
Invertible Residual Networks

Jens Behrmann^{*1,2} Will Grathwohl^{*2} Ricky T. Q. Chen² David Duvenaud² Jörn-Henrik Jacobsen^{*2}

Abstract

We show that standard ResNet architectures can be made invertible, allowing the same model to be used for classification, density estimation, and generation. Typically, enforcing invertibility requires partitioning dimensions or restricting network architectures. In contrast, our approach only requires adding a simple normalization step during training, already available in standard frameworks. Invertible ResNets define a generative model which can be trained by maximum likelihood on unlabeled data. To compute likelihoods, we introduce a tractable approximation to the Jacobian log-determinant of a residual block. Our empirical evaluation shows that invertible ResNets perform competitively with both state-of-the-art image classifiers and flow-based generative models, something that has not been previously achieved with a single architecture.

1. Introduction

One of the main appeals of neural network-based models is that a single model architecture can often be used to solve a variety of related tasks. However, many recent advances are based on special-purpose solutions tailored to particular domains. State-of-the-art architectures in unsupervised learning, for instance, are becoming increasingly domain-specific (Van Den Oord et al., 2016b; Kingma & Dhariwal, 2018; Parmar et al., 2018; Karras et al., 2018; Van Den Oord et al., 2016a). On the other hand, one of the most successful feed-forward architectures for discriminative learning are deep residual networks (He et al., 2016; Zagoruyko & Komodakis, 2016), which differ considerably from their generative counterparts. This divide makes it complicated to choose or design a suitable architecture for a given task. It also makes it hard for discriminative tasks to benefit from

^{*}Equal contribution ¹University of Bremen, Center for Industrial Mathematics ²Vector Institute and University of Toronto. Correspondence to: Jens Behrmann <jensb@uni-bremen.de>, Jörn-Henrik Jacobsen <j.jacobsen@vectorinstitute.ai>.




Figure 1. Dynamics of a standard residual network (left) and invertible residual network (right). Both networks map the interval $[-2, 2]$ to: 1) noisy x^3 -function at half depth and 2) noisy identity function at full depth. Invertible ResNets describe a bijective continuous dynamics while regular ResNets result in crossing and collapsing paths (circled in white) which correspond to non-bijective continuous dynamics. Due to collapsing paths, standard ResNets are not a valid density model.

unsupervised learning. We bridge this gap with a new class of architectures that perform well in both domains.

To achieve this, we focus on reversible networks which have been shown to produce competitive performance on discriminative (Gomez et al., 2017; Jacobsen et al., 2018) and generative (Dinh et al., 2014; 2017; Kingma & Dhariwal, 2018) tasks independently, albeit in the same model paradigm. They typically rely on fixed dimension splitting heuristics, but common splittings interleaved with non-volume conserving elements are constraining and their choice has a significant impact on performance (Kingma & Dhariwal, 2018; Dinh et al., 2017). This makes building reversible networks a difficult task. In this work we show that these exotic designs, necessary for competitive density estimation performance, can severely hurt discriminative performance.

To overcome this problem, we leverage the viewpoint of ResNets as an Euler discretization of ODEs (Haber & Ruthotto, 2018; Ruthotto & Haber, 2018; Lu et al., 2017; Ciccone et al., 2018) and prove that invertible ResNets (i-ResNets) can be constructed by simply changing the normalization scheme of standard ResNets.

As an intuition, Figure 1 visualizes the differences in the dynamics learned by standard and invertible ResNets.

This approach allows unconstrained architectures for each residual block, while only requiring a Lipschitz constant smaller than one for each block. We demonstrate that this restriction negligibly impacts performance when building image classifiers - they perform on par with their non-invertible counterparts on classifying MNIST, CIFAR10 and CIFAR100 images.

We then show how i-ResNets can be trained as maximum likelihood generative models on unlabeled data. To compute likelihoods, we introduce a tractable approximation to the Jacobian determinant of a residual block. Like FJORD (Grathwohl et al., 2019), i-ResNet flows have unconstrained (free-form) Jacobians, allowing them to learn more expressive transformations than the triangular mappings used in other reversible models. Our empirical evaluation shows that i-ResNets perform competitively with both state-of-the-art image classifiers and flow-based generative models, bringing general-purpose architectures one step closer to reality.¹

2. Enforcing Invertibility in ResNets

There is a remarkable similarity between ResNet architectures and Euler’s method for ODE initial value problems:

$$\begin{aligned} x_{t+1} &\leftarrow x_t + g_{\theta_t}(x_t) \\ x_{t+1} &\leftarrow x_t + h f_{\theta_t}(x_t) \end{aligned}$$

where $x_t \in \mathbb{R}^d$ represent activations or states, t represents layer indices or time, $h > 0$ is a step size, and g_{θ_t} is a residual block. This connection has attracted research at the intersection of deep learning and dynamical systems (Lu et al., 2017; Haber & Ruthotto, 2018; Ruthotto & Haber, 2018; Chen et al., 2018). However, little attention has been paid to the dynamics backwards in time

$$\begin{aligned} x_t &\leftarrow x_{t+1} - g_{\theta_t}(x_t) \\ x_t &\leftarrow x_{t+1} - h f_{\theta_t}(x_t) \end{aligned}$$

which amounts to the implicit backward Euler discretization. In particular, solving the dynamics backwards in time would implement an inverse of the corresponding ResNet. The following theorem states that a simple condition suffices to make the dynamics solvable and thus renders the ResNet invertible:

Theorem 1 (Sufficient condition for invertible ResNets). *Let $F_{\theta} : \mathbb{R}^d \rightarrow \mathbb{R}^d$ with $F_{\theta} = (F_{\theta}^1 \circ \dots \circ F_{\theta}^T)$ denote a ResNet with blocks $F_{\theta}^t = I + g_{\theta_t}$. Then, the ResNet F_{θ} is*

¹Official code release: <https://github.com/jhjacobson/invertible-resnet>

Algorithm 1. Inverse of i-ResNet layer via fixed-point iteration.

