrix with row vectors $\{w_j\}_{j=1}^m$, and $\text{ReLU}(y) = \max(y, 0)$. The function $\text{ReLU}(W(\cdot)): \mathbb{R}^n \rightarrow \mathbb{R}^m$ is injective if and only if W has a DSS w.r.t every $x \in \mathbb{R}^n$.*

The question of injectivity in the case when $b = 0$ and when $b \neq 0$ are very similar and, as Lemma 3 shows, the latter question is equivalent to the former on a restricted weight matrix.

Lemma 3 (Injectivity of $\text{ReLU}(Wx + b)$) *Let $W \in \mathbb{R}^{m \times n}$ and $b \in \mathbb{R}^m$. The function $\text{ReLU}(W(\cdot) + b): \mathbb{R}^n \rightarrow \mathbb{R}^m$ is injective if and only if $\text{ReLU}(W|_{b \geq 0} \cdot)$ is injective, where $W|_{b \geq 0} \in \mathbb{R}^{m \times n}$ is row-wise the same as W where $b_i \geq 0$, and is a row of zeroes when $b_i < 0$.*

Remark 4 (Injectivity of $\text{ReLU}(Wx)$, Positive x) *Without normalization strategies, inputs of all but the first layer are element-wise non-negative. One can ask whether $\text{ReLU}(Wx)$ is injective when x is restricted to be element-wise non-negative. Following the same argument as in Theorem 2, we find that W must have a DSS w.r.t. $x \in \mathbb{R}^n$ for every x that is element-wise non-negative. With normalizations strategies however (for example batch renormalization Ioffe, 2017) the full power of Theorem 2 and Lemma 3 may be necessary. In Section 2.5 we show that common normalization strategies such as batch, layer, or group normalization do not interfere with injectivity.*

Remark 5 (Global vs Restricted Injectivity) *Even when the conditions of Theorem 2 and Lemma 3 are not satisfied the network might be injective in some $\mathcal{X} \subset \mathbb{R}^n$. Indeed, $N(x) = \text{ReLU}(\text{Id}_n \cdot x)$ is in general not injective, but it is injective in $\mathcal{X} = \{x \in \mathbb{R}^n : x_i > 0 \text{ for } i = 1, \dots, n\}$, a convex and open set. Theorem 2 and Lemma 3 are the precise criteria for single-layer injectivity for all $x \in \mathbb{R}^n$.*

The above layerwise results imply a sufficient condition for injectivity of deep neural networks. If Theorem 2 applies to each layer of a network then the entire network must be injective end-to-end.

Note that layer-wise injectivity is sufficient but not necessary for injectivity of the entire network. Consider for example an injective layer $\text{ReLU}(W)$, and $N(x) = \text{ReLU}(\text{Id}_m \cdot \text{ReLU}(W(x)))$, where I is the identity matrix. Clearly,

$$\text{ReLU}(\text{Id}_m \cdot \text{ReLU}(W(x))) = \text{ReLU}(\text{ReLU}(W(x))) = \text{ReLU}(W(x))$$

so N is injective, but it fails to be injective at the $\text{ReLU}(I \cdot)$ layer.

An important implication of Theorem 2 is the following result on minimal expansivity.

Corollary 6 (Minimal Expansivity) *For any $W \in \mathbb{R}^{m \times n}$, $\text{ReLU}(W \cdot)$ is non-injective if $m < 2 \cdot n$. If $W \in \mathbb{R}^{2n \times n}$ satisfies Theorem 2 and Lemma 3, then up to row rearrangement W can be written as*

$$W = \begin{bmatrix} B \\ -DB \end{bmatrix} \tag{2}$$

where $B, D \in \mathbb{R}^{n \times n}$ and B is a basis and D a diagonal matrix with strictly positive diagonal entries.

While $m < 2n$ immediately precludes injectivity, Corollary 6 gives a simple recipe for the construction of minimally-expansive injective layers with $m = 2n$. To the best of our

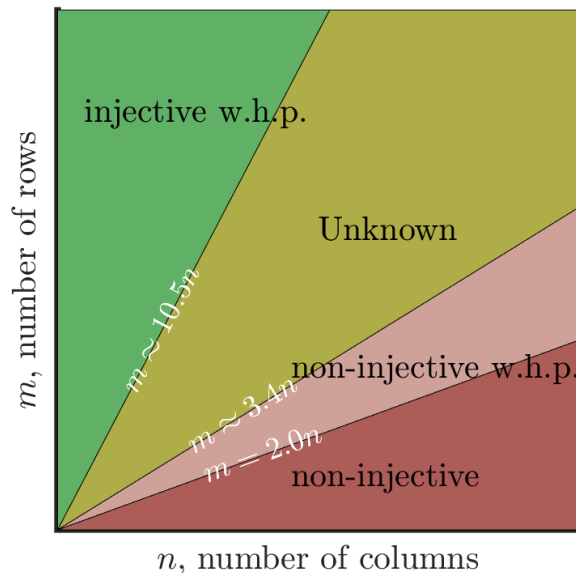


Figure 3: A visualization of regions where asymptotics of $\mathcal{I}(m, n)$ in Theorem 7 are valid. knowledge this is the first such result. Further, we can also use Corollary 6 to build weight matrices for which $m > 2n$. For example, the matrix

$$W = \begin{bmatrix} B \\ -DB \\ M \end{bmatrix},$$

where $B \in \mathbb{R}^{n \times n}$ has full rank, $D \in \mathbb{R}^{n \times n}$ is a positive diagonal matrix, and $M \in \mathbb{R}^{(m-2n) \times n}$ is arbitrary, has a DSS w.r.t. every $x \in \mathbb{R}^n$. This method can be used in practice to generate injective layers by using standard spectral regularizers to ensure that B is a basis (Cisse et al., 2017). In Section 3 we will show that, in fact, an *end-to-end* rather than layerwise doubling of dimension is sufficient for injectivity.

Previous work that sought injectivity uses random-weight models. Lei et al. (2019) show that a layer is invertible *about a point* in the range provided that $m \geq 2.1n$ and W is iid Gaussian. In Appendix 4.1 we show that $m = 2.1n$ is not enough to guarantee injectivity. In fact, we show that with high probability an iid Gaussian weight matrix W yields a *globally* invertible ReLU layer, when W is sufficiently expansive, and conversely does not satisfy Theorem 2 if it is not expansive by at least a factor of 3.4:

Theorem 7 (Injectivity for Gaussian Weights) *Let*

$\mathcal{I}(m, n) = \mathbb{P}\{x \mapsto \text{ReLU}(Wx) \text{ is injective}\}$ *with the entries of } W iid standard normal and } $c = m/n$ fixed. Then as } $n \rightarrow \infty$,*

$$\mathcal{I}(m, n) \rightarrow 1 \text{ if } c \geq 10.5 \text{ and } \mathcal{I}(m, n) \rightarrow 0 \text{ if } c \leq 3.3$$

The parameter regions defined in Theorem 7 are illustrated in Figure 3. We note that the study of neural networks with random weights is an important recent research direction (Pennington and Worah, 2017; Benigni and P  ch  , 2019). Pennington and Worah (2017)

articulate this clearly, arguing that large complex systems, of which deep networks are certainly an example, can be profitably studied by approximating their constituent parts as random variables.

2.3 Inverse Lipschitz Constant for Networks

Because $\|\text{ReLU}(Wx) - \text{ReLU}(Wy)\| \leq \|W\| \|x - y\|$, it is clear that $\text{ReLU}(W\cdot)$ is Lipschitz with constant $\|W\|$; whether the inverse is Lipschitz, and if so, with what Lipschitz constant, is less obvious. We can prove the following result (see Appendix C):

Theorem 8 (Global Inverse Lipschitz Constant) *Let $W \in \mathbb{R}^{m \times n}$ have a DSS w.r.t. every $x \in \mathbb{R}^n$. There exists a $C(W) > 0$ such that for any $x_0, x_1 \in \mathbb{R}^n$,*

$$\|\text{ReLU}(Wx_0) - \text{ReLU}(Wx_1)\|_2 \geq C(W) \|x_0 - x_1\|_2 \tag{3}$$

where $C(W) = \frac{1}{\sqrt{2m}} \min_{x \in \mathbb{R}^n} \sigma(W|_{S(x,W)})$, and σ denotes the smallest singular value.

This result immediately yields stability estimates when solving inverse problems by inverting (injective) generative networks. We note that the $(2m)^{-1/2}$ factor is essential for our estimate; the naive “smallest singular value” estimate is too optimistic.

Remark 9 *We note here that that Theorem 8 gives a bound on the stability of leaky ReLU layers. Layers of the form $\text{LReLU}_\alpha(x) = \text{ReLU}(Wx) - \alpha \text{ReLU}(-Wx)$ are always injective, provided that W is full rank, but may be unstable if $\alpha \approx 0$. To be explicit, the concept of a DSS still applies to these activation functions, and if a leaky relu layer satisfies the criterion of Theorem 2, then it automatically inherits the stability estimate of Theorem 8 which is independent of α .*

2.4 Convolutional Layer

Since convolution is a linear operator we could simply apply Theorem 2 and Lemma 3 to the convolution matrix. It turns out, however, that there exists a result specialized to convolution that is much simpler to verify. In this section we make heavy use of multi-index notation, indicated by capital letters, see Section 1.4.

Definition 10 (Convolution Operator) *Let $c \in \mathbb{R}^O$. We say that c is a convolution kernel of width O . Given c and $x \in \mathbb{R}^N$, we define the convolution operator $C \in \mathbb{R}^{N \times N}$ with stride 1 as,*

$$(Cx)_J = \sum_{I=1}^O c_{O-I+1} x_{J+I} = \sum_{I'=1+J}^{O+J} c_{O+J-I'+1} x_{I'}$$

When $1 \not\leq K$ or $K \not\leq O$ we will set $c_K = 0$. We do not specify the boundary condition on x (zero-padded, periodic, or otherwise), as our results hold generally.

Definition 11 (Convolutional Layer) *We say that a matrix $W \in \mathbb{R}^{M \times N}$ is a convolution operator, with \mathbf{n}_c convolutions, provided that W can be written (up to row rearrangement) in the form $W = [C_1^T \ C_2^T \ \dots \ C_{\mathbf{n}_c}^T]^T$ where each C_k is a convolution operator for $k = 1, \dots, \mathbf{n}_c$. A neural network layer for which W is a convolution operator is called a convolution layer.*




Figure 4: A visualization of the indices in Definition 11 in two dimensions. The blue region is a kernel c of width O , the pink region is the zero-padded box of width P , and D is the offset of the kernel c in the pink box. The entire signal is the green box of width N .

Definitions 10 and 11 automatically model the standard multi-channel convolution used in practice. For 2D images of size 512×512 , n_c input channels, and 3×3 convolutions, we simply let $x \in \mathbb{R}^N$ with $N = (512, 512, n_c)$ and $c \in \mathbb{R}^O$ with $O = (3, 3, n_c)$.

To state our main result for convolutions we also need to define the set of zero-padded kernels for $c \in \mathbb{R}^O$. To aid understanding we first give an example of a 2-dimensional zero-padded kernel, as visualized by Figure 4. The task is to zero-pad a kernel $c = \begin{bmatrix} 1 & -1 \\ -2 & 2 \end{bmatrix}$ (the blue box) of dimensions $O = (O_1, O_2) = (2, 2)$ so that it occupies the region (the pink box) of size $P = (P_1, P_2) = (3, 3)$. This can be done in four different ways, which we call $\mathcal{P}_{(3,3)}(c)$, given by

$$\mathcal{P}_{(3,3)} \left(\begin{bmatrix} 1 & -1 \\ -2 & 2 \end{bmatrix} \right) = \left\{ \begin{bmatrix} 1 & -1 & 0 \\ -2 & 2 & 0 \\ 0 & 0 & 0 \end{bmatrix}, \begin{bmatrix} 0 & 0 & 0 \\ 1 & -1 & 0 \\ -2 & 2 & 0 \end{bmatrix}, \begin{bmatrix} 0 & 1 & -1 \\ 0 & -2 & 2 \\ 0 & 0 & 0 \end{bmatrix}, \begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & -1 \\ 0 & -2 & 2 \end{bmatrix} \right\}. \quad (4)$$

Further, there is a natural way to order the elements of $\mathcal{P}_P(c)$ for any c and P . We can let $D = (D_1, D_2)$ denote the offset of c from the top-left corner. In this way, we denote the (respective) matrices on the r.h.s. of Equation 4 by $\hat{c}^{(0,0)}$, $\hat{c}^{(1,0)}$, $\hat{c}^{(0,1)}$, $\hat{c}^{(1,1)}$. We use the hat notation to remind ourselves that \hat{c}^D is defined on a different domain than c . This example can be straight-forwardly extended to any dimension and kernels of any size as given below.

Definition 12 (Zero-Padded Kernel) *Let c be a convolution kernel of width O , and let P be another multi-index. Given O and P , let $\mathcal{D}(O, P)$ be the set of admissible offsets D where $D \in \mathcal{D}(O, P)$ if $0 \leq D \leq P - O$. Each different choice of $D \in \mathcal{D}(O, P)$ corresponds to*

a different way to pad a kernel with zeros. Note that if $O \not\leq P$, then $\mathcal{D}(O, P) = \emptyset$. Likewise, let $\mathcal{T}(P)$ be the set of admissible indices T where $T \in \mathcal{T}(P)$ if $0 \leq T \leq P$. We define the set of zero-padded kernels¹ of c as the set

$$\mathcal{P}_P(c) = \left\{ \hat{c}^D \in \mathbb{R}^P : \text{for each } T, (\hat{c}^D)_T = \begin{cases} c_{D-T+O+1} & \text{if } 1 \leq T - D \leq O \\ 0 & \text{otherwise} \end{cases} \right\}. \quad (5)$$

where the $D \in \mathcal{D}(O, P)$ and $T \in \mathcal{T}(P)$ in the r.h.s. of Equation 5.

Theorem 13 (Sufficient Condition for Injectivity of Convolutional Layer) *Suppose that $W \in \mathbb{R}^{M \times N}$ is a convolution layer with convolutions $\{C_k\}_{k=1}^{n_c}$, and corresponding kernels $\{c_k\}_{k=1}^{n_c}$. If for any P ,*

$$W|_{\mathcal{P}_P} := \bigcup_{k=1}^{n_c} \mathcal{P}_P(c_k)$$

has a DSS for \mathbb{R}^P with respect to all $x \in \mathbb{R}^P$, then $\text{ReLU}(W \cdot)$ satisfies Theorem 2.

We note that Theorem 13 applies if $\bigcup_{k=1}^{n_c} \mathcal{P}_P(c_k)$ is a DSS for \mathbb{R}^P for *any* P . The proof of the theorem shows that if it holds for any P then the entire operator taken on the whole domain must be injective. See Example 1 in Appendix A.4 for an example of applying Theorem 13 in one-dimension.

Theorem 13 applies to multi-channel convolution of width $(O, n_c) = (O_1, \dots, O_p, n_c)$ provided that we choose a (P, n_c) such that $O \leq P$. This theorem shows that if a convolution operator W has a DSS of the pink region, w.r.t. vectors with support in the pink region in Figure 4, then it has a DSS of the entire green region w.r.t. vectors supported on the green region as well.

To apply Theorem 13 to a convolution layer with n_c kernels of width O , we must choose P so that $W|_{\mathcal{P}_P}$ has at least the minimal number $2 \prod_{j=1}^p P_j$ of vectors to have a DSS of a vector space of dimension $|P| = \prod_{j=1}^p P_j$. Some algebra gives that $n_c \geq 2 \prod_{j=1}^p \frac{1}{1 - O_j/P_j}$ and $P_j \geq \frac{O_j}{1 - (2/n_c)^{1/p}}$, where the last inequality holds only when $\frac{P_j}{O_j}$ is independent of j .

In this section we also present a lemma which is key in the proof of Theorem 13. We believe that this lemma in particular will be useful for future work, as it simplifies the process of showing that a matrix is a DSS of a large domain, by showing it is DSS of smaller ones.

Lemma 14 (Domain Decomposition) *Suppose that $\mathbb{R}^n = \text{span}\{\Omega_1, \dots, \Omega_K\}$ where² each Ω_k is a subspace and for each $k = 1, \dots, K$ we have a*

$$W_k = [w_{k,1}^T, \dots, w_{k,N_k}^T]^T \quad \text{and} \quad W = [W_1^T, \dots, W_K^T]$$

such that $w_{k,\ell} \in \Omega_k$ and W_k has a DSS of Ω_k w.r.t. every $x \in \Omega_k$. Then W has a DSS of \mathbb{R}^n w.r.t. every $x \in \mathbb{R}^n$.

-
1. With many convolutional neural networks, padding refers to the act of padding the image (with for example zeros), but the convolutional kernels are not padded. For our results, the variable P refers to the padding of the kernels, not the image.
 2. This is a slight abuse of traditional spanning notation. Here we mean that there is a set of vectors of $\Omega_1, \Omega_2, \dots$ such that their union spans \mathbb{R}^n .

2.5 Normalization During Runtime

Normalization strategies such as batch (Ioffe and Szegedy, 2015), layer (Ba et al., 2016), instance (Ulyanov et al., 2016), group (Wu and He, 2018), weight (Salimans and Kingma, 2016) and spectral (Miyato et al., 2018) promote convergence during training and encourage low generalization error. Normalization is a many-to-one operation. In this section we show that batch, weight, and spectral normalization do not interfere with injectivity provided that the network is injective without normalization. In the interest of making this work self-contained, we will not concern ourselves here with the detailed questions of how a network is trained, thus we are only concerned with compatibility between injectivity of a *trained* network. The question of how a network might be trained in such a way as to be injective after training is an important, but beyond the scope of this work.

Let $\{x_i\}_{i=1,\dots,m}$ represent the inputs to a given layer over a mini-batch. The batch normalization adds two learnable parameters (γ, β) and transforms x_i to y_i as

$$\mu_{\mathcal{B}} = \frac{1}{m} \sum_{i=1}^m x_i, \quad \sigma_{\mathcal{B}}^2 = \frac{1}{m} \sum_{i=1}^m (x_i - \mu_{\mathcal{B}})^2, \quad \hat{x}_i = \frac{x_i - \mu_{\mathcal{B}}}{\sqrt{\sigma_{\mathcal{B}}^2 + \epsilon}}, \quad y_i = \gamma \hat{x}_i + \beta.$$

Although for a given γ, β the relationship between x_i and y_i is not injective, batch normalization is usually present during training but not at test time. Provided that the learned weights satisfy Theorem 2 and Lemma 3, batch normalization does not spoil injectivity.

In batch renormalization (Ioffe, 2017) it is still desirable to whiten the input into each layer. During run time there may be no mini-batch, so running averages of the σ_i and μ_i (denoted $\hat{\sigma}, \hat{\mu}$) computed during training are used. Batch renormalization at test time is then $\hat{x} = \frac{x - \hat{\mu}}{\sqrt{\hat{\sigma}^2 + \epsilon}}$, $y = \gamma \hat{x} + \beta$. Since, importantly, $\hat{\sigma}$ and $\hat{\mu}$ are not functions of x , the mapping between x and y is one-to-one (when $\gamma \neq 0$), thus such normalization in an injective network does not spoil the injectivity.

In weight normalization (Salimans and Kingma, 2016) the coefficients of the weight matrices are normalized to have a given magnitude. This normalization is not a function of the input or output signals, but rather of the weight matrices themselves. This plays no role in injectivity.

Layer, instance and group normalization all take place during both training and execution and, unlike weight and spectral normalization, the normalization is done on the inputs/outputs of layers instead of the weight matrix, thus injectivity is lost. In Appendix B, we present a more general notion of injectivity (Scalar-Augmented Injective Normalization) in order to preserve some of the properties of injectivity in networks that use these normalization strategies.

3. Universality and Expansivity of Deep Injective Networks

We now consider general properties of deep injective networks. Note that a neural network $N_{\theta} \in \mathcal{NN}(n, m)$ is Lipschitz smooth: there is $L_0 > 0$ such that $|N_{\theta}(x) - N_{\theta}(y)| \leq L_0|x - y|$ for all $x, y \in \mathbb{R}^n$. If $N_{\theta} : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is also injective, results of basic topology imply that for any bounded and closed set $B \subset \mathbb{R}^n$ the map $N_{\theta} : B \rightarrow N_{\theta}(B)$ has a continuous inverse $N_{\theta}^{-1} : N_{\theta}(B) \rightarrow B$ and the sets B and its image $N_{\theta}(B)$ are homeomorphic. Thus, for example, if Z is a random variable supported on the cube $[0, 1]^n \subset \mathbb{R}^n$, the sphere $\mathbb{S}^{n-1} \subset \mathbb{R}^n$,

or the torus $\mathbb{T}^{n-1} \subset \mathbb{R}^n$, we see that $N_\theta(Z)$ is a random variable supported on a set in \mathbb{R}^m that is homeomorphic to an n -dimensional cube or an $n - 1$ dimensional sphere or a torus, respectively. This means that injective neural networks can model random distributions on surfaces with prescribed topology. Moreover, if $N_\theta : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is a decoder, all objects $x \in N_\theta(\mathbb{R}^n)$ correspond to the unique code z such that $x = N_\theta(z)$.

