
On the Universality of Volume-Preserving and Coupling-Based Normalizing Flows

Felix Draxler¹ Stefan Wahl¹ Christoph Schnörr¹ Ullrich Köthe¹

Abstract

We present a novel theoretical framework for understanding the expressive power of normalizing flows. Despite their prevalence in scientific applications, a comprehensive understanding of flows remains elusive due to their restricted architectures. Existing theorems fall short as they require the use of arbitrarily ill-conditioned neural networks, limiting practical applicability. We propose a distributional universality theorem for well-conditioned coupling-based normalizing flows such as RealNVP (Dinh et al., 2017). In addition, we show that volume-preserving normalizing flows are not universal, what distribution they learn instead, and how to fix their expressivity. Our results support the general wisdom that affine and related couplings are expressive and in general outperform volume-preserving flows, bridging a gap between empirical results and theoretical understanding.

1. Introduction

Density estimation and generative modeling of complex distributions is a fundamental problem in statistics and machine learning, with applications ranging from computer vision (Kingma & Dhariwal, 2018) to thermodynamic systems (Noé et al., 2019; Albergo et al., 2019; Nicoli et al., 2021) and uncertainty quantification (Ardizzone et al., 2018b).

Normalizing flows are a class of generative models that learn a probability density $p_\theta(x)$ from samples $x \sim p(x)$ or from a potentially unnormalized ground truth density $\hat{p}(x) \propto p(x)$. They are implemented by transporting a simple multivariate base density such as the standard normal via a learned invertible function to the distribution of interest.

Constructing flexible and tractable invertible neural net-

works is nontrivial and a significant body of work has developed a plethora of architectures, see (Kobyzev et al., 2021) for an overview. The evidence for which architecture to choose in practice is mostly limited to empirical results, however. In this work, we prove rigorous results regarding the universality of the families of volume-preserving and coupling-based normalizing flows.

First, we consider volume-preserving flows such as (Dinh et al., 2015; Sorrenson et al., 2019; Toth et al., 2020). Volume-preservation can be useful in certain applications, such as disentanglement (Sorrenson et al., 2019) or learning distributions with controllable temperature $p(x|T) \propto (\hat{p}(x))^T$ (Dibak et al., 2022). However, we show that they are not universal and learn a biased distribution in practice. We also provide a simple solution to restore universality, by adding a single one-dimensional non-volume-preserving layer.

Second, we improve on the universality theory of coupling-based normalizing flows (not preserving volume). These flows are a particularly efficient variant of parameterizing invertible neural networks, with fast training and inference. Despite their seemingly strong architectural constraints, in practice even the simple affine coupling-based normalizing flow (Dinh et al., 2017) can learn high-dimensional distributions such as images (Kingma & Dhariwal, 2018).

Theoretical explanations for this architecture’s ability to fit complex distributions are limited. Existing proofs make assumptions that are not valid in practice, as the involved constructions rely on ill-conditioned neural networks (Teshima et al., 2020a; Koehler et al., 2021) or construct a volume-preserving flow (Koehler et al., 2021). We introduce a new proof for the distributional universality of coupling-based normalizing flows that does not require ill-conditioned neural networks to converge. This proof is constructive, showing that training affine coupling blocks sequentially converges to the correct target (compare Figure 1).

In summary, we contribute:

- We show for the first time that volume-preserving flows are not universal, derive what distribution they converge to instead and provide simple fixes for their shortcomings in Section 4.

¹Heidelberg University, Germany. Correspondence to: Felix Draxler <felix.draxler@iwr.uni-heidelberg.de>.

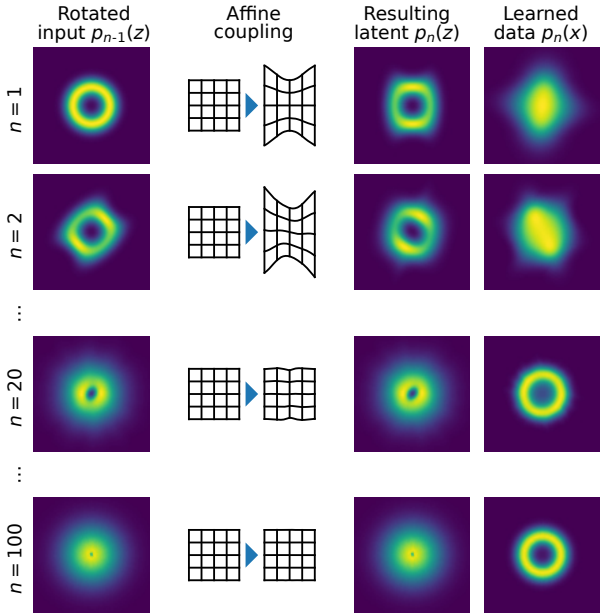


Figure 1. Our universality proof constructs a normalizing flow by iteratively adding affine coupling blocks. We illustrate this by constructing such a flow from topologically challenging toy data. Starting with the input $p(x)$, each block first rotates the latent distribution $p_{n-1}(z)$ from the previous step (first column), then applies an affine coupling layer that transforms the active dimensions to zero mean and unit variance for each passive coordinate b (second column). The resulting latent distribution $p_n(z)$ converges step by step (third column) to a standard normal distribution, until the learned additional layers essentially learn the identity (last row). The learned data distribution $p_\theta(x)$ converges in parallel (right).

- We show that whenever the target distribution is not perfectly learned, there is an affine coupling block that reduces the loss (Section 5.2).
- We use this result to give a new universality proof for coupling-based normalizing flows that is not volume-preserving, considers the full support of the distribution, and is not ill-conditioned in Section 5.3.

Our results validate insights previously observed only empirically: Affine coupling blocks are an effective foundation for normalizing flows, and volume-preserving flows have limited expressive power. We also show that the most recent distributional universality proof for affine coupling-based normalizing flows by Koehler et al. (2021) constructs such a volume-preserving flow in Section 5.1.

2. Related Work

Normalizing flows are a class of generative models based on invertible neural networks (Rezende & Mohamed, 2015). We focus on the analytical expressivity of volume-

Table 1. Our construction of a universal coupling flow overcomes important limitations of the previous work on arbitrary input $p(x)$ (Teshima et al., 2020a; Koehler et al., 2021): We use well-conditioned coupling blocks, consider convergence on the full space and allow variable volume change $|f'_\theta(x)| \neq \text{const}$, which is necessary for universality in KL divergence by Theorem 4.2.

	Teshima et al.	Koehler et al.	Thm. 5.4
Well-conditioned	✗	✗	✓
Variable $ f'_\theta(x) $	✓	✗	✓
Global support	✗	✗	✓

preserving and coupling-based normalizing flows, see Section 3.

That coupling-based normalizing flows work well in practice despite their restricted architecture has sparked the interest of several papers analyzing their distributional universality, i.e. the question whether they can approximate any target distribution to arbitrary precision (see Definition 3.1). Teshima et al. (2020a) showed that coupling flows are universal approximators for invertible functions, which results in distributional universality. Koehler et al. (2021) demonstrated that affine coupling-based normalizing flows can approximate any distribution with arbitrary precision using just three coupling blocks. However, these works rely on ill-conditioned coupling blocks and consider convergence only on a bounded subspace. Our work addresses these limitations and shows that training a normalizing flow layer by layer yields universality. In addition, we find that Koehler et al. (2021) effectively construct a volume-preserving flow, which we show not to be universal in KL divergence, the loss used in practice. This line of work is summarized in Table 1.

Some works show distributional universality of augmented affine coupling-based normalizing flows, which add at least one additional dimension usually filled with exact zeros (Huang et al., 2020; Koehler et al., 2021; Lyu et al., 2022). The problem with adding additional zeros is that the flow is not exactly invertible anymore in the data domain and usually loses tractability of the change of variables formula (Equation (1)). Lee et al. (2021) add i.i.d. Gaussians as additional dimensions, which again allows density estimation, but they only show how to approximate the limited class of log-concave distributions. Our universality proof does not rely on such a construction.

Other theoretical work on the expressivity of normalizing flows considers more expressive invertible neural networks, including SoS polynomial flows, Neural ODEs and Residual Neural Networks (Jaini et al., 2019; Zhang et al., 2020; Teshima et al., 2020b; Ishikawa et al., 2022). Another line of work found that the number of required coupling blocks

is independent of dimension D for Gaussian distributions compared to $O(D)$ Gaussianization blocks that lack couplings between dimensions (Koehler et al., 2021; Draxler et al., 2022; 2023).

Regarding volume-preserving flows, to the best of our knowledge there is no previous work showing their non-universality. Our results also complement the fact that Hamiltonian Monte Carlo (HMC) resamples momenta at every step in order to sample the correct target distribution (Neal, 2011).

3. Background

Normalizing Flows are a class of generative models that represent a distribution $p_\theta(x)$ with parameters θ by learning an invertible function $z = f_\theta(x)$ so that the latent codes $z \in \mathbb{R}^D$ obtained from the data $x \in \mathbb{R}^D$ are distributed like a standard normal distribution $p(z) = \mathcal{N}(z; 0, I)$. Via the change of variables formula, see (Köthe, 2023) for a review, this invertible function yields an explicit form for the density $p_\theta(x)$:

$$p_\theta(x) = p(z = f_\theta(x)) |f'_\theta(x)|, \quad (1)$$

where $f'_\theta(x) = \frac{\partial}{\partial x} f_\theta(x)$ is the Jacobian matrix of f_θ at x and $|f'_\theta(x)|$ is its absolute determinant. Note that we consistently denote by $p(\cdot)$ a ground truth target distribution and by $p_\theta(\cdot)$ a learned approximation.

Equation (1) allows easily evaluating the model density at points of interest. Obtaining samples from $p_\theta(x)$ can be achieved by sampling from the latent standard normal and applying the inverse $f_\theta^{-1}(z)$ of the learned transformation:

$$x = f_\theta^{-1}(z) \sim p_\theta(x) \text{ for } z \sim p(z). \quad (2)$$

The change of variables formula (Equation (1)) can be used directly to train a normalizing flow. The corresponding loss minimizes the Kullback-Leibler divergence between the true data distribution $p(x)$ and the learned distribution, which can be optimized via a Monte-Carlo estimate of the involved expectation:

$$\mathcal{L} = \mathcal{D}_{\text{KL}}(p(x) \| p_\theta(x)) \quad (3)$$

$$= \mathbb{E}_{x \sim p(x)} [\log p(x) - \log p_\theta(x)] \quad (4)$$

$$= \mathbb{E}_{x \sim p(x)} [-\log p_\theta(x)] + \text{const.} \quad (5)$$

This last variant makes clear that minimizing this loss is exactly the same as maximizing the log-likelihood of the training data. For training, the expectation value is approximated using (batches of) training samples $x_{1, \dots, N}$.

In order for Equations (1) and (2) to be useful in practice, $f_\theta(x)$ must have (i) a tractable inverse $f_\theta^{-1}(z)$ for fast sampling, and (ii) a tractable Jacobian determinant $|f'_\theta(x)|$ for

fast training while (iii) being expressive enough to model complicated distributions. These constraints are nontrivial to fulfill at the same time and significant work has been put into constructing such invertible neural networks, see (Kobyzev et al., 2021).

Coupling-Based Normalizing Flows are an invertible neural network design that lies in a sweet spot of scaling well to large dimensions yet remaining fast to sample from (Draxler et al., 2023) and exhibits a tractable Jacobian determinant. Its basic building block is the coupling layer, which consists of one invertible function $\tilde{x}_i = c(x_i; \psi_i)$ for each dimension, but with a twist: Only the second half of the dimensions $a = x_{D/2+1, \dots, D}$ (*active*) is changed in a coupling layer, and the parameters of this transformation $\psi = \psi(b)$ are predicted by a function called the *conditioner* that takes the first half of dimensions $b = x_{1, \dots, D/2}$ (*passive*) as input:

$$\tilde{x}_i = (f_{\text{cpl}}(x))_i = \begin{cases} b_i & i \leq D/2 \\ c(a_{i-D/2}; \psi_{i-D/2}(b)) & \text{else.} \end{cases} \quad (6)$$

In practice, the conditioner of each block is realized by a neural network $\psi_\varphi(b)$ with parameters φ . Calculating the inverse of the coupling layer is easy, as $b = \tilde{b}$ for the passive dimensions. This allows computing the parameters $\psi(b)$ necessary to invert the active half of dimensions:

$$x_i = (f_{\text{cpl}}^{-1}(\tilde{x}))_i = \begin{cases} \tilde{b}_i & i \leq D/2 \\ c^{-1}(\tilde{a}_{i-D/2}; \psi_{i-D/2}(\tilde{b})) & \text{else.} \end{cases} \quad (7)$$

Choosing the right one-dimensional invertible function $c(x; \psi)$ is the subject of active research, see our list in Appendix A.1 and Kobyzev et al. (2021). Many applications use affine-linear functions (Dinh et al., 2017):

$$c(x; s, t) = sx + t, \quad (8)$$

where $s > 0$ and t are predicted by the conditioner $\psi(b)$ as a function of the passive dimensions. Especially for smaller-dimensional problems it has proven useful to use more flexible c such as rational-quadratic splines (Durkan et al., 2019b). The positive and negative universality theorems in this paper apply to all coupling architectures we are aware of. At the same time, we give a direct reason for using more expressive couplings, as they can learn the same distributions with fewer layers in Section 5.4.