Input: output from residual layer y , contractive residual block g , number of fixed-point iterations n
 Init: $x^0 := y$
for $i = 0, \dots, n$ **do**
 $x^{i+1} := y - g(x^i)$
end for

invertible if

$$\text{Lip}(g_{\theta_t}) < 1, \text{ for all } t = 1, \dots, T,$$

where $\text{Lip}(g_{\theta_t})$ is the Lipschitz-constant of g_{θ_t} .

Note that this condition is not necessary for invertibility. Other approaches (Dinh et al., 2014; 2017; Jacobsen et al., 2018; Chang et al., 2018; Kingma & Dhariwal, 2018) rely on partitioning dimensions or autoregressive structures to create analytical inverses.

While enforcing $\text{Lip}(g) < 1$ makes the ResNet invertible, we have no analytic form of this inverse. However, we can obtain it through a simple fixed-point iteration, see Algorithm 1. Note, that the starting value for the fixed-point iteration can be any vector, because the fixed-point is unique. However, using the output $y = x + g(x)$ as the initialization $x^0 := y$ is a good starting point since y was obtained from x only via a bounded perturbation of the identity. From the Banach fixed-point theorem we have

$$\|x - x^n\|_2 \leq \frac{\text{Lip}(g)^n}{1 - \text{Lip}(g)} \|x^1 - x^0\|_2. \quad (1)$$

Thus, the convergence rate is exponential in the number of iterations n and smaller Lipschitz constants will yield faster convergence.

Additional to invertibility, a contractive residual block also renders the residual layer bi-Lipschitz.

Lemma 2 (Lipschitz constants of Forward and Inverse). *Let $F(x) = x + g(x)$ with $\text{Lip}(g) = L < 1$ denote the residual layer. Then, it holds*

$$\text{Lip}(F) \leq 1 + L \quad \text{and} \quad \text{Lip}(F^{-1}) \leq \frac{1}{1 - L}.$$

Hence by design, invertible ResNets offer stability guarantees for both their forward and inverse mapping. In the following section, we discuss approaches to enforce the Lipschitz condition.

2.1. Satisfying the Lipschitz Constraint

We implement residual blocks as a composition of contractive nonlinearities ϕ (e.g. ReLU, ELU, tanh) and linear mappings.

For example, in our convolutional networks $g = W_3\phi(W_2\phi(W_1))$, where W_i are convolutional layers. Hence,

$$\text{Lip}(g) < 1, \quad \text{if } \|W_i\|_2 < 1,$$

where $\|\cdot\|_2$ denotes the spectral norm. Note, that regularizing the spectral norm of the Jacobian of g (Sokoli et al., 2017) only reduces it locally and does not guarantee the above condition. Thus, we will enforce $\|W_i\|_2 < 1$ for each layer.

A power-iteration on the parameter matrix as in Miyato et al. (2018) approximates only a bound on $\|W_i\|_2$ instead of the true spectral norm, if the filter kernel is larger than 1×1 , see Tsuzuku et al. (2018) for details on the bound. Hence, unlike Miyato et al. (2018), we directly estimate the spectral norm of W_i by performing power-iteration using W_i and W_i^T as proposed in Gouk et al. (2018). The power-iteration yields an under-estimate $\tilde{\sigma}_i \leq \|W_i\|_2$. Using this estimate, we normalize via

$$\tilde{W}_i = \begin{cases} cW_i/\tilde{\sigma}_i, & \text{if } c/\tilde{\sigma}_i < 1 \\ W_i, & \text{else} \end{cases}, \quad (2)$$

where the hyper-parameter $c < 1$ is a scaling coefficient. Since $\tilde{\sigma}_i$ is an under-estimate, $\|W_i\|_2 \leq c$ is not guaranteed. However, after training Sedghi et al. (2019) offer an approach to inspect $\|W_i\|_2$ exactly using the SVD on the Fourier transformed parameter matrix, which will allow us to show $\text{Lip}(g) < 1$ holds in all cases.


3. Generative Modelling with i-ResNets

We can define a simple generative model for data $x \in \mathbb{R}^d$ by first sampling $z \sim p_z(z)$ where $z \in \mathbb{R}^d$ and then defining $x = \Phi(z)$ for some function $\Phi: \mathbb{R}^d \rightarrow \mathbb{R}^d$. If Φ is invertible and we define $F = \Phi^{-1}$, then we can compute the likelihood of any x under this model using the change of variables formula

$$\ln p_x(x) = \ln p_z(z) + \ln |\det J_F(x)|, \quad (3)$$

where $J_F(x)$ is the Jacobian of F evaluated at x . Models of this form are known as Normalizing Flows (Rezende & Mohamed, 2015). They have recently become a popular model for high-dimensional data due to the introduction of powerful bijective function approximators whose Jacobian log-determinant can be efficiently computed (Dinh et al., 2014; 2017; Kingma & Dhariwal, 2018; Chen et al., 2018) or approximated (Grathwohl et al., 2019).