A fundamental property of neural networks is that they can uniformly approximate any continuous function $f : \mathcal{C} \rightarrow \mathbb{R}^m$ defined in a bounded set $\mathcal{C} \subset \mathbb{R}^n$. For general dimensions n and m , the injective neural networks do not have an analogous property. For instance, if $f : [-\pi, \pi] \rightarrow \mathbb{R}$ is the trigonometric function $f(x) = \sin(x)$, it is easy to see that there is no injective neural network (or any other injective function) $N_\theta : [-\pi, \pi] \rightarrow \mathbb{R}$ such that $|f(x) - N_\theta(x)| < 1$. Consider, however, the following trick: add two dimensions in the output vector and consider the map $F : [-\pi, \pi] \rightarrow \mathbb{R}^3$ given by $F(x) = (0, 0, \sin(x)) \in \mathbb{R}^3$. When f is approximated by a one-dimensional, non-injective ReLU network $f_\theta : [-\pi, \pi] \rightarrow \mathbb{R}$ and $\alpha > 0$ is small, then the neural network $N_\theta(x) = (\alpha \text{ReLU}(x), \alpha \text{ReLU}(-x), f_\theta(x))$ is an injective map that approximates F .

In general, as ReLU networks are piecewise affine maps, it follows in the case $m = n$ that if a ReLU network $N_\theta : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is injective, it has to be surjective (Thm. 2.1 Scholtes, 1996). This is a limitation of injective neural networks when $m = n$. Building on these ideas we show that when the dimension of the range space is sufficiently large, $m \geq 2n + 1$, injective neural networks become sufficiently expressive to universally model arbitrary continuous maps. In other words, we show that any continuous map $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$, where the dimension of the image space satisfies $m \geq 2n + 1$, can be approximated uniformly in compact sets by injective neural networks. We emphasize that the map f does not need to be injective.

Theorem 15 (Universal Approximation with Injective Neural Networks) *Let $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$ be a continuous function, where $m \geq 2n + 1$, and $L \geq 1$. Then for any $\varepsilon > 0$ and compact subset $\mathcal{C} \subset \mathbb{R}^n$ there exists a neural network $N_\theta \in \mathcal{NN}(n, m)$ of depth L such that $N_\theta : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is injective and*

$$\|f(x) - N_\theta(x)\|_{2,m} \leq \varepsilon_1, \quad \text{for all } x \in \mathcal{C}. \quad (6)$$

Further we have the following. Let N_θ be such that there exists open and disjoint $\Omega_j \subset \mathbb{R}^n$ for $j = 1, \dots, J$ where $\mathbb{R}^n = \bigcup_{j=1}^J \bar{\Omega}_j$ and $N_\theta|_{\Omega_j}$ is affine for each j . Then there exist $C_1, C_2, \varepsilon_2(n, m) > 0$, such that for any $\varepsilon_2 \in \varepsilon_2(n, m)$ and $R > 0$ where $B(R) \subset \mathcal{C}$ there exists a $N'_{\theta'} \in \mathcal{NN}(n, m)$ where

$$\|f(x) - N'_{\theta'}(x)\|_{2,m} \leq \varepsilon_1 + \varepsilon_2, \quad \text{for all } x \in B(R)$$

and for any $x, y \in B(R)$,

$$\|N'_{\theta'}(x) - N'_{\theta'}(y)\|_{2,m} \geq C \|x - y\|_{2,n}$$

where

$$C = \frac{C_1^{(n+m)n}}{(2n + \sqrt{m+n})^n (\|N_\theta\|_{C(B(R))} + R)^n J^{2 \ln(1+n/(m-2n)) + C_2}} \frac{\varepsilon_2^n}{J^{2 \ln(1+n/(m-2n)) + C_2}}.$$

Before describing the proof, we note that this result also holds for leaky ReLU networks with no modification. The key fact used in the proof is that the network is a piecewise affine function.

To prove this result, we combine the approximation results for neural networks (for example Pinkus’s density result for shallow networks (Theorem 3.1 Pinkus, 1999) or Yarotsky’s result for deep neural networks Yarotsky, 2017), with the Lipschitz-smooth version of the generic projector technique. This technique from differential topology is used for example to prove the easy version of the Whitney’s embedding theorem (Chapter 2, Thm. 3.5 Hirsch, 2012) and Whitney’s injective immersion theorem (Thm. 2.4.3 Mukherjee, 2015). Related random projection techniques have been used earlier in machine learning and compressed sensing (Broomhead and Kirby, 2000, 2001; Baraniuk and Wakin, 2009; Hegde et al., 2008; Iwen and Maggioni, 2013).

The proof of Theorem 15 is based on applying the above classical results to locally approximate the function $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$ by some (possibly non-injective) ReLU-based neural network $F_\theta : \mathbb{R}^n \rightarrow \mathbb{R}^m$, and augment it by adding additional variables, so that $H_\theta(x) = (x, F_\theta(x))$ is an injective map $H_\theta : \mathbb{R}^n \rightarrow \mathbb{R}^{n+m}$. The image of this map, $M = H_\theta(\mathbb{R}^n)$, is an n -dimensional, Lipschitz-smooth submanifold of \mathbb{R}^{n+m} . The dimension of $m + n$ is larger than $2n + 1$, which implies that for a randomly chosen projector P_1 that maps \mathbb{R}^{n+m} to a subspace of dimension $m + n - 1$, the restriction of P_1 on the submanifold M , that is, $P_1 : M \rightarrow P_1(M)$, is injective. By applying n random projectors, P_1, \dots, P_n , we have that $N_\theta = P_n \circ P_{n-1} \circ \dots \circ P_1 \circ H_\theta$ is an injective map whose image is in a m dimensional linear space. These projectors can be multiplied together, but are generated sequentially. By choosing the projectors P_j in a suitable way, the obtained neural network N_θ approximates the map f .

We point out that in the proof of Theorem 15 it is crucial that we first approximate function f by a neural network and then apply to it a generic projection to make the neural network injective. Indeed, doing this in the opposite order may fail as an arbitrarily small deformation (in $C(\mathbb{R}^n)$) of an injective map may produce a map that is non-injective.

Note that J is estimated in theory of hyperplane arrangements (Section 3.11 Stanley, 2011). Indeed, when m hyperplanes $\mathcal{A} = \{T_1, \dots, T_m\}$ divide \mathbb{R}^n to disjoint regions, the number of the regions $J = r(\mathcal{A})$ is bounded by

$$J \leq 1 + m + \binom{m}{2} + \dots + \binom{m}{n}.$$

The proof of Theorem 15 implies the following corollary on cascaded neural networks where the dimensions of the hidden layers can both increase and decrease (see Figure 5):

Corollary 16 *Let $n, m, d_j \in \mathbb{Z}_+$, $j = 0, 1, \dots, 2k$ be such that $d_0 = n$, $d_{2k} = m \geq 2n + 1$ and $d_j \geq 2n + 1$ for even indexes $j \geq 2$. Let*

$$F_k = B_k \circ f_\theta^{(k)} \circ B_{k-1} \circ f_\theta^{(k-1)} \circ \dots \circ B_1 \circ f_\theta^{(1)}$$

where $f_\theta^{(j)} : \mathbb{R}^{d_{2j-2}} \rightarrow \mathbb{R}^{d_{2j-1}}$ are injective neural networks and $B_j \in \mathbb{R}^{d_{2j-1}} \rightarrow \mathbb{R}^{d_{2j}}$ are random matrices whose joint distribution is absolutely continuous with respect to the Lebesgue measure of $\prod_{j=1}^k (\mathbb{R}^{d_{2j} \times d_{2j-1}})$. Then the neural network $F_k : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is injective almost surely.




Figure 5: An illustration of an injective deep neural network that avoids expansivity as described by Corollary 16. White trapezoids are expansive weight matrices satisfying Theorem 2, and the blue trapezoids are random projectors that reduce dimension while preserving injectivity.

Observe that in Corollary 16 the weight matrices B_j and $B_{j'}$ may not be independent. Moreover, the distribution of the matrices B_j may be supported in an arbitrarily small set, that is, a small random perturbation of the weight matrices make the function F_k injective. Corollary 16 characterizes deep globally injective networks, while avoiding the exponential growth in dimension as indicated by Corollary 6.

4. Numerical Experiments

In this section we present numerical results. First we numerically investigate the results of Theorem 7, and then show that injective GANs improve inference.

4.1 Experiments Testing the Minimal Expansivity Threshold

Theorem 7 implies that an expansivity factor of 2.1 as suggested by Lei et al. (2019) is not enough to ensure global injectivity with Gaussian weights. Insufficient expansivity leads to failure cases in inverse problems, as we show in this section. We use a single ReLU layer, $f(x) = \text{ReLU}(Wx)$, and choose x to be an MNIST digit for ease of visualization of the “latent” space. We choose $W \in \mathbb{R}^{cn \times n}$ with iid Gaussian entries and set $c = m/n = 2.1$. Given $y = f(x)$, our goal is to reconstruct x using, for example, gradient descent. This requires at least n of the cn entries of $f(x)$ to be non-zero since $x \in \mathbb{R}^n$. As shown in Figure 7, $x_0 = -\frac{1}{m} \sum_j w_j / \|w_j\|$ violates this condition with high probability when $c = 2.1$.

We choose 2 data samples, x_1 and x_2 to be 2 “latent” codes. Note that W has a DSS with respect to x_1 . This ensures that given $f(x_1)$ one can reconstruct x_1 . Now, in order to show that injectivity around a few points is insufficient for global injectivity, we linearly interpolate between x_0 and x_1 and x_0 and x_2 in a 2D grid. We show the intermediate x s as images in Figure 6. The x s shown in red cannot be uniquely recovered given $f(x)$, that is, there exist infinitely many samples close to the red samples that all map to the same output.

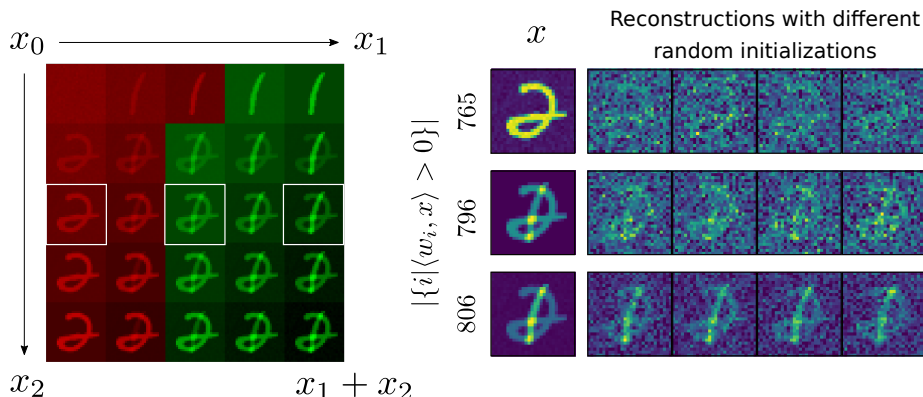


Figure 6: The mapping $\phi(x) := \text{ReLU}(Wx)$, $x \in \mathbb{R}^n$ when $W \in \mathbb{R}^{2.1n \times n}$ as suggested in Lei et al. (2019) is not injective as the samples shown in red do not have n positive inner products with the rows of W . This implies that around the red samples ϕ cannot be inverted. For the 3 samples delineated by white boxes on the left, we do a simple denoising test where $y = \phi(x) + \eta$ and η corresponds to 10dB noise. Each reconstruction reported is best out of 10 trials based on MSE error.

To further illustrate this, we also run a simple experiment where we have $y = f(x) + \eta$, with a 10dB signal-to-noise ratio. We invert f to estimate x for 3 different samples in the domain of f . We can easily see that non-injective points of the domain give poorer reconstructions due to the lower number of positive inner products with W . In order to avoid problems of non-convexity in optimization we report the best reconstruction out of 10 in 4 different trials with different random seeds, similar to experiments of Bora et al. (2017). While these results show reconstructions with single layer, the issues are only exacerbated with multiple layers.

4.2 Injective GANs Improve Inference

We devise a simple experiment to demonstrate that 1) one can construct injective networks using Corollary 6 that 2) perform as well as the corresponding non-injective networks, while 3) being better at inference problems. One way to do inference with GANs is to train a so-called inference network $I_{\text{net}} : \mathcal{X} \rightarrow \mathcal{Z}$ jointly with the generator $G : \mathcal{Z} \rightarrow \mathcal{X}$ so that $I(G(z)) \approx z$ (Donahue et al., 2016; Dumoulin et al., 2016); I_{net} is trained to invert G on its range. This is used to evaluate the likelihood of an image x in the range of G as $p_{\mathcal{Z}}(I_{\text{net}}(x))$. However, for $p_{\mathcal{Z}}(I_{\text{net}}(G(z))) = p_{\mathcal{Z}}(z)$ to hold, G must be injective—otherwise the distribution of $I_{\text{net}}(G(z))$ will be different from that of z . One would thus hope that the distribution of $I_{\text{net}}(G(z))$ will be closer to that of z if we enforce injectivity. If that is the case and the sample quality is comparable, then we have a practical tool to improve inference.

We use the DCGAN (Radford et al., 2015) architecture with the same hyperparameters, layer sizes and normalization strategies for the regular GAN and the injective GAN; see Appendix F. Note that the original DCGAN uses several ReLU conv layers with an expansivity of 2 per layer. By Corollary 6 these layers are non-injective. Therefore, the filters of the DCGAN are modified to be of the form $[C; -s^2C]$ in order to build an injective GAN. We

train the networks to generate $64 \times 64 \times 3$ images and draw the latent code $z \in \mathbb{R}^{256}$ iid from a standard normal distribution. We test on CelebA (Liu et al., 2015) and FFHQ (Karras et al., 2019) data sets. To get a performance metric, we fit Gaussian distributions $\mathcal{N}(\mu, \Sigma)$ and $\mathcal{N}(\mu_{\text{inj}}, \Sigma_{\text{inj}})$ to $G(z)$ and $G_{\text{inj}}(z)$. We then compute the Wasserstein-2 distance \mathcal{W}_2 between the distribution of $z \sim \mathcal{N}(0, I_{256})$ and the two fitted Gaussians using the closed-form expression for \mathcal{W}_2 between Gaussians, $\mathcal{W}_2^2(\mathcal{N}(\mu_1, \Sigma_1), \mathcal{N}(\mu_2, \Sigma_2)) = \|\mu_1 - \mu_2\|^2 + \text{Tr}(\Sigma_1 + \Sigma_2 - 2(\Sigma_1 \Sigma_2)^{1/2})$. We summarize the results in Table 1. Despite the restrictions on the weights of injective generators, their performance on popular GAN metrics—Fréchet inception distance (FID) (Heusel et al., 2017) and inception score (IS) (Salimans et al., 2016)—is comparable to the standard GAN while inference improves. That the generated samples are indeed comparable to the standard GAN can be also gleaned from the figure in Appendix F.

Table 1: Injectivity improves inference without sacrificing performance.

data set	Type of G	Inception Score \uparrow	Fréchet Inception Distance \downarrow	$\mathcal{W}_2^2(\mathbb{P}_z, \mathbb{P}_z) \downarrow$
<i>CelebA</i>	Injective	2.24 ± 0.09	39.33 ± 0.41	18.59
	Regular	2.22 ± 0.16	50.56 ± 0.52	33.85
<i>FFHQ</i>	Injective	2.56 ± 0.15	61.22 ± 0.51	9.87
	Regular	2.57 ± 0.16	47.23 ± 0.90	19.63

5. Conclusion

We derived and explored conditions for injectivity of ReLU neural networks. In contrast to prior work which looks at random weight matrices, our characterizations are deterministic and derived from first principles. They are also sharp in that they give sufficient and necessary conditions for layerwise injectivity. Our results apply to any network of the form (1)—they only involve weight matrices but make no assumptions about the architecture. We included explicit constructions for minimally expansive networks that are injective; interestingly, this simple criterion already improves inference in our preliminary experiments. The results on universality of injective neural networks further justify their use in applications; they also implicitly justify the various Jacobian conditioning strategies when learning to generate real-world data. Further, injective neural networks are topology-preserving homeomorphisms which opens applications in computational topology and establishes connections to tools such as self-organizing maps. Analysis of deep neural networks has an analogue in the analysis of inverse problems where one studies uniqueness, stability and reconstruction. Uniqueness coincides with injectivity, quantitative stability with the Lipschitz constant of the inverse, and, following Lei et al. (2019), from a linear program we get a reconstruction in the range.

We also remark here that we believe that many of the theoretical results from this paper can be applied to networks with other activation functions with minimal changes. For example we believe that, Theorems 2, 8, 13 and 20 can be made to apply to arbitrary activation functions that are injective on a half plane. We also believe that Theorem 15 can be applied to any model of injective networks provided that the activation functions are piece-wise affine.

Acknowledgments

We are grateful to Daniel Paluka and Charles Clum for pointing out a problem with the bound on the sum of binomial coefficients used to prove Theorem 7. We are also grateful to Teemu Saksala for the many useful conversations had, especially in the early stages of this project.

I.D. was supported by the European Research Council Starting Grant 852821—SWING. M.L. was supported by Academy of Finland, grants 284715, 312110. M.V.dH. gratefully acknowledges support from the Department of Energy under grant DE-SC0020345, the Simons Foundation under the MATH + X program, and the corporate members of the Geo-Mathematical Imaging Group at Rice University.

Appendix A. Proofs from Section 2

In this section we present all of the proof concerning layerwise injectivity.

A.1 Proofs from Subsection 2.2

Proof [Proof of Theorem 2] Suppose that W is such that the conditions of Theorem 2 hold, and that $\text{ReLU}(Wx_1) = \text{ReLU}(Wx_2) = y$. If for $j \in [[m]]$, $y|_j > 0$ then both $\langle w_j, x_1 \rangle > 0$ and $\langle w_j, x_2 \rangle > 0$. Similarly, if $y|_j \leq 0$ then both $\langle w_j, x_1 \rangle \leq 0$ and $\langle w_j, x_2 \rangle \leq 0$. In particular, this implies that

$$\langle w_j, x_1 \rangle > 0 \iff \langle w_j, x_2 \rangle > 0 \text{ and } \langle w_j, x_1 \rangle \leq 0 \iff \langle w_j, x_2 \rangle \leq 0. \quad (7)$$

If we then consider $x_\alpha := (1 - \alpha)x_1 + \alpha x_2$ where $\alpha \in (0, 1)$, then

$$\text{ReLU}(Wx_\alpha) = y = \text{ReLU}(Wx_1) = \text{ReLU}(Wx_2).$$

If $\langle w_j, x_\alpha \rangle > 0$ then at least one of $\langle w_j, x_1 \rangle > 0$ or $\langle w_j, x_2 \rangle > 0$. (7) implies that both must hold, therefore $\langle w_j, x_1 \rangle = \langle w_j, x_2 \rangle > 0$. If $\langle w_j, x_\alpha \rangle = 0$ then $\langle w_j, x_1 \rangle = \langle w_j, x_2 \rangle = 0$ (otherwise (7) is violated), thus

$$\text{ReLU}(W|_{S(x_\alpha, W)}x_1) = \text{ReLU}(W|_{S(x_\alpha, W)}x_2) \implies W|_{S(x_\alpha, W)}x_1 = W|_{S(x_\alpha, W)}x_2$$

and so because $W|_{S(x_\alpha, W)}$ is full rank, this implies that $x_1 = x_2$. This proves one direction.

The other direction follows from the following. Suppose that there exists a x such that $W|_{S(x, W)}$ don't span \mathbb{R}^n . If $S(x, W) = \emptyset$ non-injectivity trivially follows, so suppose w.l.o.g. that $S(x, W) \neq \emptyset$. Let $x^\perp \in \ker(W|_{S(x, W)})$ and $\alpha \in \mathbb{R}^+$ such that $\alpha < \min_{j \in S^c(x, W)} \frac{-\langle x, w_j \rangle}{|\langle x^\perp, w_j \rangle|} 3$. Then for $j = 1, \dots, m$ one of the following two hold

$$\begin{aligned} \text{if } j \in S(x, W) \text{ then } \langle w_j, x + \alpha x^\perp \rangle &= \langle w_j, x \rangle + \alpha \langle w_j, x^\perp \rangle = \langle w_j, x \rangle \\ \text{if } j \in S^c(x, W) \text{ then } \langle w_j, x + \alpha x^\perp \rangle &= \langle w_j, x \rangle + \alpha \langle w_j, x^\perp \rangle < 0. \end{aligned}$$

3. If $\langle x^\perp, w_j \rangle = 0$ for all $j \in S^c(x, W)$, then any $\alpha > 0$ will do.

Thus, as ReLU acts pointwise (row-wise in W), we have that

$$\text{ReLU}(W(x + \alpha x^\perp)) = \text{ReLU}(Wx) \tag{8}$$

and, hence, $\text{ReLU}(W\cdot)$ is not injective. ■

Proof [Proof of Lemma 3] First, we show that if $\text{ReLU}(W|_{b \geq 0} \cdot)$ is injective, then so is $\text{ReLU}(W \cdot + b)$. Clearly if $\text{ReLU}(Wx_1 + b) = \text{ReLU}(Wx_2 + b)$ then $\text{ReLU}(W|_{b \geq 0}x_1 + b|_{b \geq 0}) = \text{ReLU}(W|_{b \geq 0}x_2 + b|_{b \geq 0})$ as well. If we apply Lemma 32 to each component of the above equation, then we obtain that $\text{ReLU}(W|_{b \geq 0}x_1) = \text{ReLU}(W|_{b \geq 0}x_2)$ which, given the injectivity of $\text{ReLU}(W|_{b \geq 0} \cdot)$, implies that $x_1 = x_2$.