In order to be expressive, a coupling flow consists of a stack of coupling layers, each with a different active and passive subspace. Varying the subspaces is realized by an additional layer before each coupling which mixes dimensions by a rotation matrix $Q \in SO(D)$:

$$f_{\text{rot}}(x) = Qx, \quad f_{\text{rot}}^{-1}(\tilde{x}) = Q^T \tilde{x}. \quad (9)$$

The matrix Q can either be chosen to be a hard permutation or any matrix with orthonormal columns, which mixes dimensions linearly. The latter can optionally be learned (Kingma & Dhariwal, 2018). From a representational perspective, these variants are interchangeable because any soft permutation can be represented by a constant number of coupling blocks with hard permutations (Koehler et al., 2021).

A rotation layer together with a coupling layer forms a coupling block:

$$f_{\text{blk}}(x) = (f_{\text{cpl}} \circ f_{\text{rot}})(x) = f_{\text{cpl}}(Qx). \quad (10)$$

In the Section 5, we are concerned with what distributions $p(x)$ a potentially deep concatenation of coupling blocks can represent. Taken together, the parameters of the conditioners of the coupling layers φ together with the rotation matrices Q make up the learnable parameters θ of the entire normalizing flow.

Volume-Preserving Normalizing Flows or sometimes *incompressible flows* are a variant of normalizing flows that have a constant Jacobian determinant $|f'_\theta(x)| = \text{const}$. This simplifies the change of variables formula in Equation (1) above, where $C = |f'_\theta(x)|$:

$$p_\theta(x) = p(z = f_\theta(x))C. \quad (11)$$

Volume-preserving flows have been demonstrated to have useful properties in certain applications such as disentanglement (Sorensen et al., 2019), or temperature-scaling in Boltzmann generators (Dibak et al., 2022) or to preserve volume in physical state-space (Toth et al., 2020). However, we show that the volume-preserving change of variables in Equation (11) does not allow for universal normalizing flows regardless of the architecture in Section 4.

For one-dimensional functions, a constant volume change implies that $f_\theta(x) = Cx + t$ is linear. For multivariate functions, $f_\theta(x)$ can be nonlinear, only that any volume change in one dimension must be compensated by an inverse volume change in the remaining dimensions. Prominent implementations are nonlinear independent components estimation (NICE) (Dinh et al., 2015), general incompressible-flow networks (GIN) (Sorensen et al., 2019) or Neural Hamiltonian Flows (NHF) (Toth et al., 2020). For example, GIN realizes volume-preserving coupling blocks by ensuring $\sum_{i=1}^{D/2} \log s_i(b) = \text{const}$. We list common volume-preserving constructions in Appendix A.2. Note that “volume-preserving” is strictly speaking a misnomer when $C \neq 1$, but the term is commonly used also in this more general case. In Lemma B.3, we show that non-unit volume change can be absorbed into a single scaling layer.

Distributional Universality means that a certain class of generative models can represent any distribution $p(x)$. Due

to the nature of neural networks, we cannot hope for our generative model to *exactly* (i.e. exact equality in the mathematical sense) represent $p(x)$. This becomes clear via an analogue in the context of regression: A neural network with ReLU activations always models piecewise linear functions, and as such it can never *exactly* regress a parabola $y = x^2$. However, for every finite value of $\epsilon > 0$ and given more and more linear pieces, it can follow the parabola ever so closer, so that the average distance between x^2 and $f_\theta(x)$ vanishes: $\mathbb{E}_{x \sim p(x)} [|x^2 - f_\theta(x)|^2] < \epsilon$. To characterize the expressivity of a class of neural networks, it is thus instructive to call a class of networks universal if the error between the model and any target can be reduced arbitrarily.

In terms of representing distributions $p(x)$, the following definition captures universality of a class of model distributions, similar to (Teshima et al., 2020a, Definition 3):

Definition 3.1. A set of probability distributions \mathcal{P} is called a *distributional universal approximator* if for every possible target distribution $p(x)$ there is a sequence of distributions $p_n(x) \in \mathcal{P}$ such that $p_n(x) \xrightarrow{n \rightarrow \infty} p(x)$.

The formulation of universality as a convergent series is useful as it (i) captures that the distribution in question $p(x)$ may not lie in \mathcal{P} , and (ii) the series index n usually reflects a hyperparameter of the underlying model corresponding to computational requirements (for example, the depth of the network).

We have left the exact definition of the limit “ $p_n(x) \xrightarrow{n \rightarrow \infty} p(x)$ ” open as we may want to consider different variations of convergence. The existing literature on affine coupling-based normalizing flows considers weak convergence (Teshima et al., 2020a) respectively convergence in Wasserstein distance (Koehler et al., 2021). Many metrics of convergence have been proposed, see Gibbs & Su (2002) for a systematic overview.

While we consistently state our assumptions, we usually restrict ourselves to data distributions with densities $p(x)$ that are bounded and continuous, and that have infinite support and finite moments, which covers distributions of practical interest.

4. Non-Universality of Volume-Preserving Flows

In this section, we consider normalizing flows with constant Jacobian determinant $|f'_\theta(x)| = \text{const}$. We show that **volume-preserving flows are not universal in KL divergence**. We then propose how universality can be recovered.

From the volume-preserving change-of-variables in Equation (11) we can derive the KL divergence $\mathcal{D}_{\text{KL}}(p(x)||p_\theta(x))$

in the special case of $C = 1$:

$$\mathcal{L} = \int p(x) \log \frac{p(x)}{p_\theta(x)} dx \quad (12)$$

$$= -H[p(x)] - \int p(x) \log p(z = f_\theta(x)) dx. \quad (13)$$

Only the last term depends on f_θ . To derive the minimizer, consider the data $p(x)$ and latent distribution $p(z)$ on a regular grid over \mathbb{R}^D with some spacing $a > 0$. Then, define a volume-preserving flow with $C = 1$ that permutes the grid cells $B_i \mapsto B_{f(i)}$ (within the cells, keep the relative positions). Then, discretize the above integral on the grid by approximating the latent probability by the average density in each cell, that is $p(z) \approx \frac{1}{a^D} p(B_i : z \in B_i)$:

$$- \int p(x) \log p(z = f_\theta(x)) dx \quad (14)$$

$$\approx - \sum_i p(x \in B_{s_x(i)}) \log(p(z \in B_{f(i)})/|a^D|). \quad (15)$$

This is minimized by a bijective $f^* : \mathbb{N} \rightarrow \mathbb{N}$ that permutes the grid cells such that the cell with the highest probability $p(x \in B_i)$ in the data space aligns with the cell with the highest (logarithmic) probability in latent space, and so on:

$$f^*(i) = s_z(s_x^{-1}(i)), \quad (16)$$

where $s_v(i)$ is a sorting of the grid cells, determined by probability mass $p(v \in B_i)$ for $v = x$ respectively $v = z$.

The following theorem makes the above argument continuous and determines the optimal volume change $C > 0$:

Theorem 4.1. *Given a continuous bounded input density $p(x)$. Then, for any volume-preserving flow $p_\theta(x)$ with a standard normal latent distribution, the achievable KL divergence is bounded from below:*

$$\mathcal{D}_{KL}(p(x) \| p_\theta(x)) \geq \mathcal{D}_{KL}(p^*(z) \| \mathcal{N}(0, |\Sigma_{p^*(z)}|^{\frac{1}{D}} I)), \quad (17)$$

where $p^*(z)$ is constructed by decreasingly sorting the probability densities $p(x)$ from the origin with unit volume change, and $\Sigma_{p^*(z)}$ is its covariance matrix. The minimal loss is achieved for $C = |\Sigma_{p^*(z)}|^{-\frac{1}{2D}}$.

The optimal $p^*(x)$ and its latent counterpart $p^*(z)$ are constructed by sorting both the data and latent space by density and progressively assigning regions of decreasing density to each other (see proof in Appendix B.1). Figure 2 shows how this optimal distribution $p^*(x)$ differs from the target $p(x)$ for a bimodal toy distribution in 2D.

The following result formalizes that volume-preserving flows are not universal:

Theorem 4.2. *The family of normalizing flows with constant Jacobian determinant $|f'_\theta(x)| = \text{const}$ is not a universal distribution approximator under KL divergence.*

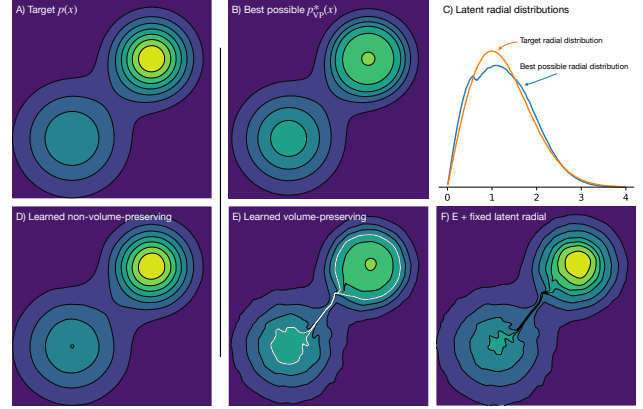


Figure 2. We reveal two limitations of volume-preserving flows: First, a 2D bimodal distribution (A) cannot be represented by a volume-preserving flow, the theoretic optimum predicted by Theorem 4.1 assigns wrong densities to both modes (B). This is because the radial part of the latent distribution $p^*(z)$ does not match radial part of the standard normal (C). In practice, learning a volume-preserving flow comes very close to the biased solution (E). A normalizing flow with variable Jacobian determinant does not have this issue (D). Our proposed fix corrects the densities at the modes (F) by correcting the latent radials, see Appendix B.2. Second, since the flow is continuous in practice, it cannot represent multi-modal distributions by Proposition 4.3, but a vanishing density bridge connecting the modes remains (E, white level set).

The proof constructs a concrete $p(x)$ and shows that the KL is bounded from below by a finite value for every possible value of C (see Appendix B.3).

The construction underlying Theorem 4.1 shows a clear path to construct a universal volume-preserving flow: The best achievable latent distribution $p^*(z)$ (the push-forward of $p(x)$ through f^*) is rotationally symmetric due to the sorting procedure. Now transform both $p^*(z)$ and the target standard normal $p(z)$ into hyperspherical coordinates (r, Ω) . As both distributions are rotationally symmetric, only their radial parts $p^*(r)$ and $p(r)$ need to be matched, which can be achieved via the addition of a single one-dimensional *non*-volume-preserving transformation (see Appendix B.2). As this fix is one-dimensional, unique, and can be applied after training, we think that it is compatible with retaining beneficial properties of volume-preserving flows.

Figure 2 also shows another shortcoming of volume-preserving flows as they are practically implemented: There is a thin bridge of density between the modes with roughly constant height, so that the lower mode in the ground truth is not a local maximum of the learned density. The reason is that flows are implemented as continuous invertible functions (as opposed to Theorems 4.1 and 4.2, which only require invertibility). This makes the learned distribution

$p_\theta(x)$ inherit the mode structure of the latent $p(z)$:

Proposition 4.3. *A normalizing flow $p_\theta(x)$ based on diffeomorphisms $f_\theta(x)$ with constant Jacobian determinant $|f'_\theta(x)| = \text{const}$ has the same number of modes as the latent distribution $p(z)$.*

The proof in Appendix B.4 uses that diffeomorphisms map open sets to open sets, and thus the neighborhoods of density maxima in the latent space remain neighborhoods of density maxima in the data space. Note that the thin bridge connecting the modes can be made arbitrarily small by an expressive enough volume-preserving flow, so that the shortcoming in Proposition 4.3 does not manifest in a bias in the KL divergence in addition to Theorem 4.1.

Together, we identify a fundamental limitation for applications based on volume-preserving flows. It explains why RealNVP significantly outperforms NICE in practice (Dinh et al., 2017). Work using volume-preserving flows must take this limited expressivity and the resulting biases in the learned distributions into account. In Section 5.1, we show that this problem also applies to the most recent universality proof for coupling-based normalizing flows by Koehler et al. (2021).

5. Universality of Coupling-Based Normalizing Flows

In this section, we present our improved universality proof for non-volume-preserving coupling-based normalizing flows. It overcomes limitations of previous constructions and relies on the simple idea of iteratively training coupling blocks.

5.1. Problems with Existing Constructions

The existing proofs (Teshima et al., 2020a; Koehler et al., 2021) that affine and more expressive coupling flows are distributional universal approximators come with several limitations. In particular, their constructions use ill-conditioned coupling blocks as the translations $t(b)$ approximate step functions, as noted by Koehler et al. (2021). Also, they give guarantees only on a compact subspace $K \subset \mathbb{R}^D$, and Teshima et al. (2020a) use only one active dimension per coupling. This limits their practical applicability.