Since i-ResNets are guaranteed to be invertible we can use them to parameterize F in Equation (3). Samples from this model can be drawn by first sampling $z \sim p(z)$ and then computing $x = F^{-1}(z)$ with Algorithm 1. In Figure 2 we



Data Samples Glow i-ResNet

Figure 2. Visual comparison of i-ResNet flow and Glow. Details of this experiment can be found in Appendix C.3.

show an example of using an i-ResNet to define a generative model on some two-dimensional datasets compared to Glow (Kingma & Dhariwal, 2018).

3.1. Scaling to Higher Dimensions

While the invertibility of i-ResNets allows us to use them to define a Normalizing Flow, we must compute $\ln |\det J_F(x)|$ to evaluate the data-density under the model. Computing this quantity has a time cost of $\mathcal{O}(d^3)$ in general which makes naïvely scaling to high-dimensional data impossible.

To bypass this constraint we present a tractable approximation to the log-determinant term in Equation (3), which will scale to high dimensions d . Previously, Ramesh & LeCun (2018) introduced the application of log-determinant estimation to non-invertible deep generative models without the specific structure of i-ResNets.

First, we note that the Lipschitz constrained perturbations $x + g(x)$ of the identity yield positive determinants, hence

$$|\det J_F(x)| = \det J_F(x),$$

see Lemma 6 in Appendix A. Combining this result with the matrix identity $\ln \det(A) = \text{tr}(\ln(A))$ for non-singular $A \in \mathbb{R}^{d \times d}$ (see e.g. Withers & Nadarajah (2010)), we have

$$\ln |\det J_F(x)| = \text{tr}(\ln J_F),$$

where tr denotes the matrix trace and \ln the matrix logarithm. Thus for $z = F(x) = (I + g)(x)$, it is

$$\ln p_x(x) = \ln p_z(z) + \text{tr}(\ln(I + J_g(x))).$$

The trace of the matrix logarithm can be expressed as a

power series (Hall, 2015)

$$\operatorname{tr}(\ln(I + J_g(x))) = \sum_{k=1}^{\infty} (-1)^{k+1} \frac{\operatorname{tr}(J_g^k)}{k}, \quad (4)$$

which converges if $\|J_g\|_2 < 1$. Hence, due to the Lipschitz constraint, we can compute the log-determinant via the above power series with guaranteed convergence.

Before we present a stochastic approximation to the above power series, we observe following properties of i-ResNets: Due to $\operatorname{Lip}(g_t) < 1$ for the residual block of each layer t , we can provide a lower and upper bound on its log-determinant with

$$\begin{aligned} d \sum_{t=1}^T \ln(1 - \operatorname{Lip}(g_t)) &\leq \ln |\det J_F(x)| \\ d \sum_{t=1}^T \ln(1 + \operatorname{Lip}(g_t)) &\geq \ln |\det J_F(x)|, \end{aligned}$$

for all $x \in \mathbb{R}$, see Lemma 7 in Appendix A. Thus, both the number of layers T and the Lipschitz constant affect the contraction and expansion bounds of i-ResNets and must be taken into account when designing such an architecture.

3.2. Stochastic Approximation of log-determinant

Expressing the log-determinant with the power series in (4) has three main computational drawbacks: 1) Computing $\operatorname{tr}(J_g)$ exactly costs $\mathcal{O}(d^2)$, or approximately needs d evaluations of g as each entry of the diagonal of the Jacobian requires the computation of a separate derivative of g (Grathwohl et al., 2019). 2) Matrix powers J_g^k are needed, which requires the knowledge of the full Jacobian. 3) The series is infinite.

Fortunately, drawback 1) and 2) can be alleviated. First, vector-Jacobian products $v^T J_g$ can be computed at approximately the same costs as evaluating g through reverse-mode automatic differentiation. Second, a stochastic approximation of the matrix trace of $A \in \mathbb{R}^{d \times d}$

$$\operatorname{tr}(A) = \mathbb{E}_{p(v)} [v^T A v],$$

known as the Hutchinsons trace estimator, can be used to estimate $\operatorname{tr}(J_g^k)$. The distribution $p(v)$ needs to fulfill $\mathbb{E}[v] = 0$ and $\operatorname{Cov}(v) = I$, see (Hutchinson, 1990; Avron & Toledo, 2011).

While this allows for an unbiased estimate of the matrix trace, to achieve bounded computational costs, the power series (4) will be truncated at index n to address drawback 3). Algorithm 2 summarizes the basic steps. The truncation turns the unbiased estimator into a biased estimator, where the bias depends on the truncation error. Fortunately, this error can be bounded as we demonstrate below.

Algorithm 2. Forward pass of an invertible ResNets with Lipschitz constraint and log-determinant approximation, SN denotes spectral normalization based on (2).

Input: data point x , network F , residual block g , number of power series terms n

for Each residual block **do**

Lip constraint: $\hat{W}_j := \operatorname{SN}(W_j, x)$ for linear Layer W_j .

Draw v from $\mathcal{N}(0, I)$

$w^T := v^T$

$\ln \det := 0$

for $k = 1$ **to** n **do**

$w^T := w^T J_g$ (vector-Jacobian product)

$\ln \det := \ln \det + (-1)^{k+1} w^T v / k$

end for

end for

To improve the stability of optimization when using this estimator we recommend using nonlinearities with continuous derivatives such as ELU (Clevert et al., 2015) or softplus instead of ReLU (See Appendix C.3).