Now suppose that $\text{ReLU}(W|_{b \geq 0} \cdot)$ is not injective. Let $x \in \mathbb{R}^n$ be such that $W|_{b \geq 0}$ doesn't have a DSS of \mathbb{R}^n w.r.t. x . Let $\beta > 0$ be small enough so that $W|_{b < 0}(\beta x) + b|_{b < 0} < 0$ component-wise, and let $x^\perp \in \mathbb{R}^n$ such that (as in (8)) $\text{ReLU}(W|_{b \geq 0}(\beta x + x^\perp)) = \text{ReLU}(W|_{b \geq 0}\beta x)$. Further let $\alpha < 1$ be small enough such that $W|_{b < 0}(\beta x + \alpha x^\perp) + b|_{b < 0} < 0$. By a component-wise analysis, we have that

$$\text{ReLU}(W(\beta x + \alpha x^\perp) + b) = \text{ReLU}(W\beta x + b)|_{b \geq 0} = \text{ReLU}(W\beta x + b);$$

thus, $\text{ReLU}(W \cdot + b)$ is not injective. ■

Proof [Proof of Corollary 6] If $W \in \mathbb{R}^{m \times n}$ is injective then consider a plane p in \mathbb{R}^n that none of the rows of W lie in. Apply Theorem 2 to both normals of the plane. The corresponding DSS' for each normal are disjoint, thus there must be at least $2n \geq m$, so $m < 2 \cdot n$ implies non-injectivity.

Now we show that if W satisfies Theorem 2, then W is of the form given by (2). Suppose that there is a row vector w_i such that there are no row vectors pointing in the $-w_i$ direction. Let p be a plane through the origin such that $w_i \in p$, but $w_{i'} \notin p$ for $i \neq i'$. By Theorem 2, there must be at least n columns that lie on each (closed) halfspace of p . Not counting w_i , one side must have $n - 1$ vectors, and the other n . Indeed one of the sides must have exactly n vectors on it (including w_i). Let θ_i be the principle angle between w_i and p . If we rotate p by an angle $\phi < \min_{i \neq i'} \theta_i$ then we can obtain a new plane p_ϕ with w_i on the side with n vectors, so that the other side of p_ϕ still only contains $n - 1$ vectors. Hence there is no DSS for p_ϕ plane's normal on that side.

Thus for $\text{ReLU } W$ to be injective, for every $i \in [[2n]]$ there must be a different $i' \in [[2n]]$ such that w_i and $w_{i'}$ are anti-parallel. This can only happen if for every $w \in W$ there is a corresponding $-dw \in W$, so W must have the form of (2). ■

A.2 Proof of Theorem 7

A.2.1 UPPER BOUND ON MINIMAL EXPANSIVITY

Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space and let $w_i : \Omega \rightarrow \mathbb{R}^n$, $1 \leq i \leq m$ be m iid Gaussian vectors on Ω stacked in a matrix $W : \Omega \rightarrow \mathbb{R}^{m \times n}$.

We identify a sufficient condition on m, n for a ReLU layer $x \mapsto \text{ReLU}(Wx)$ to be injective. Injectivity fails at p when the half-space $\{x : \langle x, p \rangle \geq 0\}$ contains fewer than n vectors w_i .

Equivalently, injectivity fails if there is a half-space which contains more than $m - n$ vectors w_i (which implies that its opposite half-space has fewer than n).

Let $k = m - n + 1$ and denote by $A \in \mathbb{R}^{k \times n}$ some $k \times n$ submatrix of W (for example, the first k rows). Our strategy is to first bound the probability that for a fixed subset of k rows of W , there exists an x having positive inner products with all k rows (which signals non-injectivity per above discussion) and then use the union bound over the $\binom{m}{k}$ subsets of k rows to upper bound the probability of non-injectivity.

For the first part we follow the proof of Bürgisser and Cucker (2013, Theorem 13.6), parts of which we reproduce for the reader's convenience. For a sign pattern $\sigma \in \{-1, 0, 1\}^k$, we denote by $R_A(\sigma)$ the set of all $x \in \mathbb{R}^n$ which produce the sign pattern σ (they belong to the σ -“wedge”, possibly empty),

$$R_A(\sigma) = \{x \in \mathbb{R}^n : \text{sign}(\langle x, a_i \rangle) = \sigma_i, i \in \{1, \dots, k\}\}.$$

For $\sigma \in \{-1, +1\}^k = \Sigma$, we define the event \mathcal{E}_σ as

$$\mathcal{E}_\sigma = \{\omega : R_{A(\omega)}(\sigma) \neq \emptyset\}.$$

We are interested in $\sigma_0 = (1, \dots, 1)$, meaning that all the inner products are positive. Note that the probability of \mathcal{E}_σ is the same for all σ due to the symmetry of the Gaussian measure. Further, note that $\sum_{\sigma \in \Sigma} \mathbf{1}_{\mathcal{E}_\sigma}(\omega) = |\{\sigma : R_{A(\omega)}(\sigma) \neq \emptyset\}|$ is the number of wedges defined by A . Then

$$\mathbb{P}(\mathcal{E}_{\sigma_0}) = \frac{1}{2^k} \sum_{\sigma \in \Sigma} \mathbb{P}(\mathcal{E}_\sigma) = \frac{1}{2^k} \sum_{\sigma \in \Sigma} \mathbf{E}(\mathbf{1}_{\mathcal{E}_\sigma}) = \frac{1}{2^k} \mathbf{E} \left(\sum_{\sigma \in \Sigma} \mathbf{1}_{\mathcal{E}_\sigma} \right) = \frac{1}{2^{k-1}} \sum_{i=0}^{n-1} \binom{k-1}{i},$$

by (Winder, 1966, Cor. 1). In Winder (1966) the corollary holds for the generic case but the equality case is generic, and so holds with probability 1, thus we use equality here.

We now have the probability that for a subset of k vectors w_i , there exists an $x \in \mathbb{R}^n$ which has positive inner products with all k vectors. We are interested in the following event which implies non-injectivity,

$$\mathcal{E}_{\text{NI}} = \{\omega : W(\omega) \text{ has a subset of } k \text{ rows } B(\omega) \text{ such that } R_{B(\omega)}(\sigma_0) \neq \emptyset\}.$$

Conversely, $\omega \in \mathcal{E}_{\text{NI}}^c$ implies almost sure injectivity.

Since there are $\binom{m}{k} = \binom{m}{m-n+1} = \binom{m}{n-1}$ different subsets of k rows, we can bound the probability of \mathcal{E}_{NI} as

$$\begin{aligned} \mathbb{P}(\mathcal{E}_{\text{NI}}) &\leq \binom{m}{n-1} \mathbb{P}(\mathcal{E}_{\sigma_0}) = \frac{1}{2^{k-1}} \binom{m}{n-1} \sum_{i=0}^{n-1} \binom{k-1}{i} \\ &\leq \left(\frac{me}{n-1} \right)^k 2^{1-k} 2^{H\left(\frac{n-1}{k-1}\right)(k-1)} \\ &\leq \left(\frac{me}{n-1} \right)^n 2^{-(m-n)} 2^{(m-n)H\left(\frac{n-1}{m-n}\right)} \\ &\leq 4(ce)^n 2^{n(c-1)[H((c-1)^{-1})-1]} \\ &= 4 \cdot 2^{-n[-\log_2(ce) - (c-1)(H((c-1)^{-1})-1)]}, \end{aligned} \tag{9}$$

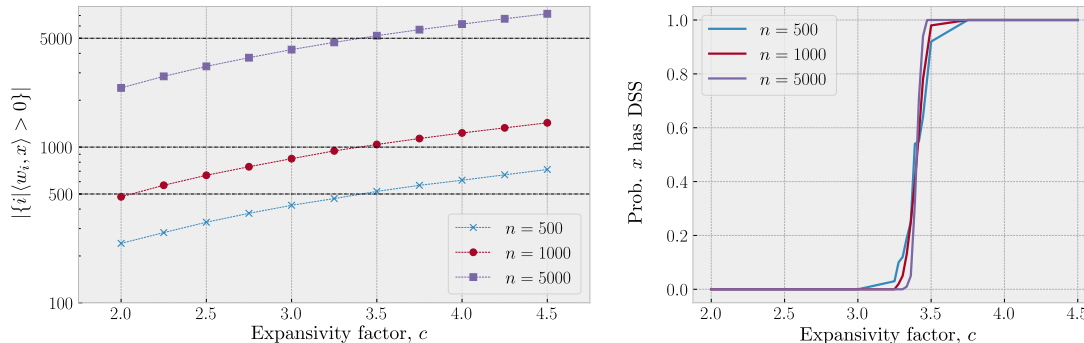


Figure 7: We empirically show that for random Gaussian matrices $\mathbf{W} \in \mathbb{R}^{m \times n}$ $m = cn$, a critical oversampling factor of at least $c > 3.5$ is required for injectivity of ReLU. We choose $x = -\frac{1}{m} \sum_j w_j / \|w_j\|$. *Left*: a plot of the number of elements of W that point in the direction of x as a function of expansivity for several different choices of n . If this quantity is less than n , then W cannot contain a DSS of \mathbb{R}^n w.r.t. x . *Right*: a plot of the empirical probability that a Gaussian matrix W contains a DSS of \mathbb{R}^n w.r.t. x as a function of expansivity for several different values of n .

where⁴ $\binom{n}{k} \leq \left(\frac{ne}{k}\right)^k$, $\sum_{i=1}^k \binom{n}{i} \leq 2^{H(k/n)n}$, and $H(\epsilon) = -\epsilon \log_2(\epsilon) - (1 - \epsilon) \log_2(1 - \epsilon)$ is the binary entropy function. The RHS in (9) converges to zero as $n \rightarrow \infty$ precisely if $[-\log_2(ce) - (c - 1)(H((c - 1)^{-1}) - 1)]$ is positive. Numerically we can see that the bound is positive if $c \geq 10.5$.

A.2.2 LOWER BOUND ON MINIMAL EXPANSIVITY OF LAYERS WITH GAUSSIAN WEIGHTS

We prove that large random Gaussian weight matrices $W \in \mathbb{R}^{m \times n}$ yield non-injective ReLU layers with high probability if $m < c^*n$, where $c^* \approx 3.4$. Injectivity fails if there exists a half-space in \mathbb{R}^n which contains less than n rows of W . Equivalently, injectivity fails if there is a half-space with more than $m - n$ rows of W , since then the opposite halfspace has less than n . The idea is to make an “educated guess” for a half-space with many vectors, and compute the probability that it has more than $m - n$. As the guess, we take the halfspace defined by the average row vector in W . Equivalently, we study the size of $S(\bar{w}, W)$ where $\bar{w} = \frac{1}{m} \sum_{k=1}^m w_k$ is the row average of a matrix W . If $|S(\bar{w}, W)| > m - n$, then $|S(-\bar{w}, W)| < n$, and W cannot have a DSS w.r.t. $-\bar{w}$. As the Figure 7 shows, when $m < c^*n$, W does not contain a DSS w.r.t. $-\bar{w}$ with high probability.

Let $W \in \mathbb{R}^{m \times n}$ be a matrix such that the rows of W are i.i.d. random vectors distributed as $\mathcal{N}(0, I_n)$. Define the event E_i as

$$E_i := \left\{ \omega : \left\langle w_i(\omega), \sum_{k=1}^m w_k(\omega) \right\rangle \geq 0 \right\}.$$

4. https://en.wikipedia.org/wiki/Binomial_coefficient#Sums_of_binomial_coefficients

Lemma 17 *Let $m, n \rightarrow \infty$ so that $\frac{m}{n} = c$. Then $\mathbb{P}(E_i) \rightarrow \frac{1}{2}\operatorname{erfc}\left(-\frac{1}{\sqrt{2c}}\right)$, where $\operatorname{erfc}(z) = 1 - \frac{2}{\sqrt{\pi}} \int_0^z e^{-t^2} dt$.*

Proof We have

$$\mathbb{P}(E_i) = \mathbb{P}\left(\|w_i\|^2 + \left\langle w_i, \sum_{i \neq j} w_j \right\rangle \geq 0\right) = \mathbb{P}\left(\|w_i\|^2 + \|w_i\| Y \geq 0\right)$$

where $Y \sim \mathcal{N}(0, m-1)$. We now use the fact that $\|w_i\|^2$ concentrates around n . Define the event

$$D_i = \left\{ \omega : \|w_i(\omega)\|^2 \in [(1-\epsilon)n, n/(1-\epsilon)] \right\}.$$

It holds by standard concentration arguments (for example Barvinok Page 6, Corollary 2.3 2005) that $\mathbb{P}(D_i) \geq 1 - 2\exp(-\epsilon^2 n/4)$, and so by the law of total probability,

$$\mathbb{P}\left\{\|w_i\|^2 + \|w_i\| Y \geq 0\right\} = \mathbb{P}\{\|w_i\| + Y \geq 0 \mid D_i\} \mathbb{P}\{D_i\} + \mathbb{P}\{\|w_i\| + Y \geq 0 \mid D_i^c\} \mathbb{P}\{D_i^c\}. \quad (10)$$

Choosing $\epsilon = \epsilon(n) = n^{-1/4}$ yields

$$\mathbb{P}\left\{\sqrt{n - n^{3/4}} + Y \geq 0\right\} \leq \mathbb{P}\{\|w_i\| + Y \geq 0 \mid D_i\} \leq \mathbb{P}\left\{\sqrt{n/(1 - n^{-1/4})} + Y \geq 0\right\}.$$

Finally, substituting $Z = \frac{Y}{\sqrt{m-1}} \sim \mathcal{N}(0, 1)$ yields

$$\mathbb{P}\left\{Z \geq -\sqrt{\frac{n - n^{3/4}}{m-1}}\right\} \leq \mathbb{P}\{\|w_i\| + Y \geq 0 \mid D_i\} \leq \mathbb{P}\left\{Z \geq -\sqrt{\frac{n/(1 - n^{-1/4})}{m-1}}\right\}$$

Both sandwiching probabilities converge to $\mathbb{P}\left\{Z \geq -\frac{1}{\sqrt{c}}\right\} = \frac{1}{2}\operatorname{erfc}\left(-\frac{1}{\sqrt{2c}}\right)$. Noting that $\mathbb{P}\{D_i\} \rightarrow 1$ and $\mathbb{P}\{D_i^c\} \rightarrow 0$ we finally have from (10) that

$$\mathbb{P}\{E_i\} \rightarrow \frac{1}{2}\operatorname{erfc}\left(-\frac{1}{\sqrt{2c}}\right).$$

■

Lemma 18 *Under the same conditions as Lemma 17 when $i \neq j$,*

$$\mathbb{P}\{E_i \cap E_j\} \rightarrow \frac{1}{4}\operatorname{erfc}\left(-\frac{1}{\sqrt{2c}}\right)^2.$$

We first prove a helper lemma.

Lemma 19 *Let $X, Y \sim \mathcal{N}(0, I_n)$ be independent Gaussian random vectors. Then*

$$\mathbb{P} \left\{ |\langle X, Y \rangle| \geq \sqrt{n^{3/2}/(1-\epsilon)} \right\} \leq e^{-\sqrt{n}/2} + 2e^{-\epsilon^2 n/4}.$$

Proof The distribution of $\langle X, Y \rangle$ is the same as that of $\|X\| \langle e_1, Y \rangle = \|X\| Z$, with $Z \sim \mathcal{N}(0, 1)$ independently from X and $e_1 = (1, 0, \dots, 0)$. Then

$$\mathbb{P} \left\{ \|X\| Z \geq \sqrt{n^{3/2}/(1-\epsilon)} \right\} \leq \mathbb{P} \left\{ Z \geq n^{1/4} \right\} + \mathbb{P} \left\{ \|X\| \geq \sqrt{n/(1-\epsilon)} \right\}.$$

The second term is bounded by $e^{-\epsilon^2 n/4}$. For the first term we have

$$\mathbb{P} \left\{ Z \geq n^{1/4} \right\} = \frac{1}{2} \operatorname{erfc} \left(n^{1/4}/\sqrt{2} \right) \leq \frac{1}{2} e^{-\sqrt{n}/2}.$$

The claim follows by noting that Z is symmetric around 0. ■

Proof of Lemma 18

$$\begin{aligned} & \mathbb{P} \left\{ \left\langle w_i, \sum_{k=1}^m w_k \right\rangle \geq 0 \text{ and } \left\langle w_j, \sum_{k=1}^m w_k \right\rangle \geq 0 \right\} \\ &= \mathbb{P} \left\{ \|w_i\|^2 + \langle w_i, w_j \rangle + \langle w_i, X \rangle \geq 0 \text{ and } \|w_j\|^2 + \langle w_i, w_j \rangle + \langle w_j, X \rangle \geq 0 \right\} \\ &= (*). \end{aligned}$$

where $X \sim \mathcal{N}(0, m-2)$ is independent from w_i and w_j . Fixing the orientation of X does not change the probability so we choose $X = \|X\| (1, 0, \dots, 0)$ to get

$$(*) = \mathbb{P} \left\{ \|w_i\|^2 + \langle w_i, w_j \rangle + \|X\| w_{i,1} \geq 0 \text{ and } \|w_j\|^2 + \langle w_i, w_j \rangle + \|X\| w_{j,1} \geq 0 \right\}$$

Similarly as in the proof of Lemma 17 define the event

$$D_i = \{ \omega : \|w_i(\omega)\|^2 \in [(1-\epsilon)n, n/(1-\epsilon)] \}$$

and

$$D_X = \left\{ \omega : \|X(\omega)\|^2 \in [(1-\epsilon)n(m-2), n(m-2)/(1-\epsilon)] \right\},$$

which individually hold with probability at least $1 - 2 \exp(-\epsilon^2 n/4)$. Let also

$$D_{\langle i,j \rangle} = \left\{ \omega : |\langle w_i(\omega), w_j(\omega) \rangle| \leq \sqrt{n^{3/2}/(1-\epsilon)} \right\}.$$

We have from Lemma 19 that

$$\mathbb{P} \{ D_{\langle i,j \rangle} \} \geq 1 - \exp(-\sqrt{n}/2) - 2 \exp(-\epsilon^2 n/4).$$

Choosing $\epsilon = \epsilon(n) = n^{-1/4}$ we get via union bound that

$$\mathbb{P}(D_i \cap D_j \cap D_X \cap D_{\langle i,j \rangle}) \geq 1 - C_1 \exp(-C_2 \sqrt{n}),$$

for some $C_1, C_2 > 0$.

Let $D = D_i \cap D_j \cap D_X \cap D_{\langle i,j \rangle}$ and denote

$$\alpha = \mathbb{P} \left\{ \|w_i\|^2 + \langle w_i, w_j \rangle + \|X\| w_{i,1} \geq 0 \quad \text{and} \quad \|w_j\|^2 + \langle w_i, w_j \rangle + \|X\| w_{j,1} \geq 0 \mid D \right\}.$$

Dividing the left-hand sides of inequalities in the argument of \mathbb{P} by m and denoting $\kappa = c^{-1}$, we have for large enough n (so that the term in the parenthesis in the next line is negative),

$$\begin{aligned} \alpha \leq \mathbb{P} \left\{ w_{i,1} \geq \sqrt{\frac{m}{m-2}} \frac{\kappa^{-1/2}}{1-n^{-1/4}} \left(-\frac{\kappa}{1-n^{-1/4}} - \kappa \sqrt{\frac{n^{-1/2}}{1-n^{-1/4}}} \right) \right. \\ \left. \text{and } w_{j,1} \geq \sqrt{\frac{m}{m-2}} \frac{\kappa^{-1/2}}{1-n^{-1/4}} \left(-\frac{\kappa}{1-n^{-1/4}} - \kappa \sqrt{\frac{n^{-1/2}}{1-n^{-1/4}}} \right) \mid D \right\} \quad (11) \end{aligned}$$

(Note that unlike in the proof of Lemma 17, conditioning affects the distribution of w_i, w_j so we need to keep it.) Similarly,

$$\begin{aligned} \alpha \geq \mathbb{P} \left\{ w_{i,1} \geq \sqrt{\frac{m}{m-2}} \kappa^{-1/2} (1-n^{-1/4}) \left(-\kappa(1-n^{-1/4}) + \kappa \sqrt{\frac{n^{-1/2}}{1-n^{-1/4}}} \right) \right. \\ \left. \text{and } w_{j,1} \geq \sqrt{\frac{m}{m-2}} (1-n^{-1/4}) \left(-\kappa(1-n^{-1/4}) + \kappa \sqrt{\frac{n^{-1/2}}{1-n^{-1/4}}} \right) \mid D \right\}. \quad (12) \end{aligned}$$

Denote the conjunctions in the arguments of the probabilities in (11) and (12) by A_u and A_ℓ . We have

$$\begin{aligned} \mathbb{P} \{A_u \mid D\} - \mathbb{P} \{A_u\} &= \frac{\mathbb{P} \{A_u \cap D\}}{\mathbb{P} \{D\}} - \mathbb{P} \{A_u\} \\ &\leq \frac{\mathbb{P} \{A_u\}}{\mathbb{P} \{D\}} - \mathbb{P} \{A_u\} \\ &= \frac{\mathbb{P} \{A_u\}}{\mathbb{P} \{D\}} (1 - \mathbb{P} \{D\}) \rightarrow 0, \end{aligned}$$

since $\mathbb{P} \{D\} \rightarrow 1$, and similarly for A_ℓ , so we can remove the conditioning in the limit. Combining all of the above, we have that

$$\alpha \rightarrow \mathbb{P} \{w_{i,1} \geq -\sqrt{\kappa} \quad \text{and} \quad w_{j,1} \geq -\sqrt{\kappa}\} = \left(\frac{1}{2} \operatorname{erfc} \left(-\sqrt{\frac{\kappa}{2}} \right) \right)^2$$

by independence of $w_{i,1}$ and $w_{j,1}$. ■

Theorem 20 *Given a Gaussian weight matrix $W \in \mathbb{R}^{m \times n}$, the layer $\operatorname{ReLU}(Wx)$ with $m/n = c$ is not injective with probability $\rightarrow 1$ as $n \rightarrow \infty$ when $c < c^*$, where c^* is the unique positive real solution to*

$$\frac{1}{2} \operatorname{erfc} \left(\frac{1}{\sqrt{2c}} \right) = \frac{1}{c}.$$

The numerical value of c^ is ≈ 3.4 .*

Proof Let $X_i = \mathbf{1}_{\langle w_i, \frac{1}{m} \sum_{j=1}^m w_j \rangle \geq 0}$, the indicator function of the event E_i . The expected number of w_i with a positive inner product with $\sum_{i=1}^m w_i$ is by the linearity of expectation equal to

$$\mathbb{E} \left(\sum_{i=1}^m X_i \right) = m \cdot \mathbb{P} \{E_i\} =: mp.$$

By Chebyshev's inequality

$$\mathbb{P} \left(\left| \frac{1}{m} \sum_{i=1}^m X_i - p \right| \geq t \right) \leq \frac{\sigma^2}{t^2}$$

where $\sigma^2 = \mathbb{V} \left(\frac{1}{m} \sum_{i=1}^m X_i \right)$ is the variance of the sum. We compute

$$m^2 \sigma^2 = \mathbb{V} \left(\sum_{i=1}^m X_i \right) = \mathbb{E} \left(\sum_{i=1}^m X_i \right)^2 - \left(\mathbb{E} \sum_{i=1}^m X_i \right)^2 = mp + \sum_{i \neq j} \mathbb{P}(E_i \cap E_j) - (mp)^2.$$

Since $\mathbb{P}(E_i \cap E_j) \rightarrow p^2$ by Lemmas 17 and 18, we have for any i and $j \neq i$

$$\sigma^2 = \frac{mp + m(m-1)\mathbb{P}(E_i \cap E_j) - (mp)^2}{m^2} \rightarrow 0.$$

Thus indeed for any $t > 0$

$$\mathbb{P} \left(\left| \frac{1}{m} \sum_{i=1}^m X_i - p \right| \geq t \right) \rightarrow 0.$$

Finally,

$$\mathbb{P}(\text{noninjectivity}) \geq \mathbb{P} \left(\sum_{i=1}^m X_i > m - n \right) = \mathbb{P} \left(\frac{1}{m} \sum_{i=1}^m X_i > 1 - \frac{1}{c} \right).$$

Combining with Lemma 17 we get that

$$\frac{1}{m} \sum_{i=1}^m X_i \xrightarrow{\mathbb{P}} \frac{1}{2} \operatorname{erfc} \left(-\frac{1}{\sqrt{2c}} \right)$$

Thus

$$\mathbb{P}(\text{noninjectivity}) \geq \mathbb{P} \left(\frac{1}{m} \sum_{i=1}^m X_i > 1 - \frac{1}{c} \right) \rightarrow \begin{cases} 0 & c > c^* \\ 1 & c < c^*. \end{cases}$$

■




Figure 8: A Proof aid of the DSS condition when $W \in \mathbb{R}^{8 \times 2}$. The blue dots are x_0 and x_1 , the pink are the weight matrix rows, black lines denote the boundaries between adjacent wedges, and the multi-colored line is $\ell^{x_0, x_1}(t)$. This line changes color each time it crosses into a new wedge.