In addition, the flow constructed in Koehler et al. (2021) is volume-preserving and thus not universal in KL divergence by our Theorem 4.2. Their proof is technically accurate, but they only show convergence under Wasserstein distance W_2 in (Koehler et al., 2021, Theorem 1), which does not imply convergence in KL (see Appendix C.4).

It is easy to see that their flow is volume-preserving by looking at the scaling functions $s(b)$ in the three affine coupling layers (Equation (8)) they use. They read: $s^{(1)} = \epsilon'$

and $s^{(2)} = s^{(3)} = \epsilon''$. This means that the overall flow has a Jacobian determinant $|f'_\theta| = (\epsilon' \epsilon''^2)^{\frac{D}{2}}$. This volume change is independent of the input, making the flow volume-preserving.

Also note how the volume change is directly tied to the guaranteed Wasserstein distance, since they guarantee that $W_2(p_\theta(x), p(x)) < \epsilon$, and the above scalings fulfill $\epsilon'' \ll \epsilon' \ll \epsilon$. Thus, the volume change $|f'_\theta| \ll \epsilon^{\frac{3D}{2}}$ vanishes, and its inverse $|(f^{-1})'_\theta| \gg \epsilon^{-\frac{3D}{2}}$ explodes as ϵ is reduced, rendering the flow ill-conditioned regardless of the distribution at hand. This is additional to the ill-condition of the translation terms $t(b)$ approximating step functions.

Together, this calls for a new universality guarantee that is based on a new coupling flow construction. Our new construction, presented in the following sections, uses a flow that is neither arbitrarily ill-conditioned nor volume-preserving. Also, it converges globally and considers vanilla affine coupling blocks.

5.2. Affine Coupling Blocks have a Unique Fixed Point

To construct our new universality theorem, we first analyze the effect of a single affine coupling on learning a target distribution $p(x)$. Our main result is that an affine coupling block always reduces the loss when it has not yet perfectly learned the target distribution.

To derive this, consider the *pushforward* of $p(x)$ through the flow $f_\theta(x)$, that is the latent distribution the flow actually creates by mapping $x \mapsto z = f_\theta(x)$:

$$p_\theta(z) = (f_{\theta\#}p(x))(z) = p(x = f_\theta^{-1}(z)) |(f_\theta^{-1})'(z)|. \quad (18)$$

Now rewrite the loss \mathcal{L} in Equation (3) into a form which compares $p_\theta(z)$ to the target latent distribution $p(z)$:

$$\mathcal{L} = \mathcal{D}_{\text{KL}}(p(x) \| p_\theta(x)) \quad (19)$$

$$= \int p(x) \log \frac{p(x)}{p(z = f_\theta(x)) |f'_\theta(x)|} dx \quad (20)$$

$$= \int p(z) |f_\theta^{-1}'(z)| \log \frac{p(f_\theta^{-1}(z)) |f_\theta^{-1}'(z)|}{p(z)} dz \quad (21)$$

$$= \mathcal{D}_{\text{KL}}(p_\theta(z) \| p(z)). \quad (22)$$

This identity shows that the divergence between the true $p(x)$ and the model $p_\theta(x)$ can equally be measured in the latent space, via the KL divergence between the current latent distribution that the model generates as the pushforward $p_\theta(z)$ and the target latent distribution $p(z)$.

Let us now consider what happens if we append one more affine coupling block $f_{\text{blk};\theta_+}$ to an existing normalizing flow $f_\theta(x)$, resulting in a flow which we call $p_{\theta \cup \theta_+}(x)$. Let us choose the parameters of the additional coupling block θ_+ such that it maximally reduces the loss without changing

the previous parameters:

$$\min_{\theta_+} \mathcal{D}_{\text{KL}}(p_{\theta \cup \theta_+}(z) \| p(z)). \quad (23)$$

Let us formalize the above construction for later use:

Definition 5.1. Given a normalizing flow $p_\theta(x)$ on a continuous probability distribution $p(x)$ with finite first and second moment and $p(x) > 0$ everywhere. Then, we define the *loss improvement by an affine coupling block* as:

$$\begin{aligned} \Delta_{\text{affine}}(p_\theta(z)) \\ := \mathcal{D}_{\text{KL}}(p_\theta(z) \| p(z)) - \min_{\theta_+} \mathcal{D}_{\text{KL}}(p_{\theta \cup \theta_+}(z) \| p(z)), \end{aligned} \quad (24)$$

where $\theta_+ = (Q, \varphi)$ parameterizes a single L -bi-Lipschitz affine coupling block whose conditioner neural network ψ_φ has at least two hidden layers of finite width and ReLU activations and $p_\theta(z) = (f_{\theta\#}p)(z)$.

The coupling block is restricted to be bi-Lipschitz in order to be well-conditioned. This means that we can choose $L > 1$ such that $L^{-1} < \|f_{\text{cpl}}(x) - f_{\text{cpl}}(y)\| / \|x - y\| < L$, making both forward and inverse passes through each coupling well-conditioned. Choosing a smaller L will result in more coupling blocks, each more numerically stable. Note that we assume ReLU networks for mathematical convenience, but think that the definition is equivalent to versions with different activation functions.

Considering this loss improvement $\Delta_{\text{affine}}(p_\theta(z))$ is a useful quantity, since we can show that is directly related to convergence of the flow:

Theorem 5.2. *With the definitions from Definition 5.1:*

$$p_\theta(z) = \mathcal{N}(z; 0, I) \iff \Delta_{\text{affine}}(p_\theta(z)) = 0. \quad (25)$$

This is a nontrivial result: One might have thought that it is possible to end up in distributions such that an affine coupling block gets stuck and cannot improve on the loss. Instead, **if adding another coupling layer has no effect, the latent distribution has converged to a standard normal**. This unique fixed point allows using $\Delta_{\text{affine}}(p_\theta(z))$ as a convergence metric for our universality theorem in Section 5.3.

In the remainder of this section, we give a sketch of the proof of Theorem 5.2, with technical details moved to Appendix C.1.

We proceed as follows: First, we use an explicit form of the maximal loss improvement $\Delta_{\text{affine}}^*(p_\theta(z))$ for infinitely expressive affine coupling blocks (Draxler et al., 2020). Then, we show in Lemma 5.3 that convergence of these unrealistic networks is equivalent to convergence of finite ReLU networks. Finally, we show that $\Delta_{\text{affine}}(p_\theta(z)) = 0$ implies $p_\theta(z) = \mathcal{N}(z; 0, I)$. The other direction is trivial, since by $p_\theta(z) = \mathcal{N}(0, I)$, no loss improvement is possible.

If we assume for a moment that neural networks can exactly represent arbitrary continuous functions, then this hypothetical maximal loss improvement was computed by Draxler et al. (2020, Theorem 1). A single affine coupling block with a fixed rotation layer Q , in order to maximally reduce the loss, will standardize the data by normalizing the first two moments of the active half of dimensions $a = (Qx)_{D/2+1, \dots, D}$ conditioned on the passive half of dimensions $b = (Qx)_{1, \dots, D/2}$. The moments before the coupling

$$\mathbb{E}_{a_i|b}[a_i] = m_i(b), \quad \text{Var}_{a_i|b}[a_i] = \sigma_i(b) \quad (26)$$

are mapped to:

$$\mathbb{E}_{\tilde{a}_i|b}[\tilde{a}_i] = 0, \quad \text{Var}_{\tilde{a}_i|b}[\tilde{a}_i] = 1. \quad (27)$$

This is achieved via the following affine transformation, shifting the conditional mean to zero and scaling the conditional standard deviation to one:

$$\tilde{a}_i(a_i; b) = \frac{1}{\sigma_i(b)}(a_i - m_i(b)). \quad (28)$$

In terms of loss, this transformation can at most achieve the following loss improvement, with a contribution from each passive coordinate b :

$$\Delta_{\text{affine}}^*(p_\theta(z)) = \max_Q \mathbb{E}_b[S(b)], \quad (29)$$

where Q enters through $(b, a) = Qx$, and the expectation goes over the component-wise *non-Standardness* (Draxler et al., 2022):

$$\begin{aligned} S(b) &= \sum_{i=1}^{D/2} \mathcal{D}_{\text{KL}}(\mathcal{N}(m_i(b), \sigma_i(b)) \| \mathcal{N}(0, 1)) \\ &= \frac{1}{2} \sum_{i=1}^{D/2} \mathbb{E}_b \left[\underbrace{m_i^2(b)}_{(A)} + \underbrace{\sigma_i^2(b) - 1 - \log \sigma_i^2(b)}_{(B)} \right], \end{aligned} \quad (30)$$

With the asterisk, we denote that this improvement cannot necessarily be reached in practice with finitely sized and well-conditioned neural networks. More expressive coupling functions can reduce the loss stronger, see Section 5.4.

What loss improvement can be achieved if we go back to finite neural networks? It turns out that $\Delta_{\text{affine}}^*(p_\theta(z)) > 0$ is equivalent to the existence of a well-conditioned coupling block as in Definition 5.1 with $\Delta_{\text{affine}}(p_\theta(z)) > 0$:

Lemma 5.3. *Given a continuous probability density $p_\theta(z)$ on $z \in \mathbb{R}^k$. Then,*

$$\Delta_{\text{affine}}^*(p_\theta(z)) > 0 \iff \Delta_{\text{affine}}(p_\theta(z)) > 0. \quad (32)$$

This says that the events $\Delta_{\text{affine}}(p_\theta(z)) = 0$ and $\Delta_{\text{affine}}^*(p_\theta(z)) = 0$ can be used interchangeably. The equivalence comes from the fact that if $\Delta_{\text{affine}}^*(p_\theta(z)) > 0$, then

we can always construct a conditioner neural network that scales the conditional standard deviations closer to one or the conditional means closer to zero, reducing the loss. In the detailed proof in Appendix C.1.2 we also make use of a classical regression universal approximation theorem (Hornik, 1991) and ensure the additional coupling block is well-conditioned.

Finally, if the first two *conditional* moments of any latent distribution $p(z)$ are normalized for all rotations Q :

$$\mathbb{E}_{a_i|b}[a_i] = m_i(b) = 0, \quad \text{Var}_{a_i|b}[a_i] = \sigma_i(b) = 1, \quad (33)$$

then the distribution must be the standard normal distribution: $p(z) = \mathcal{N}(z; 0, I)$: Equation (31) enforces two characteristics of $p_\theta(z)$ that uniquely identify the standard normal distribution: (A) It must be rotationally symmetric, since $m_i(b) = 0$ for all Q holds only for rotationally symmetric distributions (Eaton, 1986). (B) This is term non-negative and zero only for $\sigma_i(b) = 1$ for all Q , which uniquely identifies the standard normal from all rotationally symmetric distributions (Bryc, 1995).

This concludes the proof sketch of Theorem 5.2 and we are now ready to present our universality result, employing $\Delta_{\text{affine}}(p_\theta(z))$ as a convergence metric.

5.3. Affine Coupling Flows Universality

We now confirm that affine coupling flows are a distributional universal approximator in terms of the convergence metric we derived in Section 5.2:

Theorem 5.4. *For every continuous $p(x)$ with finite first and second moment with infinite support, there is a sequence of normalizing flows $p_n(x)$ consisting of n L -bi-Lipschitz affine coupling blocks such that their latent distributions converge:*

$$p_n(z) \xrightarrow{n \rightarrow \infty} \mathcal{N}(z; 0, I), \quad (34)$$

in the sense that $\Delta_{\text{affine}}(p_n(z)) \xrightarrow{n \rightarrow \infty} 0$.

This means that with increasing depth, the latent distribution of the flow converges to the standard normal. The use of $\Delta_{\text{affine}}(p_n(z))$ as a convergence metric is justified by Theorem 5.2 that $p_n(z) = \mathcal{N}(0, I) \Leftrightarrow \Delta_{\text{affine}}(p_n(z)) = 0$.

The proof of Theorem 5.4 explicitly constructs a normalizing flow by following an iterative scheme. We start with the data distribution as our original guess for the latent distribution: $p_0(z) = p(x = z)$. Then, we repeatedly append individual affine coupling blocks $f_{\text{blk}}(x)$ consisting of a rotation Q and a coupling f_{cpl} , optimizing the new parameters to maximally reduce the loss as in Equation (23).

This series of coupling blocks converges: $\Delta_{\text{affine}}(p_\theta(z))$ measures how much adding each affine coupling block reduces the loss, but the total loss that can be reduced by

the concatenation of many blocks is bounded. Since improvements $\Delta_{\text{affine}}(p_\theta(z))$ are also non-negative, they must converge to zero for the sum to be finite (Rudin, 1976, Theorems 3.14 and 3.23). By Theorem 5.2, the fixed point of this procedure is a standard normal distribution in the latent space. We give the full proof in Appendix C.2.