3.3. Error of Power Series Truncation

We estimate $\ln |\det(I + J_g)|$ with the finite power series

$$PS(J_g, n) := \sum_{k=1}^n (-1)^{k+1} \frac{\operatorname{tr}(J_g^k)}{k}, \quad (5)$$

where we have (with some abuse of notation) $PS(J_g, \infty) = \operatorname{tr}(\ln(I + J_g))$. We are interested in bounding the truncation error of the log-determinant as a function of the data dimension d , the Lipschitz constant $\operatorname{Lip}(g)$ and the number of terms in the series n .

Theorem 3 (Approximation error of Loss). *Let g denote the residual function and J_g the Jacobian as before. Then, the error of a truncated power series at term n is bounded as*

$$\begin{aligned} &|PS(J_g, n) - \ln \det(I + J_g)| \\ &\leq -d \left(\ln(1 - \operatorname{Lip}(g)) + \sum_{k=1}^n \frac{\operatorname{Lip}(g)^k}{k} \right). \end{aligned}$$

While the result above gives an error bound for evaluation of the loss, during training the error in the gradient of the loss is of greater interest. Similarly, we can obtain the following bound. The proofs are given in Appendix A.

Theorem 4 (Convergence Rate of Gradient Approximation). *Let $\theta \in \mathbb{R}^p$ denote the parameters of network F , let g, J_g be as before. Further, assume bounded inputs and a Lipschitz activation function with Lipschitz derivative. Then, we obtain the convergence rate*

$$\|\nabla_{\theta} (\ln \det(I + J_g)) - PS(J_g, n)\|_{\infty} = \mathcal{O}(c^n)$$

where $c := \text{Lip}(g)$ and n the number of terms used in the power series.

In practice, only 5-10 terms must be taken to obtain a bias less than .001 bits per dimension, which is typically reported up to .01 precision (See Appendix E).

4. Related Work

4.1. Reversible Architectures

We put our focus on invertible architectures with efficient inverse computation, namely NICE (Dinh et al., 2014), i-RevNet (Jacobsen et al., 2018), Real-NVP (Dinh et al., 2017), Glow (Kingma & Dhariwal, 2018) and Neural ODEs (Chen et al., 2018) and its stochastic density estimator FFJORD (Grathwohl et al., 2019). A summary of the comparison between different reversible networks is given in Table 1.

The dimension-splitting approach used in NICE, i-RevNet, Real-NVP and Glow allows for both analytic forward and inverse mappings. However, this restriction required the introduction of additional steps like invertible 1×1 convolutions in Glow (Kingma & Dhariwal, 2018). These 1×1 convolutions need to be inverted numerically, making Glow altogether not analytically invertible. In contrast, i-ResNet can be viewed as an intermediate approach, where the forward mapping is given analytically, while the inverse can be computed via a fixed-point iteration.

Furthermore, an i-ResNet block has a Lipschitz bound both for forward and inverse (Lemma 2), while other approaches do not have this property by design. Hence, i-ResNets could be an interesting avenue for stability-critical applications like inverse problems (Ardizzone et al., 2019) or invariance-based adversarial vulnerability (Jacobsen et al., 2019).

Neural ODEs (Chen et al., 2018) allow free-form dynamics similar to i-ResNets, meaning that any architecture could be used as long as the input and output dimensions are the same. To obtain discrete forward and inverse dynamics, Neural ODEs rely on adaptive ODE solvers, which allows for an accuracy vs. speed trade-off. Yet, scalability to very high input dimension such as high-resolution images remains unclear.

4.2. Ordinary Differential Equations

Due to the similarity of ResNets and Euler discretizations, there are many connections between the i-ResNet and ODEs, which we review in this section.

Relationship of i-ResNets to Neural ODEs: The view of deep networks as dynamics over time offers two fundamental learning approaches: 1) Direct learning of dynamics using discrete architectures like ResNets (Haber & Ruthotto,

2018; Ruthotto & Haber, 2018; Lu et al., 2017; Ciccone et al., 2018). 2) Indirect learning of dynamics via parametrizing an ODE with a neural network as in Chen et al. (2018); Grathwohl et al. (2019).

The dynamics $x(t)$ of a fixed ResNet F_θ are only defined at time points t_i corresponding to each block $g_{\theta_{t_i}}$. However, a linear interpolation in time can be used to generate continuous dynamics. See Figure 1, where the continuous dynamics of a linearly interpolated invertible ResNet are shown against those of a standard ResNet. Invertible ResNets are bijective along the continuous path while regular ResNets may result in crossing or merging paths. The indirect approach of learning an ODE, on the other hand, adapts the discretization based on an ODE-solver, but does not have a fixed computational budget compared to an i-ResNet.

Stability of ODEs: There are two main approaches to study the stability of ODEs, 1) behavior for $t \rightarrow \infty$ and 2) Lipschitz stability over finite time intervals $[0, T]$. Based on time-invariant dynamics $f(x(t))$, (Ciccone et al., 2018) constructed asymptotically stable ResNets using anti-symmetric layers such that $\text{Re}(\lambda(J_x)) < 0$ (with $\text{Re}(\lambda(\cdot))$ denoting the real-part of eigenvalues, $\rho(\cdot)$ spectral radius and $J_x g$ the Jacobian at point x). By projecting weights based on the Gershgorin circle theorem, they further fulfilled $\rho(J_x g) < 1$, yielding asymptotically stable ResNets with shared weights over layers. On the other hand, (Haber & Ruthotto, 2018; Ruthotto & Haber, 2018) considered time-dependent dynamics $f(x(t), \theta(t))$ corresponding to standard ResNets. They induce stability by using anti-symmetric layers and projections of the weights. Contrarily, initial value problems on $[0, T]$ are well-posed for Lipschitz continuous dynamics (Ascher, 2008). Thus, the invertible ResNet with $\text{Lip}(f) < 1$ can be understood as a stabilizer of an ODE for step size $h = 1$ without a restriction to anti-symmetric layers as in Ruthotto & Haber (2018); Haber & Ruthotto (2018); Ciccone et al. (2018).