A.3 Proof of Theorem 8

We divide the proof into parts for ease of understanding.

Lemma 21 (Inverse Lipschitz Constant: Face Adjacent Wedges) *Let $W \in \mathbb{R}^{m \times n}$ have a DSS w.r.t. every $x \in \mathbb{R}^n$ and $x_0, x_1 \in \mathbb{R}^n$. Define*

$$\forall t \in [0, 1], \quad \ell^{x_0, x_1}(t) = (1-t)x_0 + tx_1$$

and suppose that x_0 and x_1 are such that there is a $t' \in (0, 1)$ such that for all $t \in [0, 1]$

$$W|_{S(\ell^{x_0, x_1}(t), W)} = \begin{cases} W|_{S(x_0, W)} & \text{if } t < t' \\ W|_{S(x_1, W)} & \text{if } t' < t \end{cases}. \quad (13)$$

Suppose further that x_0, x_1 are such that there is some $\delta > 0$ such that for all $\delta x \in \mathbb{R}^n$, if $\|\delta x\|_2 < \delta$ then there is a $t' + \delta t$ such that

$$W|_{S(\ell^{x_0, x_1 + \delta x}(t), W) \geq 0} = \begin{cases} W|_{S(x_0, W)} & \text{if } t < t' + \delta t \\ W|_{S(x_1, W)} & \text{if } t' + \delta t < t \end{cases}. \quad (14)$$

Then

$$\|\text{ReLU}(Wx_0) - \text{ReLU}(Wx_1)\|_2 \geq \frac{1}{\sqrt{2}} \min(\sigma(W|_{S(x_0, W)}), \sigma(W|_{S(x_1, W)})) \|x_0 - x_1\|_2,$$

where $\sigma(M)$ is the smallest singular value of the matrix M .

Remark 22 *The conditions of Lemma 21 on x_0 and x_1 may look very odd, but they have a very natural geometric meaning. The DSS can be thought of as slicing \mathbb{R}^n into wedges by a series of hyperplanes that have the rows of W as normals.*

The condition in (13) is interpreted as that the line segment that connects x_0 to x_1 passes from x_0 's wedge into x_1 's wedge without passing through any wedges in between (see Figure 8a, as opposed to Figure 8b). This implies that x_0 and x_1 must be in wedges that share a boundary. The condition in (14) requires that the wedges of x_0 and x_1 share a face, and not just a corner.

Proof We denote

$$W|_{S(x_0, W)} = W_0, W|_{S(x_1, W)} = W_1, W_{0 \cap 1} = W_0 \cap W_1, W_0 = \begin{bmatrix} W_{0 \cap 1} \\ W_{0 \setminus 1} \end{bmatrix}, W_1 = \begin{bmatrix} W_{0 \cap 1} \\ W_{1 \setminus 0} \end{bmatrix}.$$

First we will show that if $w_i, w_{i'} \in W_{0 \setminus 1}$, then w_i and $w_{i'}$ must be parallel. From equation (13) and the continuity of $\text{ReLU}(W^{\ell^{x_0, x_1}}(t))$ w.r.t. t , we have that there is a $t' \in (0, 1)$ such that

$$\langle w_i, (1 - t')x_0 + t'x_1 \rangle = 0 = \langle w_{i'}, (1 - t')x_0 + t'x_1 \rangle.$$

If w_i and $w_{i'}$ are not parallel, then let δx be some vector $0 < \|\delta x\| < \delta$ that is perpendicular to w_i but not $w_{i'}$ then

$$\begin{aligned} \langle w_i, (1 - t')x_0 - t'\delta x + t'x_1 \rangle &= -t \langle w_i, \delta x \rangle = 0, \\ \langle w_{i'}, (1 - t')x_0 - t'\delta x + t'x_1 \rangle &= -t \langle w_{i'}, \delta x \rangle \neq 0, \end{aligned}$$

which contradicts (14). W.l.o.g. the same argument applies to $W_{1 \setminus 0}$ and also it is straight forward to see that all elements of $W_{0 \setminus 1}$ must be anti-parallel to all elements of $W_{1 \setminus 0}$. From $W_{1 \setminus 0}$ and $W_{0 \setminus 1}$ parallelism, there is a $c \geq 0$ such that for all $x \in \mathbb{R}^n$,

$$\|W_{1 \setminus 0}x\| = c^2 \|W_{0 \setminus 1}x\|.$$

Assume that $c \geq 1$, then

$$\begin{aligned} \|W_0x_0 - W_1x_1\|_2^2 &= \|W_{0 \cap 1}x_0 - W_{0 \cap 1}x_1\|_2^2 + \|W_{0 \setminus 1}x_0\|_2^2 + \|W_{1 \setminus 0}x_1\|_2^2 \\ &= \|W_{0 \cap 1}x_0 - W_{0 \cap 1}x_1\|_2^2 + \|W_{0 \setminus 1}x_0\|_2^2 + c^2 \|W_{0 \setminus 1}x_1\|_2^2 \\ &\geq \|W_{0 \cap 1}x_0 - W_{0 \cap 1}x_1\|_2^2 + \|W_{0 \setminus 1}x_0\|_2^2 + \|W_{0 \setminus 1}x_1\|_2^2 \\ &\geq \|W_{0 \cap 1}x_0 - W_{0 \cap 1}x_1\|_2^2 + \frac{1}{2} \|W_{0 \setminus 1}(x_0 - x_1)\|_2^2 \\ &\geq \frac{1}{2} \|W_0x_0 - W_0x_1\|_2^2 \\ &\geq \frac{\sigma(W_0)^2}{2} \|x_0 - x_1\|_2^2. \end{aligned}$$

The antepenultimate inequality comes as from the definition of $W_{0 \setminus 1}$, we have that $W_{0 \setminus 1}x_0$ and $-W_{0 \setminus 1}x_1$ are the same sign, thus (34) applies. In the case that $c < 1$, then the rolls of $W_{0 \setminus 1}$ and $W_{1 \setminus 0}$ can be switched, and the same result (with $\sigma(W_1)$ in place of $\sigma(W_0)$) is obtained. In either case,

$$\|W_0x_0 - W_1x_1\|_2^2 \geq \frac{1}{2} \min \left(\sigma(W|_{S(x_0, W)})^2, \sigma(W|_{S(x_1, W)})^2 \right) \|x_0 - x_1\|_2^2.$$

■

Lemma 23 (Inverse Lipschitz Constant: Connected through Faces) *Let $W \in \mathbb{R}^{m \times n}$ have a DSS w.r.t. every $x \in \mathbb{R}^n$. Let x_0, x_1 be such that the line connecting passes through n_t wedges, and through their faces (in the sense of Lemma 22). Then*

$$\|\text{ReLU}(Wx_0) - \text{ReLU}(Wx_1)\|_2 \geq \frac{1}{\sqrt{2n_t}} \min_{t \in [0,1]} \sigma(W|_{S(\ell^{x_0, x_1}(t), W)}) \|x_0 - x_1\|_2.$$

Proof Let $t_1 = 0$, $t_{n_t} = 1$, and let $T = \{t_k\}_{k=1}^{n_t}$ such that $t_k < t_{k+1}$ for $k < n_t$ and $x_k := \ell^{x_0, x_1}(t_k)$ are each in different wedges. Let $c = \frac{1}{\sqrt{2}} \min_{t \in [0,1]} \sigma(W|_{S(\ell^{x_0, x_1}(t), W)})$, then

$$c \|x_0 - x_1\|_2 \leq \sum_{k=1}^{n_t-1} c \|x_{t_k} - x_{t_{k+1}}\|_2 \quad (15)$$

then by Lemma 21,

$$\sum_{k=1}^{n_t-1} c \|x_{t_k} - x_{t_{k+1}}\|_2 \leq \sum_{k=1}^{n_t-1} \|\text{ReLU}(Wx_{t_k}) - \text{ReLU}(Wx_{t_{k+1}})\|_2$$

and by Lemma 34 we have

$$\sum_{k=1}^{n_t-1} \|\text{ReLU}(Wx_{t_k}) - \text{ReLU}(Wx_{t_{k+1}})\|_2 \leq \sqrt{n_t} \|\text{ReLU}(Wx_0) - \text{ReLU}(Wx_1)\|_2. \quad (16)$$

Combining (15) - (16) yields

$$\|\text{ReLU}(Wx_0) - \text{ReLU}(Wx_1)\|_2 \geq \frac{1}{\sqrt{2n_t}} \min_{t \in [0,1]} \sigma(W|_{S(\ell^{x_0, x_1}(t), W)}) \|x_0 - x_1\|_2.$$

■

Proof [Proof of Theorem 8] Let x_0 and x_1 be given. If the line connecting x_0 and x_1 does not pass through any wedge corners, then we can apply Lemma 23 directly, and get that

$$\|\text{ReLU}(Wx_0) - \text{ReLU}(Wx_1)\|_2 \geq \frac{1}{\sqrt{2n_t}} \min_{t \in [0,1]} \sigma(W|_{S(\ell^{x_0, x_1}(t), W)}) \|x_0 - x_1\|_2.$$

We now argue that the number of wedges that $\ell^{x_0, x_1}(t)$ passes through (and so n_t) is at most m . For each $j \in [[m]]$,

$$\text{ReLU}(W\ell^{x_0, x_1}(t))|_j = \langle w_j, \ell^{x_0, x_1}(t) \rangle$$

is monotone increasing or decreasing. This implies that each $w_j \in W$ can enter or exit the DSS w.r.t. $\ell^{x_0, x_1}(t)$ at most once, therefore the total number of unique DSS w.r.t. $\ell^{x_0, x_1}(t)$ (that is wedges $\ell^{x_0, x_1}(t)$ pass through) is at most m . Hence, $n_t \leq m$. Clearly

$$\min_{t \in [0,1]} \sigma(W|_{S(\ell^{x_0, x_1}(t), W)}) \geq \min_{x \in \mathbb{R}^n} \sigma(W|_{S(x, W)}),$$

hence we have

$$\|\text{ReLU}(Wx_0) - \text{ReLU}(Wx_1)\|_2 \geq \frac{C}{\sqrt{m}} \|x_0 - x_1\|_2.$$

Now we show that if Lemma 23 does not apply to two points (namely, the line $\ell^{x_0, x_1}(t)$ passes through a corner in the sense of Remark 22), then the two points can be perturbed an arbitrarily small amount, so that the perturbed points do satisfy Lemma 23. The corners (again, as in Remark 22) describe the points which are orthogonal to at least two $w_{j_1}, w_{j_2} \in W$. w_{j_1} and w_{j_2} must not be parallel to each other (otherwise Lemma 21 would apply), thus the set of points orthogonal to both w_{j_1} and w_{j_2} constitute a $n - 2$ dimensional linear space in \mathbb{R}^n .

Let x_0 and x_1 be such that $\ell^{x_0, x_1}(t)$ intersects one of these corners. By considering $\tilde{x}_0 = \delta x + x_0, \tilde{x}_1 = \delta x + x_1$ where δx is perpendicular to $x_1 - x_0$, we can obtain a line $\ell^{\tilde{x}_0, \tilde{x}_1}(t)$ so that $\ell^{\tilde{x}_0, \tilde{x}_1}(t)$ and $\ell^{x_0, x_1}(t)$ do not intersect. The choice of δx is $n - 1$ dimensional, thus for every $\delta > 0$ and w_{j_1} and w_{j_2} (that are non-perpendicular) there is a δx so that $\|\delta x\|_2 < \delta$ and $\ell^{\tilde{x}_0, \tilde{x}_1}(t)$ does not intersect the corner of w_{j_1} and w_{j_2} .

Consider a sequence of $\tilde{x}_0^{(i)}, \tilde{x}_1^{(i)}, i = 1, \dots$ such that $\lim_{i \rightarrow \infty} (\tilde{x}_0^{(i)}, \tilde{x}_1^{(i)}) = (x_0, x_1)$ and $\ell^{\tilde{x}_0^{(i)}, \tilde{x}_1^{(i)}}(t)$ does not pass through a corner for any i . Given that $\|\cdot\|_2$ and $\text{ReLU}(W(\cdot))$ are continuous and so by Lemma 23

$$\left\| \text{ReLU}(W\tilde{x}_0^{(i)}) - \text{ReLU}(W\tilde{x}_1^{(i)}) \right\|_2 - \frac{C}{\sqrt{2m}} \left\| \tilde{x}_0^{(i)} - \tilde{x}_1^{(i)} \right\|_2 \geq 0,$$

thus

$$\begin{aligned} & \left\| \text{ReLU}(Wx_0) - \text{ReLU}(Wx_1) \right\|_2 - \frac{C}{\sqrt{2m}} \|x_0 - x_1\|_2, \\ &= \lim_{i \rightarrow \infty} \left\| \text{ReLU}(W\tilde{x}_0^{(i)}) - \text{ReLU}(W\tilde{x}_1^{(i)}) \right\|_2 - \frac{C}{\sqrt{2m}} \left\| \tilde{x}_0^{(i)} - \tilde{x}_1^{(i)} \right\|_2 \geq 0. \end{aligned}$$

and so

$$\left\| \text{ReLU}(Wx_0) - \text{ReLU}(Wx_1) \right\|_2 \geq \frac{C}{\sqrt{2m}} \|x_0 - x_1\|_2. \quad \blacksquare$$

Remark 24 (Factor of $\frac{1}{\sqrt{m}}$ in Theorem 8.) *Throughout this paper we define the discrete norm of $y \in \mathbb{R}^d$ as*

$$\|y\|_2 = \left(\sum_{j=1}^d [y]_j^2 \right)^{\frac{1}{2}}.$$

This is to be contrasted with the norm that arise from the discretization of the L_2 function norm on a finite domain. For example, if we instead thought of y as a discrete sampling of a

continuous function $\tilde{y} \in L_2([0, 1])$, such that $\forall j = 1, \dots, d$

$$\tilde{y}\left(\frac{j-1}{m}\right) = [y]_j,$$

then we could approximate the $L_2([0, 1])$ norm of \tilde{y} by

$$\|\tilde{y}\|_{L_2([0,1])} \approx \|y\|_{l_2([0,1])} := \frac{1}{\sqrt{m}} \left(\sum_{j=1}^d [y]_j^2 \right)^{\frac{1}{2}}.$$

If we express Theorem 8 in terms of $\|\cdot\|_{l_2([0,1])}$, then it would become

$$\|\text{ReLU}(Wx_0) - \text{ReLU}(Wx_1)\|_{l_2([0,1])} \geq \frac{C(W)}{m} \|x_0 - x_1\|_2.$$

A.4 Theorem 13

Example 1 (Applying Theorem 13, One Channel) Consider a layer of the form $\text{Reshape}(\text{ReLU}(Wx))$ where $W = [C_1^T, \dots, C_d^T]^T$ and each C_k is a convolution operator with kernel c_k . Suppose further that $W \in \mathbb{R}^{4 \times 4 \times 1024 \times 100} = \mathbb{R}^{16384 \times 100}$ (as in Radford et al., 2015). The reshaping operator takes the 16384 single-channel output of W and transforms it into a multi-channel signal. This is necessary for subsequent convolutions, but plays no role in injectivity. Let $\mathbf{n}_c = 8$, and the 2×2 convolution kernels be given as

$$c_1 = \begin{bmatrix} 3 & -1 \\ -1 & -1 \end{bmatrix}, \quad c_2 = \begin{bmatrix} -1 & 3 \\ -1 & -1 \end{bmatrix}, \quad c_3 = \begin{bmatrix} -1 & -1 \\ 3 & -1 \end{bmatrix}, \quad c_4 = \begin{bmatrix} -1 & -1 \\ -1 & 3 \end{bmatrix}$$

$$c_5 = -c_1, c_6 = -c_2, c_7 = -c_3, c_8 = -c_4.$$

Directly proving that a 16384×100 dimension operator has a DSS w.r.t. every $x \in \mathbb{R}^{100}$ is daunting. However, since each layer of the operator is given by one of only 8 simple convolutions, we can leverage Theorem 13 to significantly simplify the problem. Choosing $P = (2, 2)$, implies that $\mathcal{P}_{(2,2)}(c_k) = c_k$, and so $W|_{\mathcal{P}_{(2,2)}} = \bigcup_{k=1}^8 c_k$. Further, it is easy to see, that $\{c_1, c_2, c_3, c_4\}$ is a basis for $\mathbb{R}^{2 \times 2}$, so Corollary 6 applies, $W|_{\mathcal{P}_P}$ has a DSS of \mathbb{R}^4 , and by Theorem 13, $\text{ReLU}(Wx)$ is injective.

An example when a layer is injective but P in Theorem 13 must be greater than O is when W is a convolution of 4 kernels of width 3

$$c_1 = [1 \ 0 \ -1], \quad c_2 = [1 \ 0 \ 1], \quad c_3 = [-1 \ 0 \ 1], \quad c_4 = [-1 \ 0 \ -1].$$

If we choose $P = (3)$, then $W|_{\mathcal{P}_{(3)}} = \bigcup_{k=1}^4 \{c_k\}$. only has four elements, and so cannot have a DSS of \mathbb{R}^3 w.r.t. every $x \in \mathbb{R}^3$ by Corollary 6. If we however choose $P = (4)$, then $\mathcal{P}_{(4)}(c_k) = \{[c_k \ 0], [0 \ c_k]\}$ and $W|_{\mathcal{P}_{(4)}}$ has a DSS of \mathbb{R}^4 w.r.t. all $x \in \mathbb{R}^4$ (from Corollary 6), so W has a DSS w.r.t. all $x \in \mathbb{R}^N$.

In order to prove Theorem 13, we must first prove Lemma 14, which we do here.

Proof For every $k = 1, \dots, K$ $W_k|_{S(x, W_k)}$ has a DSS of Ω_k with respect to $P_{\Omega_k}(x)$ where P_{Ω_k} is the orthogonal projection of \mathbb{R}^n onto Ω_k . For every $w_{k,\ell}$,

$$\langle w_{k,\ell}, P_{\Omega_k}(x) \rangle = \langle w_{k,\ell}, x \rangle$$

thus $S(x, W_k) \subset S(x, W)$. From $\Omega_k \subset \text{span}(W_k|_{S(x, W_k)})$ for each k we have a set spanning Ω_k that lie in $W|_{S(x, W)}$, hence

$$\mathbb{R}^n = \text{span}(\Omega_1, \dots, \Omega_K) \subset \text{span}\left(\bigcup_{k=1}^K \bigcup_{\ell=1}^{N_k} \{w_{k,\ell}\}\right) = \text{span}(W)$$

contains a DSS of \mathbb{R}^n w.r.t. x . The set $\bigcup_{k=1}^K \{w_{k,\ell}\}_{\ell=1}^{N_k}$ has no dependence on x , thus it is true for all $x \in \mathbb{R}^n$. \blacksquare

Lemma 25 *Given a convolution operator $C \in \mathbb{R}^{N \times N}$. Let $0 \leq V$ be such that $V + O \leq N$. For each $x \in C$,*

$$\text{aug}_{V:V+O}(x) \in C, \quad \text{where} \quad (\text{aug}_{V:V+O}(x))_J = \begin{cases} (x)_{J-V} & \text{if } 1 + V \leq J \leq V + O \\ 0 & \text{otherwise} \end{cases}.$$

Note that $(\text{aug}_{V:V+O}(x))_J$ restricted to the indices $1 + V$ through $V + P$ is exactly x .