Figure 1 shows an example for how Theorem 5.4 constructs the coupling flow in order to learn a toy distribution. The affine coupling flow is able to learn the distribution well, despite its difficult topology. Empirically, this is also true in terms of KL divergence: Figure 5 in Appendix E.1 shows the relation between $\Delta_{\text{affine}}(p_\theta(z))$ and the KL divergence for the flow, both of which decrease over the course of training.

Table 1 summarizes how our construction is closer to practice than previous work (Teshima et al., 2020a; Koehler et al., 2021): We only use well-conditioned L -bi-Lipschitz couplings, allow variable volume change $|f'_\theta(x)|$ (as evidenced by the rescaling term in Equation (28)) and consider the entire support of $p(x)$. We give further details on the sensitivity of $\Delta_{\text{affine}}(p_\theta(z))$ to volume-preserving transformations in Appendix B.1.3.

Limitations: Despite these advances, there are some properties we hope can be improved in the future: First, our construction shows that we can build a deep enough flow with arbitrary precision, but we have not exploited that blocks can be jointly optimized. Thus, while our construction shows universality of end-to-end training, we expect a flow trained this way to require fewer blocks than our iterative proof for the same performance.

Secondly, it is unclear how the convergence metric $\Delta_{\text{affine}}(\cdot)$ in Definition 5.1 is related to convergence in the loss used in practice, the KL divergence given in Equation (3). In practice, we find that constructing a coupling flow through iterative training converges in KL divergence (see Figure 5 in Appendix E.1), so we conjecture that our way of constructing a universal coupling flow converges in KL divergence. The reverse holds: We show in Corollary C.3 in Appendix C.3 that convergence in KL implies convergence under our new metric.

Finally, our proof gives no guarantee on the number of required coupling blocks to achieve a certain performance. Related work shows that the number of layers is constant with dimension for the special case of Gaussian data (Koehler et al., 2021; Draxler et al., 2022), but in practice is a hyperparameter that is to be tuned depending on the data and together with the complexity of the subnetworks. We hope that our contribution paves the way towards a full understanding of affine coupling-based normalizing flows.




Figure 3. Information-geometric view of couplings more expressive than affine: The conditional KL divergence $\mathcal{D}_{\text{KL}}(p(a|b) \parallel \mathcal{N}(0, I))$ can be split into two orthogonal KL divergences, the non-Standardness $S(b)$ sensitive to the first two moments, and the negentropy $J(b)$ sensitive to non-Gaussianity. Affine couplings only reduce $S(b)$, more expressive coupling also affect $J(b)$.

5.4. Expressive Coupling Flow Universality

The above Theorem 5.4 shows that affine couplings $c(a_i; \theta) = sa_i + t$ are sufficient for universal distribution approximation. As mentioned in Section 3, a plethora of more expressive coupling functions have been suggested, for example neural spline flows (Durkan et al., 2019b) that use monotone rational-quadratic splines as the coupling function. It turns out that by choosing the parameters in the right way, all coupling functions we are aware of can exactly represent an affine coupling, except for the volume-preserving variants, see Appendix A.1. For example, a rational quadratic spline can be parameterized as an affine function by using equidistant knots (a_k, \tilde{a}_k) such that $\tilde{a}_k = sa_k + t$ and fixing the derivative at each knot to s .

Thus, the universality of more expressive coupling functions follows immediately from Theorem 5.4, just like Ishikawa et al. (2022) extended their results from affine to more expressive couplings:

Corollary 5.5. *For every continuous $p(x)$ with finite first and second moment with infinite support, there is a sequence of normalizing flows $f_n(x)$ consisting of n coupling blocks with coupling functions at least as expressive as affine couplings such that:*

$$p_n(z) \xrightarrow{n \rightarrow \infty} \mathcal{N}(z; 0, I), \quad (35)$$

in the sense that $\Delta_{\text{affine}}(p_n(z)) \xrightarrow{n \rightarrow \infty} 0$.

Our proof of Theorem 5.4, constructed through layer-wise training, shows how more expressive coupling functions can outperform affine functions using the same number of blocks. Similar to the loss improvement for an affine coupling in Equation (24), let us compute the maximally possible loss improvement for an arbitrarily flexible coupling function:

$$\Delta_{\text{universal}}^* = \max_Q(\mathbb{E}_b[J(b) + S(b)]) \geq \Delta_{\text{affine}}^*, \quad (36)$$

where the expectation again goes over the passive coordinate $b = (Qz)_{1, \dots, D/2}$ and $z \sim p_\theta(z)$.

Here, the loss improvement additional to the non-Standardness $S(b)$ as given in Equation (31) is the conditional negentropy $J(b) = \sum_{i=1}^{D/2} \mathcal{D}_{\text{KL}}(p_\theta(a_i|b) \parallel \mathcal{N}(m_i(b), \sigma_i(b)))$, which measures the deviation of each active dimension from a Gaussian distribution with matching mean and variance. An affine coupling function $c(a_i; \theta) = sa_i + t$ doesn't influence this term, due to its symmetrical effect on both sides of the KL in $J(p)$ (Draxler et al., 2022, Lemma 1). More expressive coupling blocks, however, are able to tap on this loss component if the conditional distributions $p(a_i|b)$ are non-Gaussian, see Figure 3 for an example. Note that while a single affine coupling does not affect $J(b)$, subsequent blocks can because the overall loss is redistributed over the loss terms.

The impact of this gain likely varies with the dataset. For instance, in images, the distribution of one color channel of one pixel conditioned on the other color channels in the entire image, often shows a simple unimodal pattern with low negentropy. This may explain why affine coupling blocks are enough to learn the distribution of images (Kingma & Dhariwal, 2018). We give additional technical details on Equation (36) and the subsequent arguments in Appendix D.

6. Conclusion

Our new universality proofs show an intriguing hierarchy of the universality of different coupling blocks:

1. *Volume-preserving normalizing flows* are not universal in KL divergence and thus cannot learn all targets $p(x)$.
2. *Affine coupling flows* such as RealNVP (Dinh et al., 2017) are distributional universal approximators in terms of Δ_{affine} despite their seemingly restrictive architecture.
3. Coupling flows with *more expressive coupling functions* are also universal approximators, but they converge faster by tapping on an additional loss component in layer-wise training.

Our work theoretically grounds choosing coupling blocks for practical applications with normalizing flows, combined with their easy implementation and training and inference speed. We remove spurious constructions present in previous proofs and use a simple principle instead: Train a flow layer by layer.

Using volume-preserving flows may have negatively affected existing work. We show what distribution $p^*(x)$ they approximate instead of the true target $p(x)$ and propose how universality can be recovered by learning the actual latent distribution after training.

Impact Statement

This paper presents work whose goal is to advance the field of Machine Learning. There are many potential societal consequences of our work, none of which we feel must be specifically highlighted here.

Acknowledgements

This work is supported by Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy EXC-2181/1 - 390900948 (the Heidelberg STRUCTURES Cluster of Excellence). It is also supported by the Vector Stiftung in the project TRINN (P2019-0092) and the Carl-Zeiss-Stiftung. We thank Armand Rousselot and Peter Sorrenson for the fruitful discussions and feedback. We thank Vincent Souveton for the helpful comments regarding Neural Hamiltonian Flows and Hamiltonian Monte Carlo.

References

- Albergo, M. S., Kanwar, G., and Shanahan, P. E. Flow-based generative models for Markov chain Monte Carlo in lattice field theory. *Physical Review D*, 100(3):034515, August 2019. ISSN 2470-0010, 2470-0029. doi: 10.1103/PhysRevD.100.034515.
- Ardizzone, L., Bungert, T., Draxler, F., Köthe, U., Kruse, J., Schmier, R., and Sorrenson, P. Framework for Easily Invertible Architectures (FrEIA), 2018a.
- Ardizzone, L., Kruse, J., Rother, C., and Köthe, U. Analyzing Inverse Problems with Invertible Neural Networks. In *International Conference on Learning Representations*, 2018b.
- Bryc, W. *The Normal Distribution*, volume 100 of *Lecture Notes in Statistics*. Springer New York, New York, NY, 1995. ISBN 978-0-387-97990-8 978-1-4612-2560-7. doi: 10.1007/978-1-4612-2560-7.
- Cambanis, S., Huang, S., and Simons, G. On the theory of elliptically contoured distributions. *Journal of Multivariate Analysis*, 11(3):368–385, September 1981. ISSN 0047259X. doi: 10.1016/0047-259X(81)90082-8.
- Cardoso, J.-F. Dependence, Correlation and Gaussianity in Independent Component Analysis. *Journal of Machine Learning Research*, 4:1177–1203, 2003. ISSN 1532-4435.
- Chen, S. and Gopinath, R. Gaussianization. In Leen, T., Dietterich, T., and Tresp, V. (eds.), *Advances in Neural Information Processing Systems*, 2000.
- Dibak, M., Klein, L., Krämer, A., and Noé, F. Temperature steerable flows and Boltzmann generators. *Phys. Rev. Res.*, 4(4):L042005, October 2022. doi: 10.1103/PhysRevResearch.4.L042005.
- Dinh, L., Krueger, D., and Bengio, Y. NICE: Non-linear Independent Components Estimation. In *International Conference on Learning Representations, Workshop Track*, 2015.
- Dinh, L., Sohl-Dickstein, J., and Bengio, S. Density estimation using Real NVP. In *International Conference on Learning Representations*, 2017.
- Draxler, F., Schwarz, J., Schnörr, C., and Köthe, U. Characterizing the Role of a Single Coupling Layer in Affine Normalizing Flows. In *German Conference on Pattern Recognition*, 2020.
- Draxler, F., Schnörr, C., and Köthe, U. Whitening Convergence Rate of Coupling-based Normalizing Flows. In *Advances in Neural Information Processing Systems*, 2022.
- Draxler, F., Kühmichel, L., Rousselot, A., Müller, J., Schnoerr, C., and Koethe, U. On the Convergence Rate of Gaussianization with Random Rotations. In *International Conference on Machine Learning*, 2023.
- Durkan, C., Bekasov, A., Murray, I., and Papamakarios, G. Cubic-Spline Flows. In *International Conference on Machine Learning, Workshop Track*, 2019a. doi: 10.48550/ARXIV.1906.02145.
- Durkan, C., Bekasov, A., Murray, I., and Papamakarios, G. Neural Spline Flows. In *Advances in Neural Information Processing Systems*, 2019b.
- Eaton, M. L. A characterization of spherical distributions. *Journal of Multivariate Analysis*, 20(2):272–276, December 1986. ISSN 0047259X. doi: 10.1016/0047-259X(86)90083-7.
- Gibbs, A. L. and Su, F. E. On Choosing and Bounding Probability Metrics. *International Statistical Review / Revue Internationale de Statistique*, 70(3):419–435, 2002. ISSN 03067734, 17515823.
- Glorot, X. and Bengio, Y. Understanding the difficulty of training deep feedforward neural networks. In *International Conference on Artificial Intelligence and Statistics*, 2010.
- Harris, C. R., Millman, K. J., van der Walt, S. J., Gommers, R., Virtanen, P., Cournapeau, D., Wieser, E., Taylor, J., Berg, S., Smith, N. J., Kern, R., Picus, M., Hoyer, S., van Kerkwijk, M. H., Brett, M., Haldane, A., del Río, J. F., Wiebe, M., Peterson, P., Gérard-Marchant, P., Sheppard,