4.3. Spectral Sum Approximations

The approximation of spectral sums like the log-determinant is of broad interest for many machine learning problems such as Gaussian Process regression (Dong et al., 2017). Among others, Taylor approximation (Boutsidis et al., 2017) of the log-determinant similar to our approach or Chebyshev polynomials (Han et al., 2016) are used. In Boutsidis et al. (2017), error bounds on the estimation via truncated power series and stochastic trace estimation are given for symmetric positive definite matrices. However, $I + J_g$ is not symmetric and thus, their analysis does not apply here.

Recently, unbiased estimates (Adams et al., 2018) and unbiased gradient estimators (Han et al., 2018) were proposed for symmetric positive definite matrices. Furthermore, Chebyshev polynomials have been used to approximate the log-

Invertible Residual Networks

Method	ResNet	NICE/ i-RevNet	Real-NVP	Glow	FFJORD	i-ResNet
Free-form	✓	✗	✗	✗	✓	✓
Analytic Forward	✓	✓	✓	✓	✗	✓
Analytic Inverse	N/A	✓	✓	✗	✗	✗
Non-volume Preserving	N/A	✗	✓	✓	✓	✓
Exact Likelihood	N/A	✓	✓	✓	✗	✗
Unbiased Stochastic Log-Det Estimator	N/A	N/A	N/A	N/A	✓	✗

Table 1. Comparing i-ResNet and ResNets to NICE (Dinh et al., 2014), Real-NVP (Dinh et al., 2017), Glow (Kingma & Dhariwal, 2018) and FFJORD (Grathwohl et al., 2019). Non-volume preserving refers to the ability to allow for contraction and expansions and exact likelihood to compute the change of variables (3) exactly. The unbiased estimator refers to a stochastic approximation of the log-determinant, see section 3.2.

determinant of Jacobian of deep neural networks in Ramesh & LeCun (2018) for density matching and evaluation of the likelihood of GANs.

5. Experiments

We complete a thorough experimental survey of invertible ResNets. First, we numerically verify the invertibility of i-ResNets. Then, we investigate their discriminative abilities on a number of common image classification datasets. Furthermore, we compare the discriminative performance of i-ResNets to other invertible networks. Finally, we study how i-ResNets can be used to define generative models.

5.1. Validating Invertibility and Classification

To compare the discriminative performance and invertibility of i-ResNets with standard ResNet architectures, we train both models on CIFAR10, CIFAR100, and MNIST. The CIFAR and MNIST models have models have 54 and 21 residual blocks, respectively and we use identical settings for all other hyperparameters. We replace strided downsampling with “invertible downsampling” operations (Jacobsen et al., 2018) to ensure bijectivity, see Appendix C.2 for training and architectural details. We increase the number of input channels to 16 by padding with zeros. This is analogous to the standard practice of projecting the data into a higher-dimensional space using a standard convolutional layer at the input of a model, but this mapping is reversible. To obtain the numerical inverse, we apply 100 fixed point iterations (Equation (1)) for each block. This number is chosen to ensure that the poor reconstructions for vanilla ResNets (see Figure 3) are not due to using too few iterations. In practice far fewer iterations suffice, as the trade-off between reconstruction error and number of iterations analyzed in Appendix D shows.

Classification and reconstruction results for a baseline pre-activation ResNet-164, a ResNet with architecture like i-ResNets without Lipschitz constraint (denoted as vanilla)



Figure 3. Original images (top) and reconstructions from i-ResNet with $c = 0.9$ (middle) and a standard ResNet with the same architecture (bottom), showing that the fixed point iteration does not recover the input without the Lipschitz constraint.

and five invertible ResNets with different spectral normalization coefficients are shown in Table 2. The results illustrate that for larger settings of the layer-wise Lipschitz constant c , our proposed invertible ResNets perform competitively with the baselines in terms of classification performance, while being provably invertible. When applying very conservative normalization (small c), the classification error becomes higher on all datasets tested.

To demonstrate that our normalization scheme is effective and that standard ResNets are not generally invertible, we reconstruct inputs from the features of each model using Algorithm 1. Intriguingly, our analysis also reveals that unconstrained ResNets are invertible after training on MNIST (see Figure 7 in Appendix B), whereas on CIFAR10/100 they are not. Further, we find ResNets with and without BatchNorm are not invertible after training on CIFAR10, which can also be seen from the singular value plots in Appendix B (Figure 6). The runtime on 4 GeForce GTX 1080 GPUs with 1 spectral norm iteration was 0.5 sec for a forward and backward pass of batch with 128 samples, while it took 0.2 sec without spectral normalization. See section C.1 (appendix) for details on the runtime.

The reconstruction error decays quickly and the errors are already imperceptible after 5-20 iterations, which is the cost of 5-20 times the forward pass and corresponds to 0.15-0.75 seconds for reconstructing 100 CIFAR10 images.