The above is always well defined, as $c_{H-T+O+1}$ is well defined iff $1 \leq H - T + O + 1 \leq O$ (recall that the kernel c is of width O), and so

$$1 \leq H - T + O + 1 \leq O \iff -O \leq H - T \leq -1 \iff 1 \leq T - H \leq O$$

Now we proceed with the proof of Theorem 13.

Proof [Proof of Theorem 13] The strategy for this proof will be to use Lemma 14 by decomposing \mathbb{R}^n into some number of different domains $\{\Omega_v\}_{v=1}^{n_v}$, each of which are a restriction of \mathbb{R}^n to P non-zero components. For each Ω_v , the elements of W that lie in Ω_v can be identified as the elements in $W|_{\mathcal{P}_P}$. If for one v the components of $W|_{\mathcal{P}_P}$ form a DSS of Ω_v w.r.t. every $x \in \Omega_v$, then it has a DSS for every such Ω_k , so we can apply Lemma 14 and get that W has a DSS of \mathbb{R}^n w.r.t. all $x \in \mathbb{R}^n$.

Given an offset $V' \geq 0$ such that $V' + P \leq N$, define $\Omega_{V'}$ as the subspace of all vectors $x \in \mathbb{R}^N$ such that

$$\Omega_{V'} = \{x \in \mathbb{R}^N : x_J = 0 \text{ if } 1 + V' \not\leq J \text{ or } J \not\leq V' + P\}.$$

From Lemma 25, for any $k = 1, \dots, n_v$, if $x^{k,P} \in \mathcal{P}_P(c_k)$, C_k is a submatrix of W and

$$\text{aug}_{V':V'+P}(x^{k,P}) \in C_k.$$

Further, for any such $x^{k,P}$,

$$\text{aug}_{V':V'+P}(x^{k,P}) \in \Omega_{V'}.$$

If $W|_{\mathcal{P}_P}$ contains a DSS of \mathbb{R}^P w.r.t. all $x \in \mathbb{R}^P$, then

$$\text{aug}_{V':V'+P}(W|_{\mathcal{P}_P}) \text{ contains a DSS of } \Omega_{V'} \text{ w.r.t. all } x \in \Omega_{V'}. \quad (17)$$

This follows from Lemma 14. From Lemma 25 for any V' such that

$$\text{aug}_{V':V'+P}(W|_{\mathcal{P}_P}) \in W, \quad (18)$$

if $W|_{\mathcal{P}_P}$ has a DSS of \mathbb{R}^P w.r.t. all $x \in \mathbb{R}^P$, then W contains a DSS of $\Omega_{V'}$ for all $0 \leq V' \leq N - P$. Finally, note that $\text{span}(\{\Omega_{V'}\}_{V'=0}^{N-P})$, and so using (17), (18) we can apply Lemma 14 and find that W contains a DSS of \mathbb{R}^N w.r.t. all $x \in \mathbb{R}^N$. \blacksquare

For a multi-channel input (with n_c channels) $x \in \underbrace{\mathbb{R}^N \times \mathbb{R}^N \times \cdots \times \mathbb{R}^N}_{n_c \text{ times}}$, a multi-channel convolution C on x is given by $Cx = \sum_{q=1}^{n_c} C_q x_q$ where C_o is a convolution on \mathbb{R}^N (defined by Definition 10) and $x_o \in \mathbb{R}^N$ is the restriction of x to the o 'th channel. Because of the additive structure of multi-channel convolutions a n_c over $\mathbb{R}^N = \mathbb{R}^{N_1} \times \cdots \times \mathbb{R}^{N_p}$ dimensional domain of width $O = (O_1, \dots, O_p)$ with kernels c_1, \dots, c_{n_c} is equivalent to a single convolution of width $(O, n_c) = (O_1, \dots, O_p, n_c)$ over $\mathbb{R}^{(N,p)} = \mathbb{R}^{N_1} \times \cdots \times \mathbb{R}^{N_p} \times \mathbb{R}^{n_c}$. This follows from

$$(Cx)_J = \sum_{q=1}^{n_c} (C_q x_q)_J = \sum_{q=1}^{n_c} \sum_{I=1}^O (c_q)_{O-I-1} (x_q)_{J+I} = \sum_{(I,n_c)=1}^{(O,q)} c_{(O,n_c)-(I,q)-1} x_{(J,q)+(I,q)}.$$

Appendix B. Robustness to Layer, Instance and Group Normalization

As described in Section 2.5, layer, instance and group normalization take place during both training and execution and, unlike weight and spectral normalization, the normalization is done on the input/outputs of layers instead of on the weight matrices. For these normalizations (1) is modified so that it becomes

$$N(z) = \phi_L(W_L M_L(\cdots \phi_2(W_2 M_2(\phi_1(W_1 z + b_1)) + b_2) \cdots + b_L)) \quad (19)$$

where $M_\ell: \mathbb{R}^{n_{i+1}} \rightarrow \mathbb{R}^{n_{i+1}}$ are normalization functions that are many-to-one. In general $\phi_\ell(W_\ell M_\ell(\cdot) + b_\ell)$ will not be injective for any $\phi_\ell, W_\ell, b_\ell$ on account of M_ℓ , but for all of the mentioned normalization techniques we can get near injectivity. Before we descend into the particular we make the following observation about normalization methods that obey a certain structure.

Definition 26 (Scalar-Augmented Injective Normalization) *Let $M_\ell(x): \mathbb{R}^n \rightarrow \mathbb{R}^n$ be a normalization function that is understood to be many-to-one. We say that $M_\ell(x)$ is scalar-augmented injective if there exists a function $m_\ell(x): \mathbb{R}^n \rightarrow \mathbb{R}^k$ where $k \ll n$ and $\tilde{M}_\ell: \mathbb{R}^n \times \mathbb{R}^k \rightarrow \mathbb{R}^n$ such that*

$$M_\ell(x) := \tilde{M}_\ell(x; m_\ell(x))$$

and $\tilde{M}_\ell(x; m_\ell(x))$ is injective on x given $m_\ell(x)$.

An example of a normalization function that is scalar-augmented injective is

$$M_\ell(x) = \frac{x}{\|x\|_2}. \quad (20)$$

For this choice of M_ℓ , $k = 1$, and

$$\tilde{M}_\ell(x; c) = \frac{x}{c} \quad m_\ell(x) = \|x\|_2.$$

With this definition, we can prove the following trivial but useful result

Lemma 27 (Restricted Injectivity of Scalar-Augmented Normalized Networks)

Let N be a deep network of the form in (19) and let each $\phi_\ell(W_\ell \cdot)$ be layer-wise injective. Let the normalization functions $\{M_\ell\}_{\ell=1, \dots, L}$ each be scalar-augmented injective. Then given $\{m_\ell(x)\}_{\ell=1, \dots, L}$, the network

$$\tilde{N}(z; m_1, \dots, m_\ell) = \phi_L(W_L \tilde{M}_L(\dots \phi_2(W_2 \tilde{M}_2(\phi_1(W_1 z + b_1); m_2) + b_2) \dots + b_L; m_L))$$

is injective.

Proof The proof of Lemma 27 follows from a straightforward application of induction, combined with Definition 26. \blacksquare

Remark 28 Note that Lemma 27 implies that for a fixed $\{m_\ell(x)\}_{\ell=1, \dots, L}$, there is at most one value of z such that

$$\tilde{N}(z; m_1, \dots, m_\ell) = N(z), \quad (21)$$

where $N(z)$ is given by (19), and that the z 's on both sides of (21) are the same. It is still entirely possible that there are another choice of z' , $\{m_\ell(x)\}_{\ell=1, \dots, L}$ such that

$$\tilde{N}(z; m_1, \dots, m_\ell) = \tilde{N}(z'; m'_1, \dots, m'_\ell).$$

An example of this would be if M_ℓ is of the form in (20), then

$$\tilde{M}_\ell(x; m_\ell(x)) = \tilde{M}_\ell(2x; 2m_\ell(x)).$$

In other words, Lemma 27 implies that the deep network is injective (in z) for a fixed $\{m_\ell(x)\}_{\ell=1, \dots, L}$, but it may still not be injective for all z and $\{m_\ell(x)\}_{\ell=1, \dots, L}$.

With Lemma 27 in tow, we can show that layer, instance, and group normalization are all scalar-augmented injective normalizations, so Lemma 27 applies and yields a kind of injectivity. Layer, instance and group normalization are all related insofar as they can all be expressed in the same abstract form. For a given input x , all three break x up into K parts denoted $\{x|_{S_k}\}_{k=1, \dots, K}$ such that for each $k = 1, \dots, K$

$$\begin{aligned} \mu_k &= \frac{1}{m} \sum_{i=1}^m (x|_{S_k})_i & \sigma_k^2 &= \frac{1}{m} \sum_{i=1}^m ((x|_{S_k})_i - \mu_k)^2 \\ (\hat{x}|_{S_k})_i &= \frac{(x|_{S_k})_i - \mu_k}{\sqrt{\sigma_k^2 + \epsilon}} & M(x)|_{S_i} &= \gamma_k (\hat{x}|_{S_k})_i + \beta_k. \end{aligned}$$

The differences between the three normalization are how $\{S_k\}_{k=1,\dots,K}$ is chosen. For layer normalization $K = 1$ and the normalization is applied to the entire input signal. For instance normalization, there is one S_k for each channel, and the $x|_{S_k}$ restricts x to just one channel of inputs, that is the normalization is done channel-wise. Group normalization is part way between these two, where k is less than the number of channels, and channels are batched together.

In any case, for any of these normalization methods, they are all scalar-augmented injective normalization where

$$M_\ell(x) = \tilde{M}_\ell(x; \{\sigma_{k,\ell}, \mu_{k,\ell}\}_{k=1,\dots,K}).$$

Thus, by Lemma 27 their corresponding deep networks are all injective, provided that for each ℓ , $\{\sigma_{k,\ell}, \mu_{k,\ell}\}_{k=1,\dots,K}$ is saved.

B.1 Pooling Operations

Although pooling may have a different aim than typical normalization, we consider it in this section, as it is mathematically similar to (19). Pooling is similar to layer, instance and group normalization in the sense that they partition the input space into K disjoint pieces, and then output a weighted average upon each piece. Specifically, if $M_p(x): \mathbb{R}^n \rightarrow \mathbb{R}^K$ where for $k = 1, \dots, K$,

$$M_p(x)|_{S_k} = \|x|_{S_k}\|_p$$

where $\|\cdot\|_p$ is the discrete p norm of x restricted to the set S_k . For $p = 1$ this is the mean of the absolute value, for $p = 2$ this is the Euclidean mean and for $p = \infty$ it is the maximum of the absolute value. The injectivity of this operation in the cases where $p = 1, 2, \infty$ is considered in the work Bruna et al. (2013).

Appendix C. Proofs of Theorem 15

Before we present the proof of Theorem 15, we present the following stand-alone lemma which is used during the proof of Theorem 15. This lemma shows that, generically, an injective piecewise-affine neural network followed by an orthogonal projector is still injective, provided that the orthogonal projector projects onto a space of sufficient dimension.

Let $\mathcal{V}(k, D)$ denote the set of k -tuples (v_1, v_2, \dots, v_k) where v_j are orthonormal vectors in \mathbb{R}^D . Such vectors span a k -dimensional linear space. Furthermore, let $G(k, D)$ denote the set of k -dimensional linear subspaces of \mathbb{R}^D , and for $V \in G(k, D)$, $V = \text{span}(v_1, v_2, \dots, v_k)$, let $P_V = P_{(v_1, v_2, \dots, v_k)}: \mathbb{R}^D \rightarrow \mathbb{R}^D$ be an orthogonal projection which image is the space V . As the dimension of the orthogonal group $O(k)$ is $k(k-1)/2$ and by Milnor and Stasheff (1974), the set $G(k, D)$, called the Grassmannian, is a smooth algebraic variety, of dimension $k(D-k)$ and the dimension of $\mathcal{V}(k, D)$ is $k(D-k) + k(k-1)/2 = k(2D-k-1)/2$. When $V, W \in G(k, D)$ are two k -dimensional linear subspaces of \mathbb{R}^D , we define their distance in $G(k, D)$ to be the operator norm

$$d_{G(k,D)}(V, W) = \|P_V - P_W\| \tag{22}$$

where P_V and P_W are orthogonal projections in \mathbb{R}^D onto V and W , respectively, see (p. 49 Mattila, 1995).

Lemma 29 *Let $H_\theta \in \mathcal{NN}(n, D)$, $D > m \geq 2n + 1$ be a neural network such that $H_\theta : \mathbb{R}^n \rightarrow \mathbb{R}^D$ is injective. Let $X_\theta = \{V \in G(m, D) : P_V \circ H_\theta : \mathbb{R}^n \rightarrow \mathbb{R}^D \text{ is injective}\}$. Then the set X_θ is an intersection of countably many open and dense subsets of $G(m, D)$, that is, elements of X_θ are generic. Moreover, the $m(D - m)$ dimensional Hausdorff measure of the complement of X_θ in $G(m, D)$ is zero.*

Note that for $V \in X_\theta$, we have $P_V \circ H_\theta \in \mathcal{NN}(n, D)$.

Proof We use that fact that $H_\theta : \mathbb{R}^n \rightarrow \mathbb{R}^D$ is injective and locally Lipschitz-smooth. Recall that $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is locally Lipschitz if for any compact set $\mathcal{C} \subset \mathbb{R}^n$ there is $L_{\mathcal{C}} > 0$ such that $\|f(x) - f(y)\|_{2,m} \leq L_{\mathcal{C}} \|x - y\|_{2,n}$ for all $x, y \in \mathcal{C}$.

Let $v \in \mathbb{R}^m$ be a unit vector and let $Q_v : \mathbb{R}^D \rightarrow \mathbb{R}^D$ be the projection $Q_v(z) = z - (z \cdot v)v$. Let

$$\mathcal{P} = \{(x, y) \in \mathbb{R}^n \times \mathbb{R}^n : x \neq y\}.$$

As $H_\theta : \mathbb{R}^n \rightarrow \mathbb{R}^D$ is an injection, we can define $\mathbb{S}^{D-1} = \{w \in \mathbb{R}^D : \|w\|_{2,D} = 1\}$ and the map

$$s_\theta : \mathcal{P} \rightarrow \mathbb{S}^{D-1}, \quad s_\theta(x, y) = \frac{H_\theta(x) - H_\theta(y)}{\|H_\theta(x) - H_\theta(y)\|_{2,D}}. \quad (23)$$

Observe that $s_\theta(\mathcal{P}) = \{s_\theta(x, y) : (x, y) \in \mathcal{P}\} = -s_\theta(\mathcal{P}) = \{-s_\theta(x, y) : (x, y) \in \mathcal{P}\}$.

As observed in the proof of the Whitney's embedding theorem (Chapter 2, Theorem 3.5 Hirsch, 2012), the map $Q_w \circ H_\theta : \mathbb{R}^n \rightarrow \mathbb{R}^D$ is an injection when $w \notin s_\theta(\mathcal{P})$. To see this, assume that $w \in \mathbb{S}^{D-1}$ satisfies $w \notin s_\theta(\mathcal{P})$, that is, w is not in the image of s_θ . Then, if there are $x, y \in \mathbb{R}^n$, $x \neq y$ such that $Q_w H_\theta(x) = Q_w H_\theta(y)$, we see that there is $t \in \mathbb{R}$ such that $H_\theta(x) - H_\theta(y) = tw$. By changing the roles of x and y , in needed, without loss of generality we may assume that $t \geq 0$. As w is a unit vector and $H_\theta(x) \neq H_\theta(y)$, this yields $t = \|H_\theta(x) - H_\theta(y)\|_{2,D} \neq 0$ and

$$w = (H_\theta(x) - H_\theta(y))/t = (H_\theta(x) - H_\theta(y))/\|H_\theta(x) - H_\theta(y)\|_{2,D} = s_\theta(x, y),$$

which is in contradiction with the assumption that $w \notin s_\theta(\mathcal{P})$. Hence, $w \notin s_\theta(\mathcal{P})$ yields that $Q_w \circ H_\theta : \mathbb{R}^n \rightarrow \mathbb{R}^D$ is an injection.

We consider the image $s_\theta(\mathcal{P}) \subset \mathbb{S}^{D-1}$. For $h > 0$, let

$$\mathcal{P}_h = \{(x, y) \in \overline{B}^n(0, h^{-1}) \times \overline{B}^n(0, h^{-1}) : \|H_\theta(x) - H_\theta(y)\|_{2,D} \geq h\}.$$

As $H_\theta : \overline{B}^n(0, h^{-1}) \rightarrow \mathbb{R}^m$ is Lipschitz-smooth with some Lipschitz constant L_h , we see that the map $s_\theta : \mathcal{P}_h \rightarrow \mathbb{S}^{D-1}$ is Lipschitz-smooth. Since the set \mathcal{P} has the Hausdorff dimension $2n$ and the map $s_\theta : \mathcal{P}_h \rightarrow \mathbb{S}^{D-1}$ is Lipschitz-smooth, the Hausdorff dimension of the set $s_\theta(\mathcal{P}_h)$ is at most $2n$, see for example Morgan (p. 26 2016). The set \mathcal{P} is the union of all sets \mathcal{P}_{h_j} , where $h_j = 1/j$ and $j \in \mathbb{Z}$. By Mattila (p. 59 1995) the Hausdorff dimension of a countable union of sets S_j is the supremum of the Hausdorff dimension of the sets S_j . Hence $s_\theta(\mathcal{P}) = \bigcup_{j=1}^{\infty} s_\theta(\mathcal{P}_{h_j})$ has the Hausdorff dimension less or equal $2n$.

Since the dimension $D - 1$ of \mathbb{S}^{D-1} is strictly larger than $2n$, we see that the set $s_\theta(\mathcal{P}_{h_j})$ is closed, its complement is an open and dense set, and thus the set $Y_1(\theta) := \mathbb{S}^{D-1} \setminus s_\theta(\mathcal{P})$ is an intersection of countably many open and dense sets.

Observe that as Q_{w_1} is a linear map, the map $Q_{w_1} \circ H_\theta$ is also a neural network that belongs in $\mathcal{NN}(n, D)$, and we can denote $Q_{w_1} \circ H_\theta = H_{\theta_1}$ with some parameters θ_1 . Thus we can repeat the above arguments using the map $Q_{w_1} \circ H_\theta : \mathbb{R}^n \rightarrow \text{span}(w_1)^\perp \equiv \mathbb{R}^{D-1}$ instead of $H_\theta : \mathbb{R}^n \rightarrow \mathbb{R}^D$. Repeating the above arguments $D-m$ times, can choose orthonormal vectors $w_j \in \mathbb{S}^{D-1}$, $j = 1, 2, \dots, D-m$, and sets $Y_j(\theta, w_1, \dots, w_{j-1}) \subset \mathbb{S}^{D-1} \cap (w_1, \dots, w_{j-1})^\perp$, which $D-j$ dimensional Hausdorff measures vanish and which complements are intersections of countably many open and dense sets. Let B_θ to be the set of all n -tuples (w_1, \dots, w_{D-m}) where $w_1 \in Y_1(\theta)$ and $w_j \in Y_j(\theta, w_1, \dots, w_{j-1})$ for all $j = 2, \dots, D-m$. Note that all such vectors w_j , $j = 1, 2, \dots, D-m$ are orthogonal vectors spanning a $D-m$ dimensional vector space V , and the map

$$P_V \circ H_\theta : \mathbb{R}^n \rightarrow \mathbb{R}^D, \quad \text{where } P_V = Q_{w_{D-m}} \circ \dots \circ Q_{w_2} \circ Q_{w_1}$$

is injective. By the above, construction the $(D-(m+1)/2)m$ dimensional Hausdorff measure of the complement of B_θ in $\mathcal{V}(D, m)$ is zero and B_θ is generic set. As the dimension of the set of the orthogonal basis in a m -dimensional vector space is $(m-1)m/2$, we obtain the claim. ■

Now that this lemma is established, we proceed to the proof of Theorem 15.

Our proof is divided into two parts. First we prove the universality part (the statement that an N_θ exists), and then show that we can make a small perturbation to the N_θ to obtain N'_θ , with the prescribed inverse Lipschitz Constant.

To prove the universality result, we combine the approximation results for neural networks and the low regularity version of the generic orthogonal projector technique used to prove the easy version of the Whitney's embedding theorem (Hirsch, 2012, Chapter 2, Theorem 3.5) that shows that a C^2 -smooth manifold of dimension n can be embedded in \mathbb{R}^{2n+1} with an injective, C^2 -smooth map. To prove the result, we first approximate $f(x)$ by a ReLU-type neural network that is only Lipschitz-smooth, so the graph of the map F_θ is only a Lipschitz-manifold. We note that limited regularity often causes significant difficulties for embedding results, as for example for the Lipschitz-smooth manifolds it is presently known only that a n -dimensional manifold can be embedded (without preserving distances) in the Euclidean space \mathbb{R}^N of dimension $N = (n+1)^2$, and the classical Whitney problem, whether a n -dimensional manifold can be embedded in \mathbb{R}^{2n+1} , is still open Luukkainen and Väisälä (1977); Cobzaş et al. (2019). Due to this lack of smoothness, we recall the details how this generic projector technique works.