- K., Reddy, T., Weckesser, W., Abbasi, H., Gohlke, C., and Oliphant, T. E. Array programming with NumPy. *Nature*, 585(7825):357–362, 2020.
- Ho, J., Chen, X., Srinivas, A., Duan, Y., and Abbeel, P. Flow++: Improving Flow-Based Generative Models with Variational Dequantization and Architecture Design. In *International Conference on Machine Learning*, 2019.
- Hornik, K. Approximation capabilities of multilayer feed-forward networks. *Neural Networks*, 4(2):251–257, 1991. ISSN 08936080. doi: 10.1016/0893-6080(91)90009-T.
- Huang, C.-W., Krueger, D., Lacoste, A., and Courville, A. Neural Autoregressive Flows. In *International Conference on Machine Learning*, 2018.
- Huang, C.-W., Dinh, L., and Courville, A. Augmented Normalizing Flows: Bridging the Gap Between Generative Flows and Latent Variable Models. In *International Conference on Learning Representations, Workshop Track*, 2020.
- Hunter, J. D. Matplotlib: A 2D graphics environment. *Computing in Science & Engineering*, 9(3):90–95, 2007.
- Ishikawa, I., Teshima, T., Tojo, K., Oono, K., Ikeda, M., and Sugiyama, M. Universal approximation property of invertible neural networks. *arXiv preprint arXiv:2204.07415*, 2022.
- Jaini, P., Selby, K. A., and Yu, Y. Sum-of-Squares Polynomial Flow. In *International Conference on Machine Learning*, 2019.
- Kingma, D. P. and Ba, J. Adam: A Method for Stochastic Optimization, January 2017.
- Kingma, D. P. and Dhariwal, P. Glow: Generative flow with invertible 1x1 convolutions. In *Advances in Neural Information Processing Systems*, 2018.
- Kobyzev, I., Prince, S. J., and Brubaker, M. A. Normalizing Flows: An Introduction and Review of Current Methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(11):3964–3979, 2021.
- Koehler, F., Mehta, V., and Risteski, A. Representational aspects of depth and conditioning in normalizing flows. In *International Conference on Machine Learning*, 2021.
- Köthe, U. A review of change of variable formulas for generative modeling. *arXiv preprint arXiv:2308.02652*, 2023.
- Lee, H., Pabbaraju, C., Sevekari, A. P., and Risteski, A. Universal approximation using well-conditioned normalizing flows. In Ranzato, M., Beygelzimer, A., Dauphin, Y., Liang, P., and Vaughan, J. W. (eds.), *Advances in Neural Information Processing Systems*, volume 34, pp. 12700–12711. Curran Associates, Inc., 2021.
- Lyu, J., Chen, Z., Feng, C., Cun, W., Zhu, S., Geng, Y., Xu, Z., and Chen, Y. Universality of parametric Coupling Flows over parametric diffeomorphisms. *arXiv preprint arXiv:2202.02906*, 2022.
- McKinney, W. Data Structures for Statistical Computing in Python. In van der Walt, S. and Jarrod Millman (eds.), *9th Python in Science Conference*, 2010.
- Müller, T., McWilliams, B., Rousselle, F., Gross, M., and Novák, J. Neural Importance Sampling. *ACM Transactions on Graphics*, 38(5):1–19, 2019. ISSN 0730-0301. doi: 10.1145/3341156.
- Neal, R. M. MCMC Using Hamiltonian Dynamics. In Brooks, S., Gelman, A., Jones, G. L., and Meng, X.-L. (eds.), *Handbook of Markov Chain Monte Carlo*. Chapman and Hall/CRC, 2011.
- Nicoli, K. A., Anders, C. J., Funcke, L., Hartung, T., Jansen, K., Kessel, P., Nakajima, S., and Stornati, P. Estimation of Thermodynamic Observables in Lattice Field Theories with Deep Generative Models. *Physical Review Letters*, 126(3):032001, January 2021. ISSN 0031-9007, 1079-7114. doi: 10.1103/PhysRevLett.126.032001.
- Noé, F., Olsson, S., Köhler, J., and Wu, H. Boltzmann generators: Sampling equilibrium states of many-body systems with deep learning. *Science*, 365(6457):eaaw1147, 2019.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems*, 2019.
- Rezende, D. and Mohamed, S. Variational inference with normalizing flows. In *International Conference on Machine Learning*, 2015.
- Rudin, W. *Principles of Mathematical Analysis*. International Series in Pure and Applied Mathematics. McGraw-Hill, New York St. Louis San Francisco [etc.], third ed edition, 1976. ISBN 978-0-07-054235-8.
- Sorrenson, P., Rother, C., and Köthe, U. Disentanglement by nonlinear ICA with general incompressible-flow networks (GIN). In *International Conference on Learning Representations*, 2019.
- Souveton, V., Guillin, A., Jasche, J., Lavaux, G., and Michel, M. Fixed-kinetic Neural Hamiltonian Flows for enhanced interpretability and reduced complexity. In *International Conference on Artificial Intelligence and Statistics*, 2024.

- Teshima, T., Ishikawa, I., Tojo, K., Oono, K., Ikeda, M., and Sugiyama, M. Coupling-based Invertible Neural Networks Are Universal Diffeomorphism Approximators. In *Advances in Neural Information Processing Systems*, 2020a.
- Teshima, T., Tojo, K., Ikeda, M., Ishikawa, I., and Oono, K. Universal Approximation Property of Neural Ordinary Differential Equations. In *Advances in Neural Information Processing Systems, Workshop Track*, 2020b.
- The pandas development team. Pandas-dev/pandas: Pandas, February 2020.
- Toth, P., Rezende, D. J., Jaegle, A., Racanière, S., Botev, A., and Higgins, I. Hamiltonian generative networks. In *International Conference on Learning Representations*, 2020.
- Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S. J., Brett, M., Wilson, J., Millman, K. J., Mayorov, N., Nelson, A. R. J., Jones, E., Kern, R., Larson, E., Carey, C. J., Polat, İ., Feng, Y., Moore, E. W., VanderPlas, J., Laxalde, D., Perktold, J., Cimrman, R., Henriksen, I., Quintero, E. A., Harris, C. R., Archibald, A. M., Ribeiro, A. H., Pedregosa, F., van Mulbregt, P., and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020. doi: 10.1038/s41592-019-0686-2.
- Wehenkel, A. and Louppe, G. Unconstrained Monotonic Neural Networks. In *Advances in Neural Information Processing Systems*, 2019.
- Zhang, H., Gao, X., Unterman, J., and Arodz, T. Approximation Capabilities of Neural ODEs and Invertible Residual Networks. In *International Conference on Machine Learning*, 2020.
- Ziegler, Z. and Rush, A. Latent Normalizing Flows for Discrete Sequences. In *International Conference on Machine Learning*, 2019.

A. Architectures

A.1. Compatible Coupling Functions

The following lists all coupling functions $c(a; \psi)$ (see Equation (6) for its usage) we are aware of. Our universality guarantees Theorem 5.4 and Corollary 5.5 hold for all of them:

- **Affine coupling flows** as RealNVP (Dinh et al., 2017) and GLOW (Kingma & Dhariwal, 2018):

$$c(x; \psi) = sx + t. \quad (37)$$

Here, $\psi = [s; t] \in \mathbb{R}_+ \times \mathbb{R}$. Note that NICE ($s = 1$, Dinh et al. (2015)) and GIN ($\prod_{i=1}^{D/2} s_i = 1$, (Sorenson et al., 2019)) follow the same functional form, but are volume-preserving and thus not universal (see Section 4).

- **Nonlinear squared flow** (Ziegler & Rush, 2019):

$$c(x; \psi) = ax + b + \frac{c}{1 + (dx + h)^2}, \quad (38)$$

for $\psi = [a, b, c, d, h] \in \mathbb{R}_+ \times \mathbb{R}^4$. Choose $c = 0$ to obtain an affine coupling.

- **Flow++** (Ho et al., 2019):

$$c(x; \psi) = s\sigma^{-1} \left(\sum_{j=1}^K \pi_j \sigma \left(\frac{x - \mu_j}{\sigma_j} \right) \right) + t. \quad (39)$$

Here, $\psi = [s; t; (\pi_j, \mu_j, \sigma_j)_{j=1}^K] \in \mathbb{R}_+ \times \mathbb{R} \times (\mathbb{R} \times \mathbb{R} \times \mathbb{R}_+)^K$ and σ is the logistic function. Choose all $\pi_j = 0$ except for $\pi_1 = 1$, all $\mu_j = 0$ and all $\sigma_j = 1$ to obtain an affine coupling.

- **SOS polynomial flows** (Jaini et al., 2019):

$$c(x; \psi) = \int_0^x \sum_{\kappa=1}^k \left(\sum_{l=0}^r a_{l,\kappa} u^l \right)^2 du + t. \quad (40)$$

Here, $\psi = [t; (a_{l,\kappa})_{l,\kappa}] \in \mathbb{R} \times \mathbb{R}^{rk}$. Choose all $a_{l,\kappa} = 0$ except for $a_{1,0} = s$ to obtain an affine coupling.

- **Spline flows** in all variants: Cubic (Durkan et al., 2019a), piecewise-linear, monotone quadratic (Müller et al., 2019), and rational quadratic (Durkan et al., 2019b) splines. A spline is parameterized by knots θ with optional derivative information depending on the spline type, and c computes the corresponding spline function. Choose the spline knots as $y_i = sx_i + t$ for an affine coupling, choose the derivatives as $x'_i = s$ for an affine coupling.
- **Neural autoregressive flow** (Huang et al., 2018) use a feed-forward neural network to parameterize $c(x; \theta)$ by a feed-forward neural network. They show that a neural network is guaranteed to be bijective if all activation functions are strictly monotone and all weights positive. One can construct a ReLU network with a single linear region to obtain an affine coupling.
- **Unconstrained monotonic neural networks** (Wehenkel & Louppe, 2019) also use a feed-forward neural, but restrict it to have positive output. To obtain $c(x; \theta)$, this function is then numerically integrated with a learnable offset for $x = 0$. Choose a constant neural network to obtain an affine coupling.

A.2. Volume-Preserving Normalizing Flows

Below, we list the ways to construct volume-preserving flows we are aware of. Our non-universality results Theorems 4.1 and 4.2 and Proposition 4.3 hold for all of them:

- **Nonlinear independent components estimation** (NICE) (Dinh et al., 2015) is a coupling block:

$$c(x; \psi) = x + t. \quad (41)$$

Here, $\psi = t \in \mathbb{R}$.

- **General Incompressible-flow Networks (GIN)** (Sorenson et al., 2019) generalize NICE by allowing the individual dimensions to change volume, only the overall volume change is normalized:

$$c_i(x_i; \psi) = \frac{s_i}{\prod_{j=1}^{D/2} s_j} x_i + t_i, \quad (42)$$

Here, $\psi = [s, t] \in \mathbb{R}_+^{D/2} \times \mathbb{R}^{D/2}$ is jointly predicted for all active dimensions and then normalized as above.

- **Neural Hamiltonian Flows** (Toth et al., 2020) parameterize a Neural ODE as a Hamiltonian system:

$$\frac{dq}{dt} = \frac{\partial \mathcal{H}}{\partial p}, \quad \frac{dp}{dt} = -\frac{\partial \mathcal{H}}{\partial q} \quad (43)$$

The Hamiltonian $\mathcal{H}(p, q)$ is a real-valued function that is parameterized by a neural network. Its derivatives are obtained via automatic differentiation. The variant **Fixed-kinetic Neural Hamiltonian Flows** (Souveton et al., 2024) fixes the kinetic term of the Hamiltonian to $K(p) = \frac{1}{2} p^T M^{-1} p$, where the positive definite matrix M is learned, and learns the potential $V(q)$ via a neural network to obtain $\mathcal{H}(p, q) = K(p) + V(q)$. The solution to the above ODE is volume-preserving on $x = (p, q)$.

Note that some works employing volume-preserving flows such as Dibak et al. (2022); Souveton et al. (2024) consider *augmented flows* (Huang et al., 2020), where additional noise dimensions are padded $p(a|x)$ to the data distribution of interest $p(x)$. Then, the flow learns the joint distribution $p(x, a) = p(x)p(a|x)$. Depending on how $p(a|x)$ is constructed, this can positively or negatively impact the expressivity of the considered volume-preserving flow. For example, if $a \perp x$, then the joint distribution $p(x, a)$ has at least the same number of modes as $p(x)$, but the learned joint distribution $p_\theta(x, a)$ can only have a single mode by Proposition 4.3, inducing a bias. To derive the universality in terms of KL, apply Theorem 4.1 to the joint distribution at hand. On the positive side, Souveton et al. (2024) find that having $p(a|x) = \mathcal{N}(\mu(x), \sigma(x)^2 I)$ brings the obtained $p_\theta(x) = \int p_\theta(x, a) da$ closer to the target. This can be seen having an independent augmentation $a \sim \mathcal{N}(0, I) \perp x$ plus a single RealNVP coupling shifting and scaling the augmented dimensions. This effectively breaks the volume-preservation of the flow in the joint space. It is unclear, however, whether this removes all biases from the volume-preserving flow.

B. Proofs on Volume-Preserving Normalizing Flows

B.1. Minimizer of Volume-Preserving Normalizing Flows

Volume-preserving flows are not universal in KL divergence, see Theorem 4.2. In this section, we consider what distribution a volume-preserving flow converges to instead. We first construct the latent distribution a sufficiently rich volume-preserving flows converges to when trained with KL divergence (Appendix B.1.1) and then show that this actually minimizes the KL (Appendix B.1.2).

We also demonstrate in what sense Δ_{affine} is sensitive to volume-preserving transformations. We therefore show in Appendix B.1.3 that this flow converges under our convergence measure Δ_{affine} if and only if that volume-preserving flow converges to a standard normal in the latent space under KL divergence.

B.1.1. ROTATIONALLY SYMMETRIC DISTRIBUTION WITH SAME LEVEL SET STRUCTURE

Let us repeat the change-of-variables formula for a volume-preserving flow from Equation (11):

$$p_\theta(x) = p(z = f_\theta(x)) \cdot C, \quad (44)$$

where $C = |f'_\theta(x)|$ is constant with respect to x . Intuitively, this means that such a flow can *permute* the probability mass at all locations x , and apply a single global factor to scale all probability values by spreading out the distribution.

We now construct the best possible distribution learned by a volume-preserving flow to an arbitrary input $p(x)$ in terms of KL divergence. We therefore split the input distribution $p(x)$ into its level sets:

$$L_v(p(\cdot)) := \{x \in \mathbb{R}^D : p(x) = v\}. \quad (45)$$

Acting on the input distribution, a volume-preserving flow yields a latent distribution $p_\theta(z)$ whose level set structure is closely related to that of the input distribution, in the following sense:

$$|L_v(p(x))| = \frac{|L_{v/C}(p_\theta(z))|}{C}. \quad (46)$$

Intuitively, this captures that a volume-preserving flow maps the input space to the latent space such that the level set in data space for the density value v is mapped to the level set in the latent space of level v/C – and the $(D - 1)$ -dimensional volumes are scaled by the factor C . Here, we have assumed that the level sets are $(D - 1)$ -dimensional.