Invertible Residual Networks

		ResNet-164	Vanilla	$c = 0.9$	$c = 0.8$	$c = 0.7$	$c = 0.6$	$c = 0.5$
Classification	MNIST	-	0.38	0.40	0.42	0.40	0.42	0.86
	Error %							
	CIFAR10	5.50	6.69	6.78	6.86	6.93	7.72	8.71
	CIFAR100	24.30	23.97	24.58	24.99	25.99	27.30	29.45
Guaranteed Inverse		No	No	Yes	Yes	Yes	Yes	Yes

Table 2. Comparison of i-ResNet to a ResNet-164 baseline architecture of similar depth and width with varying Lipschitz constraints via coefficients c . Vanilla shares the same architecture as i-ResNet, without the Lipschitz constraint.

Computing the inverse is fast even for the largest normalization coefficient, but becomes faster with stronger normalization. The number of iterations needed for full convergence is approximately cut in half when reducing the spectral normalization coefficient by 0.2, see Figure 8 (Appendix D) for a detailed plot. We also ran an i-RevNet (Jacobsen et al., 2018) with comparable hyperparameters as ResNet-164 and it performs on par with ResNet-164 with 5.6%. Note however, that i-RevNets, like NICE (Dinh et al., 2014), are volume-conserving, making them less well-suited to generative modeling.

In summary, we observe that invertibility without additional constraints is unlikely, but possible, whereas it is hard to predict if networks will have this property. In our proposed model, we can guarantee the existence of an inverse without significantly harming classification performance.

5.2. Comparison with Other Invertible Architectures

In this section we compare i-ResNet classifiers to the state-of-the-art invertible flow-based model Glow. We take the implementation of Kingma & Dhariwal (2018) and modify it to classify CIFAR10 images (with no generative modeling component). We create an i-ResNet that is as close as possible in structure to the default Glow model on CIFAR10 (denoted as i-ResNet Glow-style) and compare it to two variants of Glow, one that uses learned (1×1 convolutions) and affine block structure, and one with reverse permutations (like Real-NVP) and additive block structure. Results of this experiment can be found in Table 3. We can see that i-ResNets outperform all versions of Glow on this

Affine Glow 1×1 Conv	Additive Glow Reverse	i-ResNet Glow-Style	i-ResNet 164
12.63	12.36	8.03	6.69

Table 3. CIFAR10 classification results compared to state-of-the-art flow Glow as a classifier. We compare two versions of Glow, as well as an i-ResNet architecture as similar as possible to Glow in its number of layers and channels, termed “i-ResNet, Glow-Style”.

discriminative task, even when adapting the network depth and width to that of Glow. This indicates that i-ResNets have a more suitable inductive bias in their block structure for discriminative tasks than Glow.

We also find that i-ResNets are considerably easier to train than these other models. We are able to train i-ResNets using SGD with momentum and a learning rate of 0.1 whereas all version of Glow we tested needed Adam or Adamax (Kingma & Ba, 2014) and much smaller learning rates to avoid divergence.

5.3. Generative Modeling

We run a number of experiments to verify the utility of i-ResNets in building generative models. First, we compare i-ResNet Flows with Glow (Kingma & Dhariwal, 2018) on simple two-dimensional datasets. Figure 2 qualitatively shows the density learned by a Glow model with 100 coupling layers and 100 invertible linear transformations. We compare against an i-ResNet where the coupling layers are replaced by invertible residual blocks with the same number of parameters and the invertible linear transformations are replaced by actnorm (Kingma & Dhariwal, 2018). This results in the i-ResNet model having slightly fewer parameters, while maintaining an equal number of layers. In this experiment we train i-ResNets using the brute-force computed log-determinant since the data is two-dimensional. We find that i-ResNets are able to more accurately fit these simple densities. As stated in Grathwohl et al. (2019), we believe this is due to our model’s ability to avoid partitioning dimensions.

Next we evaluate i-ResNets as a generative model for images on MNIST and CIFAR10. Our models consist of multiple i-ResNet blocks followed by invertible downsampling or dimension “squeezing” to downsample the spatial dimensions. We use multi-scale architectures like those of Dinh et al. (2017); Kingma & Dhariwal (2018). In these experiments we train i-ResNets using the log-determinant approximation, see Algorithm 2. Full architecture, experimental, and evaluation details can be found in Appendix C.3. Samples from our CIFAR10 model are shown in Figure 5 and samples from our MNIST model can be found in Appendix F.

Method	MNIST	CIFAR10
NICE (Dinh et al., 2014)	4.36	4.48†
MADE (Germain et al., 2015)	2.04	5.67
MAF (Papamakarios et al., 2017)	1.89	4.31
Real NVP (Dinh et al., 2017)	1.06	3.49
Glow (Kingma & Dhariwal, 2018)	1.05	3.35
FFJORD (Grathwohl et al., 2019)	0.99	3.40
i-ResNet	1.06	3.45

Table 4. MNIST and CIFAR10 bits/dim results. † Uses ZCA pre-processing making results not directly comparable.

Compared to the classification model, the log-determinant approximation with 5 series terms roughly increased the computation times by a factor of 4. The bias and variance of our log-determinant estimator is shown in Figure 4.

Results and comparisons to other generative models can be found in Table 4. While our models did not perform as well as Glow and FFJORD, we find it intriguing that ResNets, with very little modification, can create a generative model competitive with these highly engineered models. We believe the gap in performance is mainly due to our use of a biased log-determinant estimator and that the use of an unbiased method (Han et al., 2018) can help close this gap.

6. Other Applications

In many applications, a secondary unsupervised learning or generative modeling objective is formulated in combination with a primary discriminative task. i-ResNets are appealing here, as they manage to achieve competitive performance on both discriminative and generative tasks. We summarize some application areas to highlight that there is a wide variety of tasks for which i-ResNets would be promising to consider:

- Hybrid density and discriminative models for joint classification and detection or fairness applications (Nalis-




Figure 4. Bias and standard deviation of our log-determinant estimator as the number of power series terms increases. Variance is due to the stochastic trace estimator.