Proof [Proof of Theorem 15, universality result] Let $\varepsilon > 0$ and $\mathcal{C} \subset \mathbb{R}^n$ be a compact set. As a continuous function $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ can be uniformly approximated in a compact set by a C^∞ -smooth function (see for example Thm. 2.29 Adams, 1975), we can without loss of generality assume that f smooth and therefore a locally Lipschitz function.

By classical results of approximation theory for shallow neural networks, see Hornik (1991); Leshno et al. (1993); Pinkus (1999), for any $L \geq 1$ there are \mathbf{n} and a neural network $F_\theta \in \mathcal{NN}(n, m, L, \mathbf{n})$ such that

$$\|f(x) - F_\theta(x)\|_{2,m} \leq \frac{1}{2}\varepsilon_1, \quad \text{for all } x \in \mathcal{C}. \quad (24)$$

We note that by using recent results for deep neural networks, for example by Yarotsky (2017), one can obtain efficient estimates on how a given accuracy ε_1 can be obtained using sufficiently large L and \mathbf{n} . Our aim is the perturb $F_\theta : \mathbb{R}^n \rightarrow \mathbb{R}^m$ so that it becomes injective.

We assume that $\mathcal{C} \subset B^n(0, r_1)$, where $B^n(0, r_1) \subset \mathbb{R}^n$ is an open ball having centre 0 and radius $r_1 > 0$. We denote the closure of this ball by $\overline{B}^n(0, r_1)$.

Let $D = m + n$, $\alpha > 0$, and define a map $H_\theta : \mathbb{R}^n \rightarrow \mathbb{R}^{m+n}$,

$$H_\theta(x) = (\alpha x, F_\theta(x)) \in \mathbb{R}^n \times \mathbb{R}^m = \mathbb{R}^D. \quad (25)$$

Observe that the map $H_\theta : \mathbb{R}^n \rightarrow \mathbb{R}^D$ is injective.

Now we continue with the proof Theorem 15 universality result. Let $V_0 = \{(0, 0, \dots, 0)\} \times \mathbb{R}^m \subset \mathbb{R}^D$. By Lemma 29, the complement of the set X_θ in the manifold $G(m, D)$ has measure zero and thus it does not contain any open subset of the manifold $G(m, D)$. Thus for any $\varepsilon_0 > 0$ there is $V \in X_\theta$ such that $d_{G(m, D)}(V, V_0) < \varepsilon_0$, see (22). Below, we use

$$\varepsilon_0 = \min \left(\frac{1}{16(1 + \alpha + \|F_\theta\|_{\mathcal{C}(\mathcal{C})})} \varepsilon_1, \frac{1}{2} \right).$$

As the m -dimensional vector space V satisfies $V \in X_\theta$, the map $P_V \circ H_\theta$ is injective. Moreover,

$$\|(P_V - P_{V_0}) \circ H_\theta\|_{\mathcal{C}(\mathcal{C})} \leq \|P_V - P_{V_0}\| \cdot \|H_\theta\|_{\mathcal{C}(\mathcal{C})} < \frac{1}{4} \varepsilon_1.$$

Moreover, as $\|P_V - P_{V_0}\| < 1$, it is show in Kato (1995), Section I.4.6, that there is orthogonal matrix $R_V \in O(D)$ maps the subspace V onto the subspace V_0 . The matrix R_V is obtained by considering the matrices P_V and P_{V_0} as linear operators in the complex vector space \mathbb{C}^D . Then P_V and P_{V_0} are orthogonal projectors in \mathbb{C}^D onto the (complexified) linear subspaces $\hat{V} = \mathbb{C}V$ and $\hat{V}_0 = \mathbb{C}V_0$. Then, following Kato (1995), Section I.4.6, formula (4.38), one defines the operators $A, B : \mathbb{C}^D \rightarrow \mathbb{C}^D$ by setting $A = (P_V - P_{V_0})^2$ and

$$B = \left(\text{Id} - (P_V + P_{V_0})(P_{V_0} - P_V) \right) (\text{Id} - A)^{-1/2}, \quad (\text{Id} - A)^{-1/2} = \sum_{n=0}^{\infty} \binom{-\frac{1}{2}}{n} A^n,$$

where Id is the identity function and the series converges as $\|A\| \leq \varepsilon_0^2 < \frac{1}{4}$. It is shown in Kato (1995), Section I.4.6, formula (4.2) that the map B satisfies $P_{V_0} = B P_V B^{-1}$ and hence $\hat{V}_0 = \text{Ran}(P_{V_0}) = B(\text{Ran}(P_V)) = B(\hat{V})$. Moreover, by Kato (1995), Section I.6.8, it holds that $B^{-1} = B^*$. As $\|A\| \leq \varepsilon_0^2$, we see that

$$\|\text{Id} - B\| \leq 2\varepsilon_0 \left(\sum_{n=0}^{\infty} \binom{-\frac{1}{2}}{n} \|A\|^n \right) + \left(\sum_{n=1}^{\infty} \binom{-\frac{1}{2}}{n} \|A\|^n \right) \leq 2\varepsilon_0 \frac{1}{1 - \varepsilon_0^2} + \frac{\varepsilon_0^2}{1 - \varepsilon_0^2} \leq 4\varepsilon_0.$$

Let $R_V : \mathbb{R}^D \rightarrow \mathbb{R}^D$ be the restriction of B in \mathbb{R}^D . Then the operator $R_V \in O(D)$ and R_V maps the subspace $V = \mathbb{R}^D \cap \hat{V}$ onto the subspace $V_0 = \mathbb{R}^D \cap \hat{V}_0$. Moreover, it satisfies

$$\|R_V - \text{Id}\|_{\mathbb{R}^D \rightarrow \mathbb{R}^D} = \|B - \text{Id}\|_{\mathbb{C}^D \rightarrow \mathbb{C}^D} < \frac{1}{4(1 + \alpha + \|F_\theta\|_{\mathcal{C}(\mathcal{C})})} \varepsilon_1.$$

Let $\pi_0 : \mathbb{R}^D \rightarrow \mathbb{R}^m$ be the map $\pi_0(y', y'') = y''$ be the projection to the last m coordinates. Then

$$\|R_V \circ P_V \circ H_\theta - P_V \circ H_\theta\|_{C(\mathcal{C})} < \frac{1}{4}\varepsilon_1, \quad \pi_0 \circ H_\theta = F_\theta$$

This and (24) imply that the result of the first half of Theorem 15 applies to the neural network $N_\theta = \pi_0 \circ R_V \circ P_V \circ H_\theta : \mathbb{R}^n \rightarrow \mathbb{R}^m$. \blacksquare

Now we prove the latter half of Theorem 15, the construction of the N'_θ with the prescribed inverse Lipschitz constant.

Lemma 30 *Assume that $D_1 \geq 2n + 2$ and $n \geq 1$. Let $v_1 \in S^{D_1-1}$, $\mathbb{R}^n = \bigcup_{j=1}^J \overline{B}_j$, where B_j are open disjoint sets and let $H_\theta : \mathbb{R}^n \rightarrow \mathbb{R}^{D_1}$ be a continuous map whose restriction to the sets B_j are affine maps, that is $H_\theta|_{B_j}(x) = W_j x + b_j$. Further, let $\delta, \rho \in \mathbb{R}$ be such that $0 < 8\rho \leq \delta < \pi$, then there exists $c_0 \in (0, 1)$ such that if*

$$c_0 \frac{1}{J^{2/(D_1-(2n+1))}} \delta \leq \rho$$

then there exists $w \in \mathbb{S}^{D_1-1}$ such that

$$\text{dist}(w, v_1) < \delta, \quad \text{dist}(w, s_\theta(\mathcal{P})) \geq \rho,$$

where the distances are measured in \mathbb{R}^{D_1} . In other words, there exists a w that is a slight perturbation of v and is at least ρ away from $s_\theta(\mathcal{P})$.

Proof First, we describe our proof technique in words. We show Lemma 30 by using sphere packing arguments (Schneider, 2015) by bounding the size of various the manifolds L_{jk} (defined in Equation 26), which denote directions that we need to stay away from. We bound the size above by considering overlapping spheres, and then below by non-overlapping spheres of half of the radius. By showing that the size of $s_\theta(\mathcal{P}) \subset \bigcup_{jk} L_{jk}$ is ‘small enough,’ there is always room to find a w suitable for the lemma that avoids being too close to $s_\theta(\mathcal{P})$.

We write

$$s_\theta(x, y) = P(r_\theta(x, y))$$

where $P(z) = z / \|z\|_{D_1}$, $P : \mathbb{R}^{D_1} \rightarrow \mathbb{S}^{D_1-1}$ and

$$r_\theta : \mathcal{P} \rightarrow \mathbb{R}^{D_1}, \quad r_\theta(x, y) = H_\theta(x) - H_\theta(y).$$

Observe that as H_θ is a piecewise affine map, so is also r_θ . By our assumptions,

$$\mathcal{P} = \bigcup_{j,k=1}^J (B_j \times B_k) \cap \mathcal{P}$$

and r_θ is affine in all sets $(B_j \times B_k) \cap \mathcal{P}$. Let \mathcal{D}_{jk} be linear space that is spanned by the set

$$\{W_j y - W_k y' + b_j - b_k : y, y' \in \mathbb{R}^n, t \in \mathbb{R}\}, \quad j, k = 1, 2, \dots, J.$$

Note that the dimension \mathcal{D}_{jk} is at most $2n + 1$. To simplify the notations below, we consider linear subspaces $T_{jk} \subset \mathbb{R}^{D_1}$ that have dimension $2n + 1$ and satisfy $\mathcal{D}_{jk} \subset T_{jk}$. Then for all j and k the sets $r_\theta((B_j \times B_k) \cap \mathcal{P})$ and $s_\theta((B_j \times B_k) \cap \mathcal{P}) = P(r_\theta((B_j \times B_k) \cap \mathcal{P}))$ are subsets of T_{jk} . Let

$$L_{jk} = T_{jk} \cap \mathbb{S}^{D_1-1}. \quad (26)$$

Note that $s_\theta(\mathcal{P}) \subset \bigcup_{jk} L_{jk}$. Observe also that as L_{jk} is the intersection of a linear subspace $T_{jk} \subset \mathbb{R}^{D_1}$ of dimension $2n + 1$ and the unit sphere, we have that the dimension of L_{jk} is $2n + 1 - 1 = 2n$. We recall that $D_1 \geq 2n + 2$. Then dimension of L_{jk} is $2n < D_1 - 1$.

Next we estimate the maximal distance of point, in a metric δ -ball of the sphere \mathbb{S}^{D_1-1} , to the union of the sets L_{jk} .

Let $\tilde{\rho} = \frac{1}{8}\rho$ and $B_\delta = B_{\mathbb{S}^{D_1-1}}(v_1, \delta)$ be a metric ball of radius δ centred at v_1 on the sphere $\mathbb{S}^{D_1-1} \subset \mathbb{R}^{D_1}$. Let $\xi_i, i = 1, 2, \dots, K$ be a maximal $\tilde{\rho}$ -separated subset of B_δ . Then we see that $\tilde{\rho}$ -balls $B_{\mathbb{S}^{D_1-1}}(\xi_i, \tilde{\rho})$ cover the set B_δ , that is

$$B_\delta \subset \bigcup_{i=1}^K B_{\mathbb{S}^{D_1-1}}(\xi_i, \tilde{\rho}).$$

Thus the sum of the $(D_1 - 1)$ -volumes of the balls $B_{\mathbb{S}^{D_1-1}}(\xi_i, \tilde{\rho})$ is larger or equal to the volume of the ball B_δ . Using the comparison estimates for manifolds with a bounded sectional curvature (Ch. 6, Cor. 2.4 Petersen et al., 2006), we see that if $c_3 = \sin(1)$ then $c_3^{D_1-1} \delta^{D_1-1} \leq K \tilde{\rho}^{D_1-1}$, that is

$$K \geq c_3^{D_1-1} \left(\frac{\delta}{\tilde{\rho}} \right)^{D_1-1}. \quad (27)$$

In other words, a maximal $\tilde{\rho}$ -separated set in B_δ contains at least $c_3^{D_1-1} (\delta/\tilde{\rho})^{D_1-1}$ points.

On the other hand, we see that $\tilde{\rho}/2$ -balls $B_{\mathbb{S}^{D_1-1}}(\xi_i, \tilde{\rho}/2)$ are disjoint subsets of the ball $B_{\mathbb{S}^{D_1-1}}(v_1, 2\delta)$ of radius 2δ . Thus the sum of volumes of the balls $B_{\mathbb{S}^{D_1-1}}(\xi_i, \tilde{\rho}/2)$ is smaller or equal to the volume of $B_{\mathbb{S}^{D_1-1}}(v_1, 2\delta)$. Again by using (Ch. 6, Cor. 2.4 Petersen et al., 2006), we see that, when we denote $c_4 = 4c_3^{-1} \geq 1$, we have

$$K \leq c_4^{D_1-1} \left(\frac{\delta}{\tilde{\rho}} \right)^{D_1-1}. \quad (28)$$

In other words, a maximal $\tilde{\rho}$ -separated set in B_δ contains at most $c_4^{D_1-1} (\delta/\tilde{\rho})^{D_1-1}$ points.

Next, for any $j, k = 1, 2, \dots, J$, let $\{x_p^{jk} : p = 1, \dots, K_{jk}\}$, consisting of K_{jk} points, be a maximal $\tilde{\rho}/2$ separated subset set in L_{jk} . Recall that $L_{jk} = T_{jk} \cap \mathbb{S}^{D_1-1}$ and that T_{jk} are linear spaces of dimension $2n + 1$. Thus L_{jk} are isometric to the $2n$ dimensional sphere \mathbb{S}^{2n} . By the above considerations in (28), we have that

$$K_{jk} \leq c_4^{2n} (2\delta/\tilde{\rho})^{2n}, \quad j, k = 1, 2, \dots, J. \quad (29)$$

Next we show that if

$$\tilde{\rho} < \frac{1}{8} c_0 \frac{1}{J^{2/(D_1-(2n+1))}} \delta,$$

then the $\tilde{\rho}/2$ -neighborhood of the set $\bigcup_{j,k=1}^J L_{jk}$ does not contain the set B_δ . To show this, assume the opposite that the $\tilde{\rho}/2$ -neighborhood of the set $\bigcup_{j,k=1}^J L_{jk}$ contains the set B_δ . Then the $\tilde{\rho}$ -neighborhood of the set $Y = \bigcup_{j,k=1}^J \{x_p^{jk} : p = 1, \dots, K_{jk}\}$ contains the set B_δ , and hence the set Y contains a maximal $2\tilde{\rho}$ -separated subset of B_δ . Let Y_1 denote this set. Let K be the number of the points in the set Y_1 .

From Equation 27 we have $c_3^{D_1-1}(\delta/\tilde{\rho})^{D_1-1} \leq K$. From $Y_1 \subset Y$ we have $K \leq \sum_{j,k=1}^J K_{jk}$. From Equation 29 we have $\sum_{j,k=1}^J K_{jk} \leq J^2 c_4^{2n} (2\delta/\tilde{\rho})^{2n}$, this in total we have

$$c_3^{D_1-1}(\delta/\tilde{\rho})^{D_1-1} \leq J^2 c_4^{2n} (2\delta/\tilde{\rho})^{2n}$$

or equivalently,

$$c_3^{D_1}(\delta/\tilde{\rho})^{D_1-(2n+1)} \leq J^2 (2c_4)^{2n},$$

or equivalently,

$$\tilde{\rho} \geq \frac{c_3^{(D_1-1)/(D_1-(2n+1))}}{(2c_4)^{2n/(D_1-(2n+1))}} \frac{1}{J^{2/(D_1-(2n+1))}} \delta.$$

And so

$$\tilde{\rho} \geq c_0 \frac{1}{J^{2/(D_1-(2n+1))}} \delta, \quad c_0 = \frac{c_3^{(D_1-1)/(D_1-(2n+1))}}{(2c_4)^{2n/(D_1-(2n+1))}}.$$

Summarizing, the above shows that if

$$\rho < c_0 \frac{1}{J^{2/(D_1-(2n+1))}} \delta,$$

so that also $\tilde{\rho} \leq \rho < c_0 \frac{1}{J^{2/(D_1-(2n+1))}} \delta$, then there exists $w \in B_\delta$ such that w is not in the $\tilde{\rho}/2$ -neighborhood of the set $\bigcup_{j,k=1}^J L_{jk}$ in \mathbb{S}^{D_1-1} . As $\bigcup_{j,k=1}^J L_{jk}$ contains the sets $s_\theta(\mathcal{P})$ and for $w_1, w_2 \in B_\delta \subset \mathbb{S}^{D_1-1} \subset \mathbb{R}^{D_1}$,

$$\|w_1 - w_2\|_{\mathbb{R}^{D_1}} \leq \text{dist}_{\mathbb{S}^{D_1-1}}(w_1, w_2) \leq 4\|w_1 - w_2\|_{\mathbb{R}^{D_1}},$$

the claim follows. ■

Next we apply that above lemma to obtain inverse Lipschitz estimates for the map $Q_w \circ H_\theta$.

Lemma 31 *Let $H_\theta : \mathbb{R}^n \rightarrow \mathbb{R}^D$ satisfy $\|H_\theta(x) - H_\theta(y)\|_{2,D} \geq \alpha \|x - y\|_{2,n}$ and assume that $w \in \mathbb{S}^{D-1}$ satisfies $\text{dist}(w, s_\theta(\mathcal{P})) \geq \rho$. Then*

$$\|Q_w H_\theta(x) - Q_w H_\theta(y)\|_{2,D} \geq \frac{1}{\sqrt{2}} \rho \alpha \|x - y\|_{2,n}.$$

Proof Assume that $w \in \mathbb{S}^{D-1}$ satisfies

$$\text{dist}_{\mathbb{R}^D}(w, s_\theta(\mathcal{P})) \geq \rho.$$

Let $\beta \geq \frac{1}{\sqrt{2}}\rho\alpha$. Let $x, y \in \mathbb{R}^n$, $x \neq y$ and

$$v = Q_w H_\theta(x) - Q_w H_\theta(y) \in \mathbb{R}^D.$$

Note that $\|v\|_{2,D} < \beta \|x - y\|_{2,n}$. Then, we see that there is $t \in \mathbb{R}$

$$H_\theta(x) - H_\theta(y) = tw + v.$$

Let $r = \|H_\theta(x) - H_\theta(y)\|_{2,D} \neq 0$. Again, by changing roles of x and y , if needed, we can assume that $t \geq 0$. Note that $w \perp v$.

The above yields

$$\frac{t}{r}w = (H_\theta(x) - H_\theta(y) - v)/r = s_\theta(x, y) - \frac{v}{r}.$$

Then,

$$\frac{t}{r}w + \frac{v}{r} = s := s_\theta(x, y) \in s_\theta(\mathcal{P}).$$

Assume first that $\langle s, w \rangle \geq 0$. Then, as w and s are unit vectors, we have for any $b \in \mathbb{R}$

$$\|s - bw\|_{2,D}^2 \geq \|s - \langle s, w \rangle w\|_{2,D}^2 = \|s\|_{2,D}^2 - 2\langle s, \langle s, w \rangle w \rangle + \langle s, w \rangle^2 = 1 - \langle s, w \rangle^2$$

and as $0 \leq \langle s, w \rangle \leq 1$,

$$\|s - w\|_{2,D}^2 = \|s\|_{2,D}^2 - 2\langle s, w \rangle + \|w\|_{2,D}^2 = 2(1 - \langle s, w \rangle) \leq 2(1 - \langle s, w \rangle^2)$$

and hence $\|s - w\|_{2,D}^2 \leq 2 \inf_{b \in \mathbb{R}} \|s - bw\|_{2,D}^2$. Thus,

$$\|w - s\|_{2,D} \leq \sqrt{2} \inf_{b \in \mathbb{R}} \|s - bw\|_{2,D} \leq \sqrt{2} \left\| s - \frac{t}{r}w \right\|_{2,D} = \sqrt{2} \left\| \frac{v}{r} \right\|_{2,D}.$$

Hence

$$\begin{aligned} \rho &\leq \text{dist}(w, s_\theta(\mathcal{P})) \\ &\leq \sqrt{2} \left\| \frac{v}{r} \right\|_{2,D} \\ &\leq \frac{\sqrt{2}}{r} \beta \|x - y\|_{2,n}. \end{aligned} \tag{30}$$

Recall that by our assumptions in the claim

$$\|H_\theta(x) - H_\theta(y)\|_{2,D} \geq \alpha \|x - y\|_{2,n}$$

Hence, by (30),

$$\rho < \frac{\sqrt{2}}{r} \beta \|x - y\|_{2,n} \leq \frac{\sqrt{2}}{\alpha \|x - y\|_{2,n}} \beta \|x - y\|_{2,n} = \frac{\sqrt{2}\beta}{\alpha}.$$

This shows that for all $x, y \in \mathbb{R}^n$

$$\|Q_w H_\theta(x) - Q_w H_\theta(y)\|_{2,D} \geq \frac{1}{\sqrt{2}} \rho \alpha \|x - y\|_{2,n}.$$

■

Next we apply Lemma 31 for the neural networks $H_\theta^{(0)} = H_\theta$ and

$$H_\theta^{(j)} = Q_{w_j} \circ \dots \circ Q_{w_1} H_\theta : \mathbb{R}^n \rightarrow \mathbb{R}^D, \quad j = 1, 2, \dots, n, \quad D = n + m,$$

with suitably chosen orthogonal vectors w_j . We choose these vectors using a recurrent procedure: Let v_1, \dots, v_{n+m} be the unit coordinate vectors in \mathbb{R}^{n+m} . Our initial choice of w_1 is arbitrary. Once the first w_1, w_2, \dots, w_{k-1} are chosen, we choose $w_k \in \{w_1, w_2, \dots, w_{k-1}\}^\perp \subset \mathbb{R}^D$. Next, we use the isomorphism

$$A_k : \{w_1, w_2, \dots, w_{k-1}\}^\perp \rightarrow \mathbb{R}^{D-(k-1)}, \quad A_1 = I : \mathbb{R}^D \rightarrow \mathbb{R}^D$$

where $\{w_1, w_2, \dots, w_{k-1}\}^\perp \subset \mathbb{R}^D$. The map A_k is used to identify the linear space $\{w_1, w_2, \dots, w_{k-1}\}^\perp \subset \mathbb{R}^D$ with a $D - (k - 1)$ dimensional Euclidean space, and it can be also considered as a map that gives coordinates in the space $\{w_1, w_2, \dots, w_{k-1}\}^\perp$. By applying Lemma 30 for the map $A_k \circ H_\theta^{(k-1)} : \mathbb{R}^n \rightarrow \mathbb{R}^{D-(k-1)}$, that is a map from n dimensional space to an Euclidean space of dimension $d_{k-1} = D - (k - 1) \geq n + 2n + 1 - (n - 1) \geq 2n + 2$. In this, we use the values

$$\rho_k = \frac{c_0}{2J^{2/((D-(k-1))-(2n+1))}} \delta = \frac{c_0}{2J^{2/(m-n-k)}} \delta$$

and

$$\alpha_k = \left(\prod_{j=1}^k \frac{1}{\sqrt{2}} \rho_j \right) \alpha,$$

so that $\alpha_k = (\frac{1}{\sqrt{2}} \rho_k) \alpha_{k-1}$ and $\alpha_0 = \alpha$.