In the following, we use these level sets to construct the distribution $p^*(z)$ a volume-preserving flow converges to in the latent space. This allows us to specify the best solution a volume-preserving flow can converge to. We first consider a fixed C and then later solve the optimization over C separately.

Lemma B.1. *Let $p(x)$ be a bounded continuous probability density with $(D - 1)$ -dimensional level sets almost everywhere. Then, a unique continuous probability density $p^*(z)$ with the following properties exists:*

1. *Its level sets have equal volume: $|L_v(p)| = |L_v(p^*)|$,*
2. *p^* is rotationally symmetric: $p^*(z) = p^*(Qz)$ for all $Q \in SO(D)$,*
3. *$p^*(z_1, 0, \dots, 0)$ is strictly monotonically decreasing in $0 \leq z_1 < \infty$.*

Proof. We write $p^*(x) = p^*(r)p^*(\Omega|r)$, where (r, Ω) are hyper-spherical coordinates. Since $p^*(x)$ should be rotationally symmetric, the distribution of the solid angle Ω is isotropic, and equal to one over the surface $A_{D-1}(r)$ of the $(D - 1)$ -dimensional hypersphere:

$$p^*(\Omega|r) = \frac{1}{A_{D-1}(r)} = \frac{\Gamma(D/2)}{2\pi^{\frac{D}{2}} r^{D-1}} \quad (47)$$

This makes p^* rotationally symmetric and leaves us with constructing $p^*(r)$.

Define the superlevel sets v for $p(x)$ as follows:

$$L_v^+(p) = \{x \in \mathbb{R}^D : p(x) \geq v\}. \quad (48)$$

Their volume $|L_v^+(p)|$ as measured in $(D - 1)$ dimensions is monotonically decreasing in v and its derivative yields the volume of the level set:

$$\frac{\partial}{\partial v} |L_v^+(p)| = |L_v(p)|. \quad (49)$$

We now demand that

$$|L_v^+(p)| = |L_v^+(p^*)| \quad (50)$$

and integrate out Ω :

$$|L_v^+(p^*)| = \int \mathbf{1}[p^*(x) \geq v] dx \quad (51)$$

$$= \int d\Omega \int r^{D-1} \mathbf{1}[p^*(r) \geq v] dr \quad (52)$$

$$= A_{D-1}(1) \int r^{D-1} \mathbf{1}[p^*(r) \geq v] dr. \quad (53)$$

Since $p^*(r)$ should decrease monotonically with r , we can replace the indicator function by integral boundaries, where the upper limit depends on the target density value v . We identify $\max_x p(x)$ with $R = 0$:

$$|L_v^+(p^*)| = A_{D-1}(1) \int_0^{R(v)} r^{D-1} dr \quad (54)$$

$$= A_{D-1}(1) \frac{1}{D} R(v)^D \quad (55)$$

Rearranging yields $R(v)$ from $|L_v^+(p)|$:

$$R(v) = \left(\frac{D|L_v^+(p)|}{A_{D-1}(1)} \right)^{\frac{1}{D}} \quad (56)$$

As $R(v)$ is monotonous in $|L_v^+(p)|$, $|L_v^+(p)|$ is continuous and monotonous in v , and $R(v)$ is invertible, its inverse can be used to define $p^*(r)$:

$$p^*(r) = R^{-1}(r). \quad (57)$$

By choosing $R(v)$ to fulfill Equation (50), their derivatives also match:

$$|L_v(p)| = \frac{\partial}{\partial v} |L_v^+(p)| = \frac{\partial}{\partial v} |L_v^+(p^*)| = |L_v^+(p^*)|. \quad (58)$$

The other properties of p^* follow directly from the construction above. The density is unique (up to zero sets) since a rotationally symmetric distribution is uniquely defined by a one-dimensional ray (Eaton, 1986). \square

We now show that the latent distribution $p^*(z)$ is actually attainable by a volume-preserving flow:

Lemma B.2. *Under the assumptions of Lemma B.1. There exists a function $f_\theta : \mathbb{R}^D \rightarrow \mathbb{R}^D$ that is bijective, continuous and volume-preserving with unit volume change ($|f'_\theta(x)| = 1$) almost everywhere and $p_\theta(z) = p^*(x = z)$.*

This means that there is a volume-preserving flow that exactly pushes $p(x)$ to its respective $p^*(x)$. Note that this volume-preserving flow has $|f'_\theta(x)| = 1$.

Proof. First, note that $p_\theta(z) = p^*(x = z)$ as constructed above can be achieved by a volume-preserving bijection f_θ . To see this, divide the space into the level sets $L_p(v)$ of $p(x)$. Since $p(x)$ is continuous, there is a countable sequence of thresholds v_i at which the number of connected components in the level set $L_p(v)$ jumps: The first jump is at $v_{\max} = \max_x p(x)$, below which find as many connected components as there are maximal modes in $p(x)$. The number of connected components changes whenever there is a saddle point or maximum in $p(x)$. Between each subsequent pair of jumps (v_i, v_{i+1}) , each connected component can be continuously assigned to a countable cluster number. This yields two tessellations of the entire space: One into level sets, and one into continuous connected components. To construct f_θ , we assign the highest points of $x^* \in \mathbb{R}^D$ to $f_\theta(x^*) = 0$. Then, we continuously arrange the finite number of components until the next jump in $(\max_x p(x), v_1)$ around the origin such that the resulting level sets are concentric circles. This constructs f_θ that pushes $p(x)$ to $p^*(x)$ and fulfills the above constraints. \square

B.1.2. BEST VOLUME-PRESERVING NORMALIZING FLOW UNDER KL DIVERGENCE

We are going to make use of the following identity that a volume-preserving flow with $|f'_\theta(x)| \neq 1$ with a standard normal in the latent space can be written as a volume-preserving flow with $|\tilde{f}'_\theta(x)| = 1$ and an alternative latent distribution $\tilde{p}(z) = \mathcal{N}(0, C^{-2}I)$:

Lemma B.3. *Given a volume-preserving bijection f_θ that is diffeomorphic almost everywhere, and with $|f'_\theta(x)| = C$ for some $C > 0$. Then, there exists a volume-preserving bijection \tilde{f}_θ with $|\tilde{f}'_\theta(x)| = 1$ such that:*

$$((f_\theta^{-1})_\# \mathcal{N}(0, I))(x) = ((\tilde{f}_\theta^{-1})_\# \mathcal{N}(0, C^{-2}I))(x) \quad (59)$$

In other words, the global volume change of a volume-preserving flow can be absorbed into a single scaling last layer at the latent end of the flow.

Proof. Let $\tilde{f}_\theta(x) = C^{-1}f_\theta(x)$, which has $|\tilde{f}'_\theta(x)| = C/C = 1$ and write $\mathcal{N}(0, C^{-2}I)$ as the push-forward through $z \mapsto C^{-1}z$. \square

Now we show our main Theorem 4.1 on volume-preserving flows:

Theorem 4.1. *Given a continuous bounded input density $p(x)$. Then, for any volume-preserving flow $p_\theta(x)$ with a standard normal latent distribution, the achievable KL divergence is bounded from below:*

$$\mathcal{D}_{\text{KL}}(p(x)||p_\theta(x)) \geq \mathcal{D}_{\text{KL}}(p^*(z)||\mathcal{N}(0, |\Sigma_{p^*(z)}|^{\frac{1}{D}} I)), \quad (60)$$

where $p^*(z)$ is constructed as in Lemma B.1, and $\Sigma_{p^*(z)}$ is its covariance matrix. The minimal loss is achieved for $C = |\Sigma_{p^*(z)}|^{-\frac{1}{2D}}$.

Proof. By Equation (22), $\mathcal{D}_{\text{KL}}(p(x)||p_\theta(x)) = \mathcal{D}_{\text{KL}}(p_\theta(z)||\mathcal{N}(0, I))$. The variant in the latent space can be rewritten using the entropy of $p_\theta(z)$ as:

$$\mathcal{D}_{\text{KL}}(p_\theta(z)||\mathcal{N}(0, I)) = \mathcal{D}_{\text{KL}}(p_\theta(z)||\mathcal{N}(0, I)) = -H[p_\theta(x)] - \int p_\theta(z) \log \mathcal{N}(z; 0, I) dz. \quad (61)$$

The entropy of the latent distribution of a volume-preserving flow only depends on the volume change constant C , but not on the exact choice of f_θ :

$$H[p_\theta(z)] = - \int p_\theta(z) \log p_\theta(z) dz \quad (62)$$

$$= - \int p(x) \log(p_\theta(z = f_\theta(x))C) dx - \log C \quad (63)$$

$$= - \int p(x) \log p(x) dx - \log C \quad (64)$$

$$= H[p(x)] - \log C. \quad (65)$$

Inserting into Equation (61):

$$\mathcal{D}_{\text{KL}}(p_\theta(z)||\mathcal{N}(0, I)) = \mathcal{D}_{\text{KL}}(p_\theta(z)||\mathcal{N}(0, I)) = -H[p(x)] + \log C - \int p_\theta(z) \log \mathcal{N}(z; 0, I) dz. \quad (66)$$

Using Lemma B.3, rewrite the last term in Equation (66) as an integral over x :

$$- \int p_\theta(z) \log \mathcal{N}(z; 0, C^{-2}I) dz = - \int p(x) \log \mathcal{N}(\tilde{f}_\theta(x); 0, C^{-2}I) dx. \quad (67)$$

This reveals the KL is minimized by assigning the highest values of $p(x)$ to the highest values of $\mathcal{N}(0, C^{-2}I)$. Since the order of densities $\mathcal{N}(z; 0, C^{-2}I)$ is the same regardless of C , the assignment $\tilde{f}_\theta(x)$ does not depend on C , which can be estimated separately via

$$\mathcal{D}_{\text{KL}}(p^*(z)||\mathcal{N}(0, C^{-2}I)). \quad (68)$$

Adapting (Draxler et al., 2022, Proposition 1), the KL divergence Equation (68) can be decomposed as follows:

$$\mathcal{D}_{\text{KL}}(p^*(z)||\mathcal{N}(0, C^{-2}I)) = \mathcal{D}_{\text{KL}}(p^*(z)||\mathcal{N}(0, |\Sigma_z|^{\frac{1}{D}} I)) + \mathcal{D}_{\text{KL}}(\mathcal{N}(0, |\Sigma_z|I)||\mathcal{N}(0, C^{-2}I)), \quad (69)$$

where $|\Sigma_z|$ is the determinant of the covariance matrix of the latent codes of p^* : $\Sigma_z = \text{Cov}_{z \sim p^*(z)}[z]$. The first term is invariant under scaling the latent codes, and the second term is minimal for $C^{-2} = |\Sigma_z|^{\frac{1}{D}}$. □

The remaining loss $\mathcal{D}_{\text{KL}}(p^*(z)||\mathcal{N}(0, |\Sigma_{p^*(z)}|^{\frac{1}{D}} I))$ can be reduced to zero by switching to spherical coordinates and learning $p^*(r)$ via a non-volume-preserving one-dimensional distribution, see Appendix B.2.

B.1.3. SENSITIVITY OF Δ_{affine} TO VOLUME-PRESERVING FLOWS

We now confirm that if a volume-preserving flow is not universal under KL divergence, it is also not universal under Δ_{affine} .

Lemma B.4. *For a family of normalizing flows with constant Jacobian determinant $|f'_\theta(x)| = \text{const}$, such that $\mathcal{D}_{\text{KL}}(p_\theta(z)||p^*(z)) \rightarrow 0$, it holds that $\Delta_{\text{affine}} \rightarrow 0$ if and only if $\mathcal{D}_{\text{KL}}(p^*(z)||\mathcal{N}(0, I)) = 0$.*

Proof. By Lemma 5.3, we can use $\Delta_{\text{affine}}^* > 0$ and $\Delta_{\text{affine}} > 0$ interchangeably. By its definition in Equation (31),

$$\Delta_{\text{affine}}^* = \max_Q \frac{1}{2} \sum_i^{D/2} \mathbb{E}_b [m_i^2(b) + \sigma_i^2(b) - 1 - \log \sigma_i^2(b)] \quad (70)$$

According to Theorem 4.1, the latent distribution $p_\theta(z) = p^*(z = x)$ minimizes the KL in the latent space: $\mathcal{D}_{\text{KL}}(p_\theta(z)||\mathcal{N}(0, I))$ (the exact assignment between data and latent codes is not unique, but all lead to the same latent estimate).