Figure 5. CIFAR10 samples from our i-ResNet flow. More samples can be found in Appendix F.

nick et al., 2018; Louizos et al., 2016)

- Unsupervised learning for downstream tasks (Hjelm et al., 2019; Van Den Oord et al., 2018)
- Semi-supervised learning from few labeled examples (Oliver et al., 2018; Kingma et al., 2014)
- Solving inverse problems with hybrid regression and generative losses (Ardizzone et al., 2019)
- Adversarial robustness with likelihood-based generative models (Schott et al., 2019; Jacobsen et al., 2019)

Finally, it is plausible that the Lipschitz bounds on the layers of the i-ResNet could aid with the stability of gradients for optimization, as well as adversarial robustness.

7. Conclusions

We introduced a new architecture, i-ResNets, which allow free-form layer architectures while still providing tractable density estimates. The unrestricted form of the Jacobian allows expansion and contraction via the residual blocks, while partitioning-based models (Dinh et al., 2014; 2017; Kingma & Dhariwal, 2018) must include affine blocks and scaling layers to be non-volume preserving.

Several challenges remain to be addressed in future work. First, our estimator of the log-determinant is biased. However, there have been recent advances in building unbiased estimators for the log-determinant (Han et al., 2018), which we believe could improve the performance of our generative model. Second, learning and designing networks with a Lipschitz constraint is challenging. For example, we need to constrain each linear layer in the block instead of being able to directly control the Lipschitz constant of a block, see Anil et al. (2018) for a promising approach for addressing this problem.

Acknowledgments

We thank Rich Zemel for very helpful comments on an earlier version of the manuscript. We thank Yulia Rubanova

for spotting a mistake in one of the proofs. We also thank everyone else at Vector for helpful discussions and feedback.

We gratefully acknowledge the financial support from the German Science Foundation for RTG 2224 "π³: Parameter Identification - Analysis, Algorithms, Applications"

References

- Adams, R. P., Pennington, J., Johnson, M. J., Smith, J., Ovadia, Y., Patton, B., and Saunderson, J. Estimating the spectral density of large implicit matrices. *arXiv preprint arXiv:1802.03451*, 2018.
- Anil, C., Lucas, J., and Grosse, R. Sorting out lipschitz function approximation. *arXiv preprint arXiv:1811.05381*, 2018.
- Ardizzone, L., Kruse, J., Rother, C., and Köthe, U. Analyzing inverse problems with invertible neural networks. In *International Conference on Learning Representations*, 2019.
- Ascher, U. *Numerical methods for evolutionary differential equations*. Computational science and engineering. Society for Industrial and Applied Mathematics, 2008.
- Avron, H. and Toledo, S. Randomized algorithms for estimating the trace of an implicit symmetric positive semi-definite matrix. *J. ACM*, 58(2):8:1–8:34, 2011.
- Boutsidis, C., Drineas, P., Kambadur, P., Kontopoulou, E.-M., and Zouzias, A. A randomized algorithm for approximating the log determinant of a symmetric positive definite matrix. *Linear Algebra and its Applications*, 533: 95 – 117, 2017.
- Chang, B., Meng, L., Haber, E., Ruthotto, L., Begert, D., and Holtham, E. Reversible architectures for arbitrarily deep residual neural networks. *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- Chen, T. Q., Rubanova, Y., Bettencourt, J., and Duvenaud, D. Neural ordinary differential equations. *Advances in Neural Information Processing Systems*, 2018.
- Ciccone, M., Gallieri, M., Masci, J., Osendorfer, C., and Gomez, F. Nais-net: Stable deep networks from non-autonomous differential equations. In *Advances in Neural Information Processing Systems 31*, pp. 3029–3039. 2018.
- Clevert, D., Unterthiner, T., and Hochreiter, S. Fast and accurate deep network learning by exponential linear units (elus). *arXiv preprint arXiv:1511.07289*, 2015.
- Dinh, L., Krueger, D., and Bengio, Y. Nice: Non-linear independent components estimation. *arXiv preprint arXiv:1410.8516*, 2014.
- Dinh, L., Sohl-Dickstein, J., and Bengio, S. Density estimation using real nvp. *International Conference on Learning Representations*, 2017.
- Dong, K., Eriksson, D., Nickisch, H., Bindel, D., and Wilson, A. G. Scalable log determinants for gaussian process kernel learning. In *Advances in Neural Information Processing Systems 30*, pp. 6327–6337. 2017.
- Germain, M., Gregor, K., Murray, I., and Larochelle, H. MADE: Masked autoencoder for distribution estimation. In *ICML*, volume 37 of *JMLR Workshop and Conference Proceedings*, pp. 881–889. JMLR.org, 2015.
- Gomez, A. N., Ren, M., Urtasun, R., and Grosse, R. B. The reversible residual network: Backpropagation without storing activations. *Advances in Neural Information Processing Systems*, 2017.
- Gouk, H., Frank, E., Pfahringer, B., and Cree, M. Regularisation of neural networks by enforcing lipschitz continuity. *arXiv preprint arXiv:1804.04368*, 2018.
- Grathwohl, W., Chen, R. T. Q., Bettencourt, J., and Duvenaud, D. Ffjord: Scalable reversible generative models with free-form continuous dynamics. In *International Conference on Learning Representations*, 2019.
- Haber, E. and Ruthotto, L. Stable architectures for deep neural networks. *Inverse Problems*, 34(1):014004, 2018.
- Hall, B. C. Lie groups, lie algebras, and representations: An elementary introduction. *Graduate Texts in Mathematics, 222 (2nd ed.)*, Springer, 2015.
- Han, I., Malioutov, D., Avron, H., and Shin, J. Approximating the spectral sums of large-scale matrices using chebyshev approximations. *SIAM Journal on Scientific Computing*, 39, 06 2016.
- Han, I., Avron, H., and Shin, J. Stochastic chebyshev gradient descent for spectral optimization. In *Advances in Neural Information Processing Systems 31*, pp. 7397–7407. 2018.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Hjelm, R. D., Fedorov, A., Lavoie-Marchildon, S., Grewal, K., Bachman, P., Trischler, A., and Bengio, Y. Learning deep representations by mutual information estimation and maximization. In *International Conference on Learning Representations*, 2019.
- Hutchinson, M. A stochastic estimator of the trace of the influence matrix for laplacian smoothing splines. *Communications in Statistics - Simulation and Computation*, 19(2):433–450, 1990.