By replacing the vector v_1 by $A_k v_k$, ρ by ρ_k , and α by α_{k+1} , respectively, Lemma 30 implies that there is a vector $w \in \mathbb{S}^{d_{k-1}-1}$ such that

$$\text{dist}(w, A_k v_k) < \delta, \quad \text{dist}(w, s_\theta(\mathcal{P})) \geq \rho_k, \quad (31)$$

where distances are measured in $\mathbb{R}^{d_{k-1}}$ and $s_\theta(\mathcal{P})$ is defined in formula (23) with H_θ replaced by the neural network $A_k \circ H_\theta^{(k-1)} : \mathbb{R}^n \rightarrow \mathbb{R}^{d_{k-1}}$. We define $w_k = A_k^{-1} w$. Then by Lemma 31, $H_\theta^{(k)} = Q_{w_k} \circ H_\theta^{(k-1)}$ satisfies

$$\left\| H_\theta^{(k)}(x) - H_\theta^{(k)}(y) \right\|_{2,D} \geq \alpha_k \|x - y\|_{2,n}. \quad (32)$$

Observe that as Q_{w_j} are linear maps, the restrictions of all maps $H_\theta^{(j)}$ are affine in the the same sets B_1, \dots, B_J . Then, apply map the vectors $v_{n+m}, \dots, v_{n+1}, w_n, \dots, w_1$ to the vectors v_{n+m}, \dots, v_1 . The above yields that

$$\left\| H_\theta^{(n)}(x) - H_\theta^{(n)}(y) \right\|_{2,D} \geq a \|x - y\|_{2,n}, \quad a = \alpha_n, \quad C_0 = \frac{1}{2\sqrt{2}} c_0.$$

Moreover, we have

$$\|H_\theta(x)\|_{2,D} \leq (C_R + \alpha R), \quad C_R = \|F_\theta\|_{C(B(R))},$$

and as $\|Q_{w_j}\| \leq 1$,

$$\left\| H_\theta^{(j)}(x) \right\|_{2,D} \leq (C_R + \alpha R).$$

Moreover, for $j = 1, 2, \dots, n$ we have

$$\begin{aligned} \left\| H_\theta^{(j)}(x) - Q_{v_j} H_\theta^{(j-1)}(x) \right\|_{2,D} &\leq \|Q_{w_j} - Q_{v_j}\| \left\| H_\theta^{(j-1)}(x) \right\|_{2,D} \\ &\leq 2 \|w_j - v_j\|_{2,D} \left\| H_\theta^{(j-1)}(x) \right\|_{2,D}. \end{aligned}$$

Let $R_{w_n, \dots, w_1, v_n, \dots, v_1} \in \mathbb{R}^{(n+m) \times (n+m)}$ be an invertable matrix that maps the vectors $v_{n+m}, v_{n+m-1}, \dots, v_{n+1}, w_n, \dots, w_1$ to $v_{n+m}, v_{n+m-1}, \dots, v_{n+1}, v_n, \dots, v_1$, respectively. We recall that $\|w_j - v_j\|_{2,n} < \delta$. For any $x \in \mathbb{R}^{m+n}$, let $x = \sum_{i=1}^n \alpha_i v_i + \sum_{i=n+1}^{m+n} \alpha_i v_i$, then

$$\begin{aligned} &\left\| (R_{w_n, \dots, w_1, v_n, \dots, v_1}^{-1} - I) x \right\|_{2,D} \\ &= \left\| \sum_{i=1}^n \alpha_i (R_{w_n, \dots, w_1, v_n, \dots, v_1}^{-1} v_i - v_i) + \sum_{i=n+1}^{m+n} \alpha_i (R_{w_n, \dots, w_1, v_n, \dots, v_1}^{-1} v_i - v_i) \right\|_{2,D} \\ &= \left\| \sum_{i=1}^n \alpha_i (w_i - v_i) \right\|_{2,D} \leq \sum_{i=1}^n \|\alpha_i (v_i - w_i)\|_{2,D} \\ &< \sum_{i=1}^n \alpha_i \delta \leq \delta \|x\|_{l^1(\mathbb{R}^{m+n})} \\ &\leq \sqrt{m+n} \delta \|x\|_{2,D} \end{aligned}$$

where $\|x\|_{l^1(\mathbb{R}^{m+n})}$ denotes the discrete L_1 norm of x in \mathbb{R}^{m+n} . Further

$$\begin{aligned} \|R_{w_n, \dots, w_1, v_n, \dots, v_1} - I\|_{2,D} &= \left\| (I - R_{w_n, \dots, w_1, v_n, \dots, v_1}^{-1}) R_{w_n, \dots, w_1, v_n, \dots, v_1} \right\|_{2,D} \\ &\leq \|R_{w_n, \dots, w_1, v_n, \dots, v_1}^{-1} - I\|_{2,D} \|R_{w_n, \dots, w_1, v_n, \dots, v_1}\|_{2,D} \\ &\leq \frac{\sqrt{m+n}}{1 - \sqrt{m+n}\delta} \delta \end{aligned}$$

where the last inequality comes from $\|R_{w_n, \dots, w_1, v_n, \dots, v_1}\|_{2,D} = \left\| (I + (R_{w_n, \dots, w_1, v_n, \dots, v_1}^{-1} - I))^{-1} \right\|_{2,D}$. Since $Q_{v_n} \dots Q_{v_1} H_\theta(x) = \tilde{F}_\theta(x) := (F_\theta(x), 0, \dots, 0)$ we see that

$$\left\| \tilde{F}_\theta(x) - H_\theta^{(n)}(x) \right\|_{2,m} \leq 2n\delta(C_R + \alpha R).$$

Then,

$$f_\theta(x) = P_0 R_{w_n, \dots, w_1, v_n, \dots, v_1} H_\theta^{(n)}(x) \tag{33}$$

where $P_0 : \mathbb{R}^m \times \mathbb{R}^n \rightarrow \mathbb{R}^m$ maps $P_0(x, z) = x$. The above yields

$$\begin{aligned} \|f_\theta(x) - F_\theta(x)\|_{2,m} &\leq 2n\delta(C_R + \alpha R) + \frac{\sqrt{m+n}}{1 - \sqrt{m+n}\delta} \delta(C_R + \alpha R) \\ &= \left(2n + \frac{\sqrt{m+n}}{1 - \sqrt{m+n}\delta}\right) \delta(C_R + \alpha R) \end{aligned}$$

for $x \in B(R)$.

Further, from $\|R\|_{2,D} \geq \|I + (R^{-1} - I)\|_{2,D}$, the fact that $R_{w_n, \dots, w_1, v_n, \dots, v_1} H_\theta^{(n)}(x)$ belongs in the image of P_0 , and from (32) & (33) we have

$$\|f_\theta(x) - f_\theta(y)\|_{2,m} \geq a \|x - y\|_{2,n}, \quad a = \frac{\alpha_n}{1 + \sqrt{m+n}\delta} \geq \alpha_n.$$

Above, using the definition Euler's constant and the harmonic series, we see that there are $c_6, c_5 \in \mathbb{R}_+$ such that

$$\ln(m) + c_6 \leq \sum_{j=1}^m \frac{1}{j} \leq \ln(m) + c_7$$

and as $\ln(m+1) + c_6 - \ln 2 \leq \ln(2m) + c_6 - \ln 2 \leq \ln(m) + c_6$ and $\ln(m) + c_7 \leq \ln(m+1) + c_7$, we see that for $C_2 = 2(c_7 + c_6 + \ln 2)$

$$\sum_{j=1}^n \frac{1}{m-n-j} = \sum_{j=1}^{m-n-1} \frac{1}{j} - \sum_{j=1}^{m-2n-1} \frac{1}{j} \leq \ln(m-n) - \ln(m-2n) + C_2/2$$

and

$$a = \frac{\alpha}{1 + \sqrt{m+n}\delta} C_0^{nD} \delta^n J^{-2(\sum_{j=1}^n \frac{1}{m-n-j})} \geq \alpha C_0^{n(n+m)} \delta^n J^{-2 \ln((m-n)/(m-2n)) - C_2}.$$

We choose

$$\varepsilon_2 = \left(2n + \frac{\sqrt{m+n}}{1 - \sqrt{m+n}\delta}\right) \delta(C_R + \alpha R)$$

so that $\delta = \left(\left(2n + \frac{\sqrt{m+n}}{1 - \sqrt{m+n}\delta}\right) (C_R + \alpha R)\right)^{-1} \varepsilon_2$ and

$$\begin{aligned} &\left(\frac{C_0^D}{(2n + \sqrt{m+n})(C_R + \alpha R)} \varepsilon_2\right)^n J^{-2 \ln((m-n)/(m-2n)) - C_2} \alpha \\ &\leq \left(\frac{C_0^D}{\left(2n + \frac{\sqrt{m+n}}{1 - \sqrt{m+n}\delta}\right) (C_R + \alpha R)} \varepsilon_2\right)^n J^{-2 \ln((m-n)/(m-2n)) - C_2} \alpha \\ &\leq a. \end{aligned}$$

When we use above $\alpha = 1$ and $C_1 = C_0/10$, this yields the final part of Theorem 15.

Lemma 29 used in the above proof yields also Corollary 16.

Proof [Proof of Corollary 16] Observe that a measure μ that is absolutely continuous with respect to the Lebesgue $\prod_{j=1}^k \mathbb{R}^{d_{2j} \times d_{2j-1}}$ is absolutely continuous also with respect to the normalized Gaussian distribution, and that if the set $S = \{F_k \text{ is not injective}\}$ has measure zero with respect to the normalized Gaussian distribution, then its μ -measure is also zero. Thus we can assume without loss of generality that the elements of matrices B_j are independent and have normalized Gaussian distributions.

Assume next that we have shown that $F_{j-1} : \mathbb{R}^n \rightarrow \mathbb{R}^{d_{2j-2}}$ is injective almost surely. Then, $f_\theta^{(j)} \circ F_{j-1} : \mathbb{R}^n \rightarrow \mathbb{R}^{d_{2j-1}}$ is injective almost surely. If $d_{2j} > d_{2j-1}$, the matrix B_j is almost surely injective and so is $F_j = B_j \circ f_\theta^{(j)} \circ F_{j-1} : \mathbb{R}^n \rightarrow \mathbb{R}^{d_{2j}}$. Thus, it is enough to consider the case when $d_{2j-1} \geq d_{2j} \geq 2n + 1$. Then the matrix $B_j : \mathbb{R}^{d_{2j-1}} \rightarrow \mathbb{R}^{d_{2j}}$ has almost surely rank d_{2j} . By using the singular value decomposition, we can write $B_j = R_j^1 D_j R_j^2$ where $R_j^1 \in O(d_{2j})$, $R_j^2 \in O(d_{2j-1})$, $D_j \in \mathbb{R}^{d_{2j} \times d_{2j-1}}$ is matrix which principal diagonal elements are almost surely strictly positive and the other elements are zeros. Let $V_j = (R_j^2)^{-1}(\mathbb{R}^{d_{2j}} \times \{(0, 0, \dots, 0)\}) \subset \mathbb{R}^{d_{2j-1}}$. Let P_{V_j} be an orthogonal projector in $\mathbb{R}^{d_{2j-1}}$ onto the space V_j of dimension d_{2j} . As the distribution of the matrix B_j is invariant in rotations of the space, so is the distribution of linear space V_j in $G(d_{2j-1}, d_{2j})$. By Lemma 29, we see that the map $P_{V_j} \circ f_\theta^{(j)} \circ F_{j-1} : \mathbb{R}^n \rightarrow \mathbb{R}^{d_{2j-1}}$ is injective almost surely. This implies that the map $B_j \circ f_\theta^{(j)} \circ F_{j-1} : \mathbb{R}^n \rightarrow \mathbb{R}^{d_{2j}}$ is injective almost surely, that is, the map $F_j : \mathbb{R}^n \rightarrow \mathbb{R}^{d_{2j}}$ is injective almost surely. The claim follows by induction. \blacksquare

In Baraniuk and Wakin (2009), see also Hegde et al. (2008); Iwen and Maggioni (2013), and Broomhead and Kirby (2001, 2000), the authors study manifold learning using random projectors. These results are related to the proof of the first, universal, half of Theorem 15 above. Let H_θ be given by (25), a ReLU-based neural network whose graph $M \subset \mathbb{R}^d$, $d = 2n + m$. When P_V is a random projector in \mathbb{R}^d onto a m -dimensional linear subspace V , the injectivity of the neural network $P_V \circ H_\theta$ is closely related to the property that $P_V(M)$ is an n -dimensional submanifold with a large probability. In Broomhead and Kirby (2001, 2000) the authors use Whitney embedding results for C^2 -smooth manifold for dimension reduction of data. Our proof applies similar techniques for Lipschitz-smooth maps. In (Baraniuk and Wakin, 2009), (Hegde et al., 2008), and (Iwen and Maggioni, 2013), the authors apply the result that when $M \subset \mathbb{R}^D$ is a submanifold and D is large enough, a random m -dimensional projector P_V satisfies on M the restricted isometry property with a large probability. In this case, $P_V \circ H_\theta$ is not only an injection but its inverse map is also a local Lipschitz map. In this sense, the techniques in Baraniuk and Wakin (2009), Hegde et al. (2008), and Iwen and Maggioni (2013) would give improved results to the generic projection technique used in this paper. The results in Baraniuk and Wakin (2009), Hegde et al. (2008), and Iwen and Maggioni (2013), however, require that the dimension m of image space of the map $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$ satisfies $m \geq C \log(\varepsilon^{-1})$, where ε is the precision parameter in the inequality (6). Our result, Theorem 15 requires only that $m \geq 2n + 1$.

Appendix D. Miscellaneous Lemmas

Lemma 32 *Let $x, y, z \in \mathbb{R}$ and $z \geq 0$. Then*

$$\text{ReLU}(x + z) = \text{ReLU}(y + z) \implies \text{ReLU}(x) = \text{ReLU}(y).$$

Lemma 33 (Useful Inequalities) *The following geometric inequalities are useful and have straight forward proofs. Suppose that $a, b, c \in \mathbb{R}^n$, then*

1. *For $i = 1, \dots, k$, let $a_i \in \mathbb{R}^m$. If*

$$\sum_{i=1}^k \|a_i\|_2^2 \leq \left\| \sum_{i=1}^k a_i \right\|_2^2,$$

then if for each $j = 1, \dots, m$, $a_i|_j \cdot a_{i'}|_j \geq 0$ for each pair $i, i' \in [[k]]$, then

$$\sum_{i=1}^k \|a_i\|_2 \leq \sqrt{k} \left\| \sum_{i=1}^k a_i \right\|_2. \quad (34)$$

2. *If $a, b \in \mathbb{R}^n$ and $\langle a, b \rangle \geq 0$, then*

$$\|a\|_2^2 + \|b\|_2^2 \leq \|a + b\|_2^2. \quad (35)$$

Lemma 34 (Co-linear Additivity of $\text{ReLU}(W \cdot)$) *Let $W \in \mathbb{R}^{m \times n}$, $x_1, x_2 \in \mathbb{R}^n$, $\ell^{x_1, x_2}(t) = (1 - t)x_1 + tx_2$. Let there be*

$$0 = t_1 \leq t_2 \leq \dots \leq t_{n_t-1} \leq t_{n_t} = 1$$

then

$$\sum_{k=1}^{n_t-1} \|\text{ReLU}(Wx_{t_k}) - \text{ReLU}(Wx_{t_{k+1}})\|_2^2 \leq \|\text{ReLU}(Wx_1) - \text{ReLU}(Wx_2)\|_2^2. \quad (36)$$

and

$$\sum_{k=1}^{n_t-1} \|\text{ReLU}(Wx_{t_k}) - \text{ReLU}(Wx_{t_{k+1}})\|_2 \leq \sqrt{n_t} \|\text{ReLU}(Wx_1) - \text{ReLU}(Wx_2)\|_2. \quad (37)$$

Proof Let $x_t = \ell^{x_1, x_2}(t)$, then as a function of t , the j 'th component of $\text{ReLU}(Wx_t)$ is either increasing (if $\langle w_j, x_2 - x_1 \rangle \geq 0$) or decreasing (if $\langle w_j, x_2 - x_1 \rangle \leq 0$). In either case, it is clear that for each $j = 1, \dots, m$

$$(\text{ReLU}(Wx_1) - \text{ReLU}(Wx_2))|_j \cdot (\text{ReLU}(Wx_2) - \text{ReLU}(Wx_3))|_j \geq 0$$

hence clearly

$$\langle (\text{ReLU}(Wx_1) - \text{ReLU}(Wx_2)), (\text{ReLU}(Wx_2) - \text{ReLU}(Wx_3)) \rangle \geq 0$$

thus we can apply (35) and we obtain (36). Applying (34) then yields (37). ■

Appendix E. Detailed Comparison to Prior Work

In this section we detail how our work is related to the works of Bruna et al. (2013) and Mallat et al. (2018)

E.1 Comparison to Bruna et al. (2013)

In Bruna et al. (Proposition 2.2. 2013) the authors give a result invoking a condition similar to our DSS condition (Definition 1). It also concerns injectivity of a ReLU layer in terms of the injectivity of the weight matrix restricted to certain rows. The authors also compute a bi-Lipschitz bound for a layer (similar to our Theorem 8), though as we show in the following examples their analysis is in some cases not precisely aligned with injectivity.

Their criterion is given in two parts. For a weight matrix, they first define a notion of admissible set which indicates the points where the weight matrix's injectivity must be tested. Injectivity follows provided that the weight matrix is non-singular when restricted to each admissible set. Given a weight matrix $W \in \mathbb{R}^{M \times N}$ and bias b , the authors say that $\Omega \subset \{1, \dots, M\}$ is admissible if

$$\bigcap_{i \in \Omega} \{x: \langle x, w_i \rangle > b_i\} \cap \bigcap_{i \notin \Omega} \{x: \langle x, w_i \rangle < b_i\} \quad (38)$$

is not empty. For our analysis we focus on the case when $b \equiv 0$. In this case Ω is admissible if and only if

$$\exists x \in \mathbb{R}^n \quad \text{such that} \quad \langle x, w_i \rangle \begin{cases} > 0 & \text{if } i \in \Omega \\ < 0 & \text{if } i \notin \Omega \end{cases} \quad (39)$$

Note that the inequality in (39) is strict, unlike (1). If, for example, W has a column that is the zero vector, then there are no admissible Ω . The authors use the notation $\overline{\Omega}$ to denote all admissible sets for a given weight matrix. In their notation F is the transpose of our weight matrix W , F_Ω are the Ω rows of the weight matrix, $F_\Omega|_{V_\Omega}$ is the subspace generated by the Ω rows of W . The authors also call the ReLU function the half-rectification function. $\lambda_-(F)$ and $\lambda_+(F)$ denote the lower and upper frame bounds of F respectively. The injectivity criterion from Bruna et al. (2013) is

Proposition 35 *Let $A_0 = \min_{\Omega \in \overline{\Omega}} \lambda_-(F_\Omega|_{V_\Omega})$. Then the half-rectification operator $M_b(x) = \text{ReLU}(F^T x + b)$ is injective if and only if $A_0 > 0$. Moreover, it satisfies*

$$\forall x, x', A_0 \|x - x'\| \leq \|M_b(x) - M_b(x')\| \leq B_0 \|x - x'\|$$

with $B_0 = \max_{\Omega \in \overline{\Omega}} \lambda_+(F_\Omega) \leq \lambda_+(F)$.