At this minimum, $p_\theta(z) = p^*(x = z)$. Since $p^*(x)$ is symmetric under rotations, it holds that, it holds that $m_i(b) = 0$ for $(a, b) = Qx$ in all rotations Q . However, since $\text{Var}_{a \sim (Q \# p^*)(a)[a_i|b]} \neq 1$ for all Q , it holds that

$$\Delta_{\text{affine}}^* = \frac{1}{2} \sum_i^{D/2} \mathbb{E}_b [\sigma_i^2(b) - 1 - \log \sigma_i^2(b)], \quad (71)$$

which evaluates to the same value regardless of Q since the distribution is rotationally symmetric.

While this minimum may not be exactly achieved by a continuous volume-preserving flow, a sufficiently rich architecture is able to achieve universality $\mathcal{D}_{\text{KL}}(p_\theta(z)||\mathcal{N}(0, I)) \rightarrow \mathcal{D}_{\text{KL}}(p^*(z)||\mathcal{N}(0, I))$. Since the KL divergence implies the convergence of expectation values (Gibbs & Su, 2002), it holds that $\Delta_{\text{affine}}^* \rightarrow \frac{1}{2} \sum_i^{D/2} \mathbb{E}_b [\sigma_i^2(b) - 1 - \log \sigma_i^2(b)] \geq 0$. Equality holds if and only if $\mathcal{D}_{\text{KL}}(p^*(z)||\mathcal{N}(0, I)) = 0$. \square

Note that the same argument also applies to Wasserstein distance and weak convergence, so this does not indicate that Δ_{affine} is more informative about convergence than these convergence metrics.

B.2. Fixing Volume-Preserving Flows by Learning the Latent Radial Distribution

By Theorem 4.1, the global minimizer of a volume-preserving flow is given by $p^*(z)$ in the latent space, as characterized by Lemma B.1. Both $p^*(z)$ and the standard normal distribution $p(z) = \mathcal{N}(0, C^2 I)$ are rotationally symmetric, so it is useful to make a change of variables to hyperspherical coordinates (r, Ω) :

$$p^*(z) = p^*(r(z))p^*(\Omega(z)|r(z)) \left| \frac{d(r, \Omega)}{dz} \right|, \quad (72)$$

$$p(z) = p(r(z))p(\Omega(z)|r(z)) \left| \frac{d(r, \Omega)}{dz} \right|. \quad (73)$$

The rotational symmetry implies that:

$$p^*(\Omega|r) = p(\Omega|r) = \frac{1}{A_{D-1}(r)}, \quad (74)$$

where $A_{D-1}(r)$ is the surface of the $(D - 1)$ -dimensional sphere of radius r .

This means that we only need to match $p^*(r)$ and $p(r)$. This can be achieved by a one-dimensional transformation of r .

Fixing the latent distribution can be done after training, as already training the volume-preserving flow with a standard normal in the latent space sorts the data such that $p(x_a) \geq p(x_b)$ implies that $r_a \leq r_b$, that is points of higher ground truth density are mapped to points closer to the origin. If we now replace the latent distribution with another distribution that fulfills the same constraint by fitting $p^*(r)$, the minimizer with this latent distribution remains the same. Thus, we can train with a standard normal and later fix the one-dimensional distribution.

B.3. Proof of Theorem 4.2

Theorem 4.2. *The family of normalizing flows with constant Jacobian determinant $|f'_\theta(x)| = \text{const}$ is not a universal distribution approximator under KL divergence.*

Proof. In this section, we want to present a two-dimensional example, for which no normalizing flow with constant Jacobian determinant can be constructed such that the KL-divergence between the data distribution and the distribution defined by the normalizing flow is zero.

$$p(x, y) = \begin{cases} 0.9 & \text{if } (x, y) \in [-0.5, 0.5] \times [-0.5, 0.5] \\ 0.9 - k \cdot (|x| - 0.5) & \text{if } |x| \in [0.5, \frac{0.9}{k} + 0.5] \wedge |y| \in [0, |x|] \\ 0.9 - k \cdot (|y| - 0.5) & \text{if } |y| \in [0.5, \frac{0.9}{k} + 0.5] \wedge |x| \in [0, |y|] \\ 0 & \text{else} \end{cases} \quad (75)$$

The data distribution $p(x, y)$ which has to be approximated by the model is defined in Equation (75). This data distribution has a constant value of 0.9 in a box centered around the origin with a side length of one. This region of constant density is skirted by a margin where the density decreases linearly to zero. Outside the decreasing region, the density is zero. The linear decline is governed by the constant k in Equation (75) which has to be chosen such that the density integrates to one. Since our example only requires the region of constant density but not the decaying tails of it, the exact functional form of the decaying regions are not relevant as long as they lead to a properly normalized distribution. Equation (75) only provides a possible definition of such a density.

To approximate this data distribution, a normalizing flow as defined in Section 3 is considered. In this example, we focus on normalizing flows with constant Jacobian determinant. To simplify notation, we define $J = |f'_\theta(x)| = \text{const}$.

$$A = \{(x, y) \in \mathbb{R}^2 : 0.9 - \epsilon < p_\theta(x, y)\} \quad (76)$$

$$B = [-0.5, 0.5] \times [-0.5, 0.5] \quad (77)$$

$$\bar{A} = B \setminus A \quad (78)$$

We choose $\epsilon = 0.1$ and use this constant to define the set A (see Equation (76)). In addition we define B which is the region of the data space, where the data distribution has a constant value of 0.9 (see Equation (77)). \bar{A} is the complement of A on B (see Equation (78)).

The aim of this example, is to find lower bounds for the KL-divergence between the data distribution and the distribution defined by the normalizing flow. To find these bounds we use Pinsker's inequality (Gibbs & Su, 2002) which links the total variation distance to the Kullback-Leibler divergence:

$$\delta_{\text{TV}}(p, p_\theta) = \sup_{A \text{ measurable}} |P(A) - P_\theta(A)| \leq \sqrt{\frac{1}{2} D_{\text{KL}}(p || p_\theta)} \quad (79)$$

It is worth mentioning that constructing one measurable event for which $|P(A) - P_\theta(A)| > 0$ provides a lower bound for the total variation distance and therefore for the KL divergence.

To construct such an event, we consider two distinct cases, which consider different choice for the normalizing flow, characterized by the value of the absolute Jacobian determinant.

Case 1: $A = \emptyset$

This case arises if the absolute Jacobian determinant is so small, that the distribution defined by the normalizing flow never exceeds the limit defining A or if it is chosen so large, that the volume of A vanishes.

In this case, we find $\bar{A} = B$ and $|\bar{A}| = 1$ where $|\bar{A}|$ denotes the volume of the data space occupied by \bar{A} . Using the fact that the data distribution has a constant value of 0.9 in B and that $p_\theta < 0.9 - \epsilon$ in $\bar{A} = B$,

$$|P(\bar{A}) - P_\theta(\bar{A})| = |0.9 - P_\theta(\bar{A})| \quad (80)$$

$$\geq |0.9 \cdot 1 - (0.9 - \epsilon) \cdot 1| \quad (81)$$

$$= |\epsilon| = \epsilon \quad (82)$$

Using Equation (82) as a lower bound for the total variation distance Equation (82) we can apply Equation (79) to find Equation (85) as a lower bound for the KL divergence.

$$D_{KL}(p||p_\theta) \geq 2 \cdot \delta_{TV}(p, p_\theta)^2 \quad (83)$$

$$\geq 2 \cdot \epsilon^2 \quad (84)$$

$$= 0.02 \quad (85)$$

Case 2: $A \neq \emptyset$

Inserting the definition of $p_\theta(x, y)$ as given in Equation (1) into the definition of A (see Equation (76)) and rewriting the condition defining the set yields Equation (86).

$$A = \left\{ (x, y) \in \mathbb{R}^2 : \frac{0.9 - \epsilon}{J} < p(z = f_\theta(x, y)) \right\} \quad (86)$$

This defines a set C in the latent space which is defined in Equation (87).

$$C = \left\{ z \in \mathbb{R}^2 : \frac{0.9 - \epsilon}{J} < p(z) \right\} \quad (87)$$

Since the normalizing flows considered in this example have a constant Jacobian, the volume of A in the data space is directly linked to the volume of C in the latent space via Equation (88).

$$|A| = \frac{1}{J} \cdot |C| \quad (88)$$

The definition of C Equation (87) shows, that C is a circle around the origin of the latent space. To determine the volume of C we compute the radius of this circle. This is done by inserting the definition of the latent distribution, which is a two-dimensional standard normal distribution, into the condition defining C (see Equation (87)). This yields Equation (89). Since the latent distribution is rotational invariant, one can simply look at it as a function of the distance r from the origin. Solving the for r leads to Equation (90).

$$\frac{0.9 - \epsilon}{J} < \frac{1}{2\pi} \cdot \exp\left(-\frac{r^2}{2}\right) \quad (89)$$

$$\Rightarrow r = \sqrt{-2 \cdot \log\left(\frac{2\pi \cdot (0.9 - \epsilon)}{J}\right)} \quad (90)$$

Inserting Equation (90) into the formula for the area of a circle and using Equation (88), yields Equation (92) as an expression for the volume of A . The lower bound for the volume of A arises from finding the local maximum (which is also the global maximum) of Equation (92) with respect to the absolute Jacobian determinant J .

$$|A| = \frac{1}{J} \cdot \pi \cdot r^2 \quad (91)$$

$$= \frac{1}{J} \cdot 2\pi \cdot \log\left(\frac{J}{2\pi \cdot (0.9 - \epsilon)}\right) \quad (92)$$

$$\leq \frac{1}{e \cdot (0.9 - \epsilon)} \quad (93)$$

$$(94)$$

As in the previous case, we now compute $|P(\bar{A}) - P_\theta(\bar{A})|$.

$$|P(\bar{A}) - P_\theta(\bar{A})| = ||\bar{A}| \cdot 0.9 - P_\theta(\bar{A})| \quad (95)$$

$$\geq ||\bar{A}| \cdot 0.9 - |\bar{A}| \cdot (0.9 - \epsilon)| \quad (96)$$

$$= |\bar{A}| \cdot \epsilon \quad (97)$$

$$\geq (1 - |A|) \cdot \epsilon \quad (98)$$

$$\geq \left(1 - \frac{1}{e \cdot (0.9 - \epsilon)}\right) \cdot \epsilon \quad (99)$$

Using Equation (83) and Equation (99) as a lower bound for the total variation distance and inserting our choice for ϵ yields Equation (100) as a lower bound for the KL divergence between the data distribution and the distribution defined by the normalizing flow.

$$D_{KL}(p||p_\theta) \geq 2 \cdot \left(\epsilon \cdot \left(1 - \frac{1}{e \cdot (0.9 - \epsilon)}\right)\right)^2 \quad (100)$$

$$\approx 0.0058 \quad (101)$$

We can conclude that we have derived lower bounds for the KL divergence between the data distribution and the distribution defined by the normalizing flow, which cannot be undercut by any normalizing flow with a constant absolute Jacobian determinant. Therefore, we have proven that the class of normalizing flows with constant (absolute) Jacobian determinant cannot approximate arbitrary continuous distributions if one uses the KL divergence as a convergence measure. \square

B.4. Proof of Proposition 4.3

Definition B.5. Given a probability density $p(x)$ and a connected set $M \subset \mathbb{R}^D$. Then, M is called a *mode* of $p(x)$ if

$$p(x) = p(y) \quad \forall x, y \in M, \quad (102)$$

and there is a neighborhood U of M such that:

$$p(x) > p(y) \quad \forall x \in M, y \in U \setminus M. \quad (103)$$

With this definition of a mode, let us characterize the correspondence between modes of $p_\theta(x)$ and $p(z)$ for a volume-preserving flow:

Lemma B.6. Given a latent probability density $p(z)$, a diffeomorphism $f_\theta : \mathbb{R}^D \rightarrow \mathbb{R}^D$ with constant Jacobian determinant $|f'_\theta(x)| = \text{const}$ and a mode $M \subset \mathbb{R}^D$. Then, $f(M)$ is a mode of $p_\theta(x)$.

Proof. We show that $f(M)$ fulfills Definition B.5. First, for every $x, y \in f(M)$: The pre-images of x, y are unique in M as f is bijective, that is: $f^{-1}(x), f^{-1}(y) \in M$. As M is a mode:

$$p(f^{-1}(x)) = p(f^{-1}(y)). \quad (104)$$

We follow:

$$(f_\#p)(x) = p(f^{-1}(x))|J| = p(f^{-1}(y))|J| = (f_\#p)(y), \quad (105)$$

where we have used the change-of-variables formula for bijections and that $|J| = \text{const}$.