- Jacobsen, J.-H., Smeulders, A. W., and Oyallon, E. i-revnet: Deep invertible networks. In *International Conference on Learning Representations*, 2018.
- Jacobsen, J.-H., Behrmann, J., Zemel, R., and Bethge, M. Excessive invariance causes adversarial vulnerability. In *International Conference on Learning Representations*, 2019.
- Karras, T., Laine, S., and Aila, T. A style-based generator architecture for generative adversarial networks. *arXiv preprint arXiv:1812.04948*, 2018.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Kingma, D. P. and Dhariwal, P. Glow: Generative flow with invertible 1x1 convolutions. *Advances in Neural Information Processing Systems*, 2018.
- Kingma, D. P., Mohamed, S., Rezende, D. J., and Welling, M. Semi-supervised learning with deep generative models. In *Advances in neural information processing systems*, pp. 3581–3589, 2014.
- Louizos, C., Swersky, K., Li, Y., Welling, M., and Zemel, R. The variational fair autoencoder. *International Conference on Learning Representations*, 2016.
- Lu, Y., Zhong, A., Li, Q., and Dong, B. Beyond finite layer neural networks: Bridging deep architectures and numerical differential equations. *arXiv preprint arXiv:1710.10121*, 2017.
- Miyato, T., Kataoka, T., Koyama, M., and Yoshida, Y. Spectral normalization for generative adversarial networks. In *International Conference on Learning Representations*, 2018.
- Nalisnick, E., Matsukawa, A., Teh, Y. W., Gorur, D., and Lakshminarayanan, B. Hybrid models with deep and invertible features. *NeurIPS workshop on Bayesian deep learning*, 2018.
- Oliver, A., Odena, A., Raffel, C. A., Cubuk, E. D., and Goodfellow, I. Realistic evaluation of deep semi-supervised learning algorithms. In *Advances in Neural Information Processing Systems 31*. 2018.
- Papamakarios, G., Murray, I., and Pavlakou, T. Masked autoregressive flow for density estimation. In *Advances in Neural Information Processing Systems*, 2017.
- Parmar, N., Vaswani, A., Uszkoreit, J., Kaiser, L., Shazeer, N., Ku, A., and Tran, D. Image transformer. In *ICML*, volume 80 of *JMLR Workshop and Conference Proceedings*, pp. 4052–4061. JMLR.org, 2018.
- Ramesh, A. and LeCun, Y. Backpropagation for implicit spectral densities. *arXiv preprint arXiv:1806.00499*, 2018.
- Rezende, D. J. and Mohamed, S. Variational inference with normalizing flows. *Proceedings of the 32nd International Conference on International Conference on Machine Learning*, 2015.
- Ruthotto, L. and Haber, E. Deep neural networks motivated by partial differential equations. *arXiv preprint arXiv:1804.04272*, 2018.
- Schott, L., Rauber, J., Bethge, M., and Brendel, W. Towards the first adversarially robust neural network model on MNIST. In *International Conference on Learning Representations*, 2019.
- Sedghi, H., Gupta, V., and Long, P. M. The singular values of convolutional layers. In *International Conference on Learning Representations*, 2019.
- Sokoli, J., Giryes, R., Sapiro, G., and Rodrigues, M. R. D. Robust large margin deep neural networks. *IEEE Transactions on Signal Processing*, 65(16):4265–4280, 2017.
- Tsuzuku, Y., Sato, I., and Sugiyama, M. Lipschitz-margin training: Scalable certification of perturbation invariance for deep neural networks. *Advances in Neural Information Processing Systems*, 2018.
- Van Den Oord, A., Dieleman, S., Zen, H., Simonyan, K., Vinyals, O., Graves, A., Kalchbrenner, N., Senior, A., and Kavukcuoglu, K. Wavenet: A generative model for raw audio. *arXiv preprint arXiv:1609.03499*, 2016a.
- Van Den Oord, A., Kalchbrenner, N., and Kavukcuoglu, K. Pixel recurrent neural networks. In *Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48*, pp. 1747–1756, 2016b.
- Van Den Oord, A., Li, Y., and Vinyals, O. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.
- Withers, C. S. and Nadarajah, S. $\log \det a = \text{tr} \log a$. *International Journal of Mathematical Education in Science and Technology*, 41(8):1121–1124, 2010.
- Zagoruyko, S. and Komodakis, N. Wide residual networks. *arXiv preprint arXiv:1605.07146*, 2016.