We now show that Proposition 35 does not precisely align with injectivity of $\text{ReLU}(W(\cdot))$. We construct a weight matrix for which $A_0 > 0$, but does not yield an injective $\text{ReLU}(W(\cdot))$. If

$$W = \begin{bmatrix} 1 & 0 \\ 0 & 1 \\ -1 & 0 \end{bmatrix}$$

then clearly $\text{ReLU}(Wx)$ is not injective (for all $\alpha < 0$, $\text{ReLU}(W \begin{bmatrix} 0 \\ \alpha \end{bmatrix}) = \bar{0}$). The only admissible sets are $\bar{\Omega} = \{\{1\}, \{3\}, \{1, 2\}, \{2, 3\}\}$ (notably $\{2\}$ is not admissible). W is full rank on all $\Omega \in \bar{\Omega}$, so $A_0 > 0$ so Proposition 35 implies that $\text{ReLU}(W(\cdot))$ is injective. Now consider the case when

$$W = \begin{bmatrix} B \\ -DB \\ 0 \end{bmatrix} \quad (40)$$

where B is a basis of \mathbb{R}^n , D is a strictly positive diagonal matrix, and 0 is the zero row vector. From Corollary 6, W satisfies Theorem 2, and so $\text{ReLU}(W(\cdot))$ is injective. On account of the zero row vector in (40), $\forall x \in \mathbb{R}^n$, $\langle x, 0 \rangle = 0$ so there are no Ω that are admissible Ω according to (38). Thus A_0 is undefined.

Now we construct an example of a W and $x, x' \in \mathbb{R}^n$ for which $A_0 \|x - x'\| > \|\text{ReLU}(Wx) - \text{ReLU}(Wx')\|$. Let

$$W = \begin{bmatrix} 1 & 0 \\ 0 & 1 \\ -1 & 0 \\ 0 & -1 \end{bmatrix}, \quad x = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ 1 \end{bmatrix}, \quad x' = \frac{1}{\sqrt{2}} \begin{bmatrix} -1 \\ 1 \end{bmatrix}.$$

Clearly on every admissible set $\lambda_-(W_\Omega|_{V_\Omega}) = 1$, so

$$\begin{aligned} \text{ReLU}(Wx) &= \frac{1}{\sqrt{2}} \text{ReLU} \left(\begin{bmatrix} 1 \\ 1 \\ -1 \\ -1 \end{bmatrix} \right) = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ 1 \\ 0 \\ 0 \end{bmatrix}, \\ \text{ReLU}(Wx') &= \frac{1}{\sqrt{2}} \text{ReLU} \left(\begin{bmatrix} -1 \\ 1 \\ 1 \\ -1 \end{bmatrix} \right) = \frac{1}{\sqrt{2}} \begin{bmatrix} 0 \\ 1 \\ 1 \\ 0 \end{bmatrix}, \end{aligned}$$

hence,

$$\|\text{ReLU}(Wx) - \text{ReLU}(Wx')\| = \frac{1}{\sqrt{2}} \left\| \begin{bmatrix} 1 \\ 0 \\ -1 \\ 0 \end{bmatrix} \right\| = 1$$

and

$$\|x - x'\| = \frac{1}{\sqrt{2}} \left\| \begin{bmatrix} 2 \\ 0 \end{bmatrix} \right\| = \sqrt{2}.$$

From this we have

$$\sqrt{2} = A_0 \|x - x'\| > \|\text{ReLU}(Wx) - \text{ReLU}(Wx')\| = 1.$$

On the other hand, substituting this into (3) yields

$$\frac{C}{\sqrt{m}} \|x - x'\| = \frac{1}{\sqrt{8}} \sqrt{2} = \frac{1}{2} \leq 1,$$

which does hold, suggesting that the lower bound in Proposition 1 is not pessimistic enough.

E.2 Relationship to Mallat et al. (2018)

In Mallat et al. (2018) the authors consider a construction analogous to our convolutional construction (in Definition 11) defined on a continuum (that is infinite-dimensional function defined on an interval) rather than on a vector (that is discrete finite dimensional function) defined on a subset of \mathbb{R}^n . The authors posit that CNNs first learn a layer of filters localized in frequency varied in phase. The authors also show that a ReLU activation function acts as a filter on the phase of the convolution of the filters against the input signal and that, provided that the filters are sufficiently different in phase and satisfy a frame condition then the layer is bi-Lipschitz, and hence is injective. Their analysis a particularized version of ours, and can be straight forwardly subsumed by our work.

The frame condition is given by Proposition 2.6 in Mallat et al. (2018) that the weight matrix must satisfy in order ensure that W is invertible and stable. In the notation of Mallat et al. (2018) the filters $\hat{\psi}_\lambda$ are analogous to the Fourier transform of the kernels in Definition 10, and the condition in (2.25) in Mallat et al. (2018) is one natural way to generalize the notion of a basis to a continuous signal. Hence, Proposition 2.6 in Mallat et al. (2018) can be loosely interpreted as a statement that the kernels in a given layer of width P form a basis of \mathbb{R}^P .

The second condition is that kernels are given in terms of belonging to a family, and members of this family are related to each other in the sense that members of the same family are centered in the same Fourier domain, and act as phase offsets of differing phase. Equation 2.14 of Mallat et al. (2018) describes that a phase filter $H: \mathbb{C} \times [0, 2, \pi] \rightarrow \mathbb{C}$ is defined by

$$\forall z \in \mathbb{C}, \alpha \in [0, 2\pi], \quad Hz(\alpha) = |z|h(\alpha - \varphi(z))$$

where $|z|$ standard modulus of a complex number, $\varphi(z)$ is the complex phase, and $h(\alpha) = \text{ReLU}(\cos(\alpha))$. If we consider just the value of $Hz(0)$ and $Hz(\pi)$, then we find

$$Hz(0) = |z| \text{ReLU}(\cos(\varphi(z))) = \text{ReLU}(|z| \cos(\varphi(z))) = \text{ReLU}(\mathcal{R}z), \quad (41)$$

$$Hz(\pi) = |z| \text{ReLU}(-\cos(\varphi(z))) = \text{ReLU}(-|z| \cos(\varphi(z))) = \text{ReLU}(-\mathcal{R}z), \quad (42)$$

where $\mathcal{R}z$ is the real part of z . If, as in Mallat et al. (2018), z is given by $z = x \star c_k$, where x is a real signal and $\star c_k$ denotes the convolution against a kernel c_k , then (41) and (42) imply that

$$\begin{aligned} Hz(0) &= \text{ReLU}(x \star \mathcal{R}c_k), \\ Hz(\pi) &= \text{ReLU}(x \star (-\mathcal{R}c_k)), \end{aligned}$$

that is, that for every kernel c_k in a layer, the kernel $-c_k$ is also in that layer.-

Combining the two logical conditions above implies that the kernels of width c_k form a basis of \mathbb{R}_P and that for every kernel c_k there is also a kernel $-c_k$. Together these two (by Corollary 6) that the c_k form a DSS of \mathbb{R}^P , and thus by Theorem 13 the entire layer is injective.

Appendix F. Architecture Details for Experiments

Generator network: We train a generator with 5 convolutional layers. The input latent code is 256-dimensional which is treated by the network as a $1 \times 1 \times 256$ size tensor. The first layer is a transposed convolution with a kernel size of 4×4 with stride 1 and 1024 output channels. This is followed by a leaky ReLU. We follow this up by 3 conv layers each of which halve the number of channels and double the image size (that is we go from $N/2 \times N/2 \times C$ to $N \times N \times C/2$ tensor) giving an expansivity of two, the minimum required for injectivity of ReLU networks. Each of these 3 convolution layers has kernel size 3, stride 2 and is followed by the ReLU activation. These layers are made injective by having half the filters as w and the other half as $-s^2w$. Here, w and s are trainable parameters. The biases in these layers are kept at zero. We do not employ any normalization schemes. Lastly, we have a convolution layer at the end to get to 3 channels and required image size. This layer is followed by the sigmoidal activation. We compare this to a regular GAN which has all the same architectural components including nonlinearities except the filters are not chosen as w and $-s^2w$ and we also allow biases. We found that our injective architecture was able to achieve the same performance visually as a ‘off the shelf’ Bora et al. (2017) non-injective counterpart.

Critic network: The discriminator has 5 convolution layers with 128, 256, 512, 1024 and 1 channels per layer. Each convolution layer has 4×4 kernels with stride 2. Each layer is followed by the leaky-ReLU activation function. The last layer of the network is followed by identity.

Inference network: The inference network has the same architecture as the first 4 convolution layers of the discriminator. This is followed by 3 fully-connected layers of size 512, 256 and 256. The first 2 fully-connected layers have a Leaky ReLU activation while the last layer has identity activation function. The inference net is trained in tandem with the GAN.

We use the Wasserstein loss with gradient penalty (Gulrajani et al., 2017) to train our networks. We train for 40 epochs on a data set of size 80000 samples. We use a batch size of 64 and Adam optimizer for training with learning rate of 10^{-4} .

We report FID (Heusel et al., 2017) and Inception score (Salimans et al., 2016) using 10000 generated samples. The standard deviation was calculated using 5 sets of 10000 generated samples. In order to calculate the mean and covariance of generated distributions, we sample 50000 codes.

References

- Aviad Aberdam, Dror Simon, and Michael Elad. When and how can deep generative models be inverted? *arXiv preprint arXiv:2006.15555*, 2020.
- Robert A. Adams. *Sobolev spaces*. Academic Press [A subsidiary of Harcourt Brace Jovanovich, Publishers], New York-London, 1975. Pure and Applied Mathematics, Vol. 65.
- Lynton Ardizzone, Jakob Kruse, Sebastian Wirkert, Daniel Rahner, Eric W Pellegrini, Ralf S Klessen, Lena Maier-Hein, Carsten Rother, and Ullrich Köthe. Analyzing inverse problems with invertible neural networks. *arXiv preprint arXiv:1808.04730*, 2018.

- Simon Arridge, Peter Maass, Ozan Öktem, and Carola-Bibiane Schönlieb. Solving inverse problems using data-driven models. *Acta Numerica*, 28:1–174, 2019.
- Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.
- Yu Bai, Tengyu Ma, and Andrej Risteski. Approximability of discriminators implies diversity in gans. *arXiv preprint arXiv:1806.10586*, 2018.
- Randall Balestriero, Sebastien Paris, and Richard Baraniuk. Max-affine spline insights into deep generative networks. *arXiv preprint arXiv:2002.11912*, 2020.
- Richard G. Baraniuk and Michael B. Wakin. Random projections of smooth manifolds. *Found. Comput. Math.*, 9(1):51–77, 2009. ISSN 1615-3375. doi: 10.1007/s10208-007-9011-z. URL <https://doi.org/10.1007/s10208-007-9011-z>.
- Alexander Barvinok. Math 710: Measure concentration. *Lecture notes*, 2005. URL <http://www.math.lsa.umich.edu/~barvinok/total710.pdf>.
- Lucas Benigni and Sandrine Péché. Eigenvalue distribution of nonlinear models of random matrices. *arXiv preprint arXiv:1904.03090*, 2019.
- Ashish Bora, Ajil Jalal, Eric Price, and Alexandros G Dimakis. Compressed sensing using generative models. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 537–546. JMLR. org, 2017.
- Andrew Brock, Theodore Lim, James M Ritchie, and Nick Weston. Neural photo editing with introspective adversarial networks. *arXiv preprint arXiv:1609.07093*, 2016.
- D. S. Broomhead and M. Kirby. A new approach to dimensionality reduction: theory and algorithms. *SIAM J. Appl. Math.*, 60(6):2114–2142, 2000. ISSN 0036-1399. doi: 10.1137/S0036139998338583. URL <https://doi.org/10.1137/S0036139998338583>.
- David S Broomhead and Michael J Kirby. The whitney reduction network: A method for computing autoassociative graphs. *Neural Computation*, 13(11):2595–2616, 2001.
- Joan Bruna, Arthur Szlam, and Yann LeCun. Signal recovery from pooling representations. *arXiv preprint arXiv:1311.4025*, 2013.
- Peter Bürgisser and Felipe Cucker. *Condition: The geometry of numerical algorithms*, volume 349. Springer Science & Business Media, 2013.
- Xi Chen, Yan Duan, Rein Houthoofd, John Schulman, Ilya Sutskever, and Pieter Abbeel. Infogan: Interpretable representation learning by information maximizing generative adversarial nets. In *Advances in neural information processing systems*, pages 2172–2180, 2016.
- Moustapha Cisse, Piotr Bojanowski, Edouard Grave, Yann Dauphin, and Nicolas Usunier. Parseval networks: Improving robustness to adversarial examples. *arXiv preprint arXiv:1704.08847*, 2017.

- Ştefan Cobzaş, Radu Miculescu, and Adriana Nicolae. *Lipschitz functions*, volume 2241 of *Lecture Notes in Mathematics*. Springer, Cham, 2019. ISBN 978-3-030-16488-1; 978-3-030-16489-8. doi: 10.1007/978-3-030-16489-8. URL <https://doi.org/10.1007/978-3-030-16489-8>.
- Constantinos Daskalakis, Dhruv Rohatgi, and Manolis Zampetakis. Constant-expansion suffices for compressed sensing with generative priors. *arXiv preprint arXiv:2006.04237*, 2020.
- Laurent Dinh, David Krueger, and Yoshua Bengio. Nice: Non-linear independent components estimation. *arXiv preprint arXiv:1410.8516*, 2014.
- Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using real nvp. *arXiv preprint arXiv:1605.08803*, 2016.
- Jeff Donahue, Philipp Krähenbühl, and Trevor Darrell. Adversarial feature learning. *arXiv preprint arXiv:1605.09782*, 2016.
- Vincent Dumoulin, Ishmael Belghazi, Ben Poole, Olivier Mastropietro, Alex Lamb, Martin Arjovsky, and Aaron Courville. Adversarially learned inference. *arXiv preprint arXiv:1606.00704*, 2016.
- Simon Foucart and Holger Rauhut. *A Mathematical Introduction to Compressive Sensing*. Birkhäuser Basel, 2013. ISBN 0817649476.
- Aidan N Gomez, Mengye Ren, Raquel Urtasun, and Roger B Grosse. The reversible residual network: Backpropagation without storing activations. In *Advances in neural information processing systems*, pages 2214–2224, 2017.
- Will Grathwohl, Ricky TQ Chen, Jesse Bettencourt, Ilya Sutskever, and David Duvenaud. Ffjord: Free-form continuous dynamics for scalable reversible generative models. *arXiv preprint arXiv:1810.01367*, 2018.
- Aditya Grover, Manik Dhar, and Stefano Ermon. Flow-gan: Combining maximum likelihood and adversarial learning in generative models. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville. Improved training of wasserstein gans. In *Advances in neural information processing systems*, pages 5767–5777, 2017.
- Paul Hand and Vladislav Voroninski. Global guarantees for enforcing deep generative priors by empirical risk. *arXiv preprint arXiv:1705.07576*, 2017.
- Paul Hand, Oscar Leong, and Vlad Voroninski. Phase retrieval under a generative prior. In *Advances in Neural Information Processing Systems*, pages 9136–9146, 2018.
- Chinmay Hegde, Michael Wakin, and Richard Baraniuk. Random projections for manifold learning. In J. C. Platt, D. Koller, Y. Singer, and S. T. Roweis, editors, *Advances in Neural Information Processing Systems 20*, pages

- 641–648. Curran Associates, Inc., 2008. URL <http://papers.nips.cc/paper/3191-random-projections-for-manifold-learning.pdf>.
- Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in neural information processing systems*, pages 6626–6637, 2017.
- Morris W Hirsch. *Differential topology*, volume 33. Springer Science & Business Media, 2012.
- Kurt Hornik. Approximation capabilities of multilayer feedforward networks. *Neural networks*, 4(2):251–257, 1991.
- Sergey Ioffe. Batch renormalization: Towards reducing minibatch dependence in batch-normalized models. In *Advances in neural information processing systems*, pages 1945–1953, 2017.
- Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015.
- Mark A. Iwen and Mauro Maggioni. Approximation of points on low-dimensional manifolds via random linear projections. *Inf. Inference*, 2(1):1–31, 2013. ISSN 2049-8764. doi: 10.1093/imaiai/iat001. URL <https://doi.org/10.1093/imaiai/iat001>.
- Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4401–4410, 2019.
- Tosio Kato. *Perturbation theory for linear operators*. Classics in Mathematics. Springer-Verlag, Berlin, 1995. ISBN 3-540-58661-X. Reprint of the 1980 edition.
- Durk P Kingma and Prafulla Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. In *Advances in Neural Information Processing Systems*, pages 10215–10224, 2018.
- Abhishek Kumar, Ben Poole, and Kevin Murphy. Regularized autoencoders via relaxed injective probability flow. *arXiv preprint arXiv:2002.08927*, 2020.
- Qi Lei, Ajil Jalal, Inderjit S Dhillon, and Alexandros G Dimakis. Inverting deep generative models, one layer at a time. In *Advances in Neural Information Processing Systems*, pages 13910–13919, 2019.
- Moshe Leshno, Vladimir Ya Lin, Allan Pinkus, and Shimon Schocken. Multilayer feedforward networks with a nonpolynomial activation function can approximate any function. *Neural networks*, 6(6):861–867, 1993.
- Zinan Lin, Kiran Koshy Thekumparampil, Giulia Fanti, and Sewoong Oh. Infogan-cr: Disentangling generative adversarial networks with contrastive regularizers. *arXiv preprint arXiv:1906.06034*, 2019.
- Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, 12 2015.

- J. Luukkainen and J. Väisälä. Elements of Lipschitz topology. *Ann. Acad. Sci. Fenn. Ser. A I Math.*, 3(1):85–122, 1977. ISSN 0066-1953. doi: 10.5186/aasfm.1977.0315. URL <https://doi.org/10.5186/aasfm.1977.0315>.
- Stéphane Mallat, Sixin Zhang, and Gaspar Rochette. Phase harmonic correlations and convolutional neural networks. *arXiv preprint arXiv:1810.12136*, 2018.
- Morteza Mardani, Enhao Gong, Joseph Y Cheng, Shreyas S Vasanaawala, Greg Zaharchuk, Lei Xing, and John M Pauly. Deep generative adversarial neural networks for compressive sensing mri. *IEEE transactions on medical imaging*, 38(1):167–179, 2018.
- Pertti Mattila. *Geometry of sets and measures in Euclidean spaces*, volume 44 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 1995. ISBN 0-521-46576-1; 0-521-65595-1. doi: 10.1017/CBO9780511623813. URL <https://doi.org/10.1017/CBO9780511623813>. Fractals and rectifiability.
- John W. Milnor and James D. Stasheff. *Characteristic classes*. Princeton University Press, Princeton, N. J.; University of Tokyo Press, Tokyo, 1974. Annals of Mathematics Studies, No. 76.
- Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral normalization for generative adversarial networks. *arXiv preprint arXiv:1802.05957*, 2018.
- Frank Morgan. *Geometric measure theory: a beginner’s guide*. Academic press, 2016.
- Amiya Mukherjee. *Differential topology*. Hindustan Book Agency, New Delhi; Birkhäuser/Springer, Cham, second edition, 2015. ISBN 978-3-319-19044-0; 978-3-319-19045-7. doi: 10.1007/978-3-319-19045-7. URL <https://doi.org/10.1007/978-3-319-19045-7>.
- Augustus Odena, Jacob Buckman, Catherine Olsson, Tom B Brown, Christopher Olah, Colin Raffel, and Ian Goodfellow. Is generator conditioning causally related to gan performance? *arXiv preprint arXiv:1802.08768*, 2018.
- Jeffrey Pennington and Pratik Worah. Nonlinear random matrix theory for deep learning. In *Advances in Neural Information Processing Systems*, pages 2637–2646, 2017.
- Peter Petersen, S Axler, and KA Ribet. *Riemannian geometry*, volume 171. Springer, 2006.
- Allan Pinkus. Approximation theory of the mlp model in neural networks. *Acta numerica*, 8: 143–195, 1999.
- Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015.
- Tim Salimans and Durk P Kingma. Weight normalization: A simple reparameterization to accelerate training of deep neural networks. In *Advances in neural information processing systems*, pages 901–909, 2016.

- Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. In *Advances in neural information processing systems*, pages 2234–2242, 2016.
- Rolf Schneider. Curvatures of typical convex bodies—the complete picture. *Proceedings of the American Mathematical Society*, 143(1):387–393, 2015.
- Stefan Scholtes. Homeomorphism conditions for coherently oriented piecewise affine mappings. *Mathematics of operations research*, 21(4):955–978, 1996.
- Viraj Shah and Chinmay Hegde. Solving linear inverse problems using gan priors: An algorithm with provable guarantees. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4609–4613. IEEE, 2018.
- Richard P Stanley. Enumerative combinatorics volume 1 second edition. *Cambridge studies in advanced mathematics*, 2011.
- Andreas M Tillmann and Marc E Pfetsch. The computational complexity of the restricted isometry property, the nullspace property, and related concepts in compressed sensing. *IEEE Transactions on Information Theory*, 60(2):1248–1259, 2013.
- Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Instance normalization: The missing ingredient for fast stylization. *arXiv preprint arXiv:1607.08022*, 2016.
- RO Winder. Partitions of n-space by hyperplanes. *SIAM Journal on Applied Mathematics*, 14(4):811–818, 1966.
- Yan Wu, Mihaela Rosca, and Timothy Lillicrap. Deep compressed sensing. *arXiv preprint arXiv:1905.06723*, 2019.
- Yuxin Wu and Kaiming He. Group normalization. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 3–19, 2018.
- Dmitry Yarotsky. Error bounds for approximations with deep relu networks. *Neural Networks*, 94:103–114, 2017.