Let U be a neighborhood of M such that Equation (103) is fulfilled. As f is continuous, there is a neighborhood V of $f(M)$ such that $V \subseteq f(U)$. Consider $x \in f(M), y \in V \setminus f(M)$. As M is a mode:

$$p(f^{-1}(x)) > p(f^{-1}(y)). \quad (106)$$

Multiplying both sides by $|J|$, we find:

$$(f_\#p)(x) = p(f^{-1}(x))|J| > p(f^{-1}(y))|J| = (f_\#p)(y). \quad (107)$$

Thus, $f(M)$ is a mode of $(f_\#p)(x)$ by Definition B.5. \square

This makes us ready for the proof:

Proposition 4.3. *A normalizing flow $p_\theta(x)$ with constant Jacobian determinant $|f'_\theta(x)| = \text{const}$ has the same number of modes as the latent distribution $p(z)$.*

Proof. By Lemma B.6, every mode of $p(x)$ implies a mode of $(f_{\#}p)(x)$. Also, every mode of $(f_{\#}p)(x)$ implies a mode of $(f_{\#}^{-1}f_{\#}p)(x) = p(x)$. Therefore, there is a one-to-one correspondence of modes between $p(x)$ and $(f_{\#}p)(x)$. \square

C. Proofs on Affine Coupling Flows

C.1. Proof of Unique Fixed Point of Affine Coupling Blocks

C.1.1. UNDERLYING RESULTS FROM PREVIOUS WORK

Here, we restate the results from the literature that our main proof is based on:

First, (Eaton, 1986) show that if for some vector-valued random variable X and every pair of orthogonal projections the mean of one projection conditioned on the other is zero, then X follows a spherical distribution:

Theorem C.1 (Eaton (1986)). *Suppose the random vector $X \in \mathbb{R}^D$ has a finite mean vector. Assume that for each vector $v \neq 0$ and for each vector u perpendicular to v (i.e. $u \cdot v = 0$):*

$$\mathbb{E}[u \cdot X | v \cdot X] = 0. \quad (108)$$

Then X is spherical and conversely.

Secondly, Cambanis et al. (1981, Corollary 8a) identifies the Gaussian from all elliptically contoured (which includes spherical) distributions. We write it in the form of Bryc (1995, Theorem 4.1.4):

Theorem C.2 (Bryc (1995)). *Let $p(x)$ be radially symmetric with $\mathbb{E}[\|x\|^\alpha] < \infty$ for some $\alpha > 0$. If*

$$\mathbb{E}[\|x_{1,\dots,m}\|^\alpha | x_{m+1,\dots,n}] = \text{const}, \quad (109)$$

for some $1 \leq m < n$, then $p(x)$ is Gaussian.

Finally, Draxler et al. (2020, Theorem 1) show that the explicit form of the maximally achievable loss improvement by an affine coupling block $\Delta_{\text{affine}}^*(p_\theta(z))$ if the data is rotated by a fixed rotation layer Q is given by (omitting the dependence on $p_\theta(z)$ to avoid clutter):

$$\Delta_{\text{affine}}^*(Q) = \mathcal{D}_{\text{KL}}(p_0(z) \| p(z)) - \min_{s,t} \mathcal{D}_{\text{KL}}(p_{s,t|Q}(z) \| p(z)) \quad (110)$$

$$= \frac{1}{2} \mathbb{E}_b [m_i(b)^2 + \sigma_i(b)^2 - 1 - \log \sigma_i(b)^2]. \quad (111)$$

Here, s, t are the scaling and translation in an affine coupling block (see Equation (8)), and we optimize over continuous functions for now. By $p_{s,t|Q}(z)$ we denote the latent distribution achieved if $\tilde{a}(a; b) = s(b) \odot a + t(b)$ is applied to $p_0(a, b)$, the rotated version of the incoming $p_0(z)$. The symbols $m_i(b), \sigma_i(b)^2$ are conditional moments of the active dimensions a_i conditioned on the passive dimensions b :

$$m_i(b) = \mathbb{E}_{a_i|b}[a_i], \quad \sigma_i(b) = \mathbb{E}_{a_i|b}[a_i^2] - m_i(b)^2. \quad (112)$$

These conditional moments are continuous functions of b if $p(x)$ is a continuous distribution and $p(b) > 0$ for all passive $b \in \mathbb{R}^{D/2}$. The improvement in Equation (111) is achieved by the affine coupling block with the following subnetwork:

$$s_i^*(b) = \frac{1}{\sigma_i(b)}, \quad t_i^*(b) = -\frac{m_i(b)}{\sigma_i(b)}. \quad (113)$$

Note that $s^*(b)$ and $t^*(b)$ are continuous functions and not actual neural networks. In the next section, we show that a similar statement on practically realizable neural networks that is sufficient for our universality.

C.1.2. RELATION TO PRACTICAL NEURAL NETWORKS (LEMMA 5.3)

Before moving to the proof of Theorem 5.2, we show the helper statement Lemma 5.3. For reference, let us repeat the definition of the loss improvement by an affine coupling block $\Delta_{\text{affine}}(p_\theta(z))$:

Definition 5.1. *Given a normalizing flow $p_\theta(x)$ on a continuous probability distribution $p(x)$ with finite first and second moment and $p(x) > 0$ everywhere. Then, we define the loss improvement by an affine coupling block as:*

$$\Delta_{\text{affine}}(p_\theta(z)) := \mathcal{D}_{\text{KL}}(p_\theta(z) \| p(z)) - \min_{\theta_+} \mathcal{D}_{\text{KL}}(p_{\theta \cup \theta_+}(z) \| p(z)), \quad (114)$$

where $\theta_+ = (Q, \varphi)$ parameterizes a single L -bi-Lipschitz affine coupling block whose conditioner neural network ψ_φ has at least two hidden layers of finite width and ReLU activations and $p_\theta(z) = (f_{\theta\#}p)(z)$.

To relate Equations (111) and (113) to actually realizable networks, which cannot *exactly* follow the arbitrary continuous functions $s_i^*(b), t_i^*(b)$, the following statement asserts that the fix point of adding coupling layers with infinitely expressive conditioner functions is the same as actually realizable and well-conditioned coupling blocks:

Lemma 5.3. *Given a continuous probability density $p(z)$ on $z \in \mathbb{R}^k$. Then,*

$$\Delta_{\text{affine}}^*(p_\theta(z)) > 0 \quad (115)$$

if and only if:

$$\Delta_{\text{affine}}(p_\theta(z)) > 0. \quad (116)$$

Proof. First, note that $\Delta_{\text{affine}}^*(Q) \geq \Delta_{\text{affine}}(Q) \geq 0$ since no practically realizable coupling block can achieve better than Equation (111). Thus, if $\Delta_{\text{affine}}^*(Q) = 0$, so is $\Delta_{\text{affine}}(Q) = 0$.

For the reverse direction, we fix $Q = I$, and otherwise consider a rotated version of p . Also, without loss of generalization, we consider one single active dimension a_i in the following, but the construction can then be repeated for each other active dimension.

If we apply any affine coupling layer $f_{\text{cpl},\varphi}(a; b) = s_\varphi(b)a + t_\varphi(b)$, the loss change by this layer can be computed from the theoretical maximal improvement $\Delta_{\text{affine}}^*(Q)$ before and after adding this layer $\tilde{\Delta}_{\text{affine}}^*(I)$:

$$\Delta_{\text{affine}}(I) = \Delta_{\text{affine}}^*(I) - \tilde{\Delta}_{\text{affine}}^*(I) = \frac{1}{2} \mathbb{E}_b [m_i(b)^2 + \sigma_i(b)^2 - 1 - \log \sigma_i(b)^2] - \frac{1}{2} \mathbb{E}_b [\tilde{m}_i(b)^2 + \tilde{\sigma}_i(b)^2 - 1 - \log \tilde{\sigma}_i(b)^2]. \quad (117)$$

The moments after the affine coupling layers read:

$$\tilde{m}_i(b) = s_\varphi(b)m_i(b) + t_\varphi(b), \quad \tilde{\sigma}_i(b) = s_\varphi(b)\sigma_i(b). \quad (118)$$

Case 1: $\mathbb{E}_b [\sigma_i(b)^2 - 1 - \log \sigma_i(b)^2] > 0$:

Then, without loss of generality, by continuity and positivity of p and consequential continuity of $\sigma_i(b)$ in b , there is a convex open set $A \subset \mathbb{R}^{D/2}$ with non-zero measure $p(A) > 0$ where $\sigma_i(b) > 1$. If $\sigma_i(b) < 1$ everywhere, apply the following argument flipped around $\sigma_i(b) = 1$.

Denote by $\sigma_{\max} = \max_{b \in A} \sigma_i(b)$. Then, by continuity of $\sigma_i(b)$ there exists $B \subset A$ so that $\sigma_i(b) > (\sigma_{\max} - 1)/2 + 1 =: \sigma_{\max/2}$ for all $b \in B$. Let $C \subset B$ be a multidimensional interval $[l_1, r_1] \times \dots \times [l_{D/2}, r_{D/2}]$ with $p(C) > 0$ inside B .

Now, we construct a ReLU neural network with two hidden layers with the following property, where $F \subset E \subset C$ are specified later with $p(F) > p(E) > 0$:

$$\begin{cases} f_\varphi(x) = \frac{1}{\sigma_{\max/2}} & x \in E \subset D \\ \frac{1}{\sigma_{\max/2}} \leq f_\varphi(x) < 1 & x \in D \\ f_\varphi(x) = 0 & \text{else.} \end{cases} \quad (119)$$

To do so, we make four neurons for each dimension $i = 1, \dots, D/2$:

$$\text{ReLU}(x_i - l_i), \text{ReLU}(x_i - l_i - \delta), \text{ReLU}(x_i - r_i), \text{ReLU}(x_i - r_i + \delta), \quad (120)$$

where $0 < \delta < \min_i(r_i - l_i)/4$. If we add these four neurons with weights $1, -1, -1, 1$, we find the following piecewise function:

$$\begin{cases} 0 & x \leq l_i \\ x - l_i & l_i < x < l_i + \delta \\ \delta & l_i + \delta \leq x \leq r_i - \delta \\ r_i - x & r_i - \delta < x < r_i \\ 0 & r_i \leq x. \end{cases} \quad (121)$$

If we repeat this for each dimension and add together all neurons with the corresponding weights into a single neuron in the second layer, then only inside $D = (l_1 + \delta, r_1 - \delta) \times \cdots \times (l_{D/2} + \delta, r_{D/2} - \delta) \subset C$ the weighted sum would equal $\delta D/2$. By choosing δ as above, this region has nonzero volume. We thus equip the single neuron in the second layer with a bias of $-\delta D/2 + \epsilon$ for some $\epsilon < \delta$, so that it is constant with value ϵ inside $E = (l_1 + \delta - \epsilon, r_1 - \delta + \epsilon) \times \cdots \times (l_{D/2} + \delta + \epsilon, r_{D/2} - \delta - \epsilon) \subset D$ and smoothly interpolates to zero in the rest of D .

For the output neuron of our network, we choose weight $(\sigma_{\max/2} - 1)/\epsilon$ and bias 1. By inserting the above construction, we find the network specified in Equation (119).

Now, for all $b \in D$,

$$1 < \tilde{\sigma}_i(b) < \sigma_i(b), \quad (122)$$

so that

$$\tilde{m}_i(b)^2 + \tilde{\sigma}_i(b)^2 - 1 - \log \tilde{\sigma}_i(b)^2 < m_i(b)^2 + \sigma_i(b)^2 - 1 - \log \sigma_i(b)^2. \quad (123)$$

Thus, parameters φ exist that improve on the loss. (Note that this construction can be made more effective in practice by identifying the sets where $\sigma > 1$ resp. $\sigma < 1$ and then building neural networks that output one or scale towards $\tilde{\sigma}(b) = 1$ everywhere. Because we are only interested in identifying improvement, the above construction is sufficient.)

Now, regrading t_φ , we focus on $\mathbb{E}_b[m_i(b)^2] > 0$ (otherwise choose $t_\varphi = 0$ as a constant, which corresponds to a ReLU network with all weights and biases set to zero):

$$\mathbb{E}_b[m_i(b)^2] > \mathbb{E}_b[(s_\varphi(b)m_i(b) + t_\varphi(b))^2]. \quad (124)$$

By (Hornik, 1991, Theorem 1) there always is a t_φ that fulfills this relation.

Case 2: $\mathbb{E}_b[\sigma_i(b)^2 - 1 - \log \sigma_i(b)^2] = 0$. Then, choose the neural network $s_\varphi(b) = 1$ as a constant. As $\Delta_{\text{affine}} > 0$, $\mathbb{E}_{b \sim p(a,b)}[m_i(b)^2] > 0$ and we can use the same argument for the existence of t_φ as before.

It is left to show that a L -bi-Lipschitz coupling block can be constructed. To achieve this, replace the action of the coupling block $\tilde{a}_i = s(b)a_i + t(b)$ by $\tilde{a}_i = \alpha(s(b)a_i + t(b)) + (1 - \alpha)a_i$. Since $s(b)$ and $t(b)$ above were constructed to move the data in the right direction, we obtain a finite loss improvement, since $\alpha > 0 \Leftrightarrow \Delta_{\text{affine}} > 0$. The Jacobian J of the restricted coupling block is $|f'_{\text{cpl}}| = \alpha|f'_{\text{original}}| + (1 - \alpha)I$. Since the eigenvectors are unchanged, all eigenvalues $\lambda_{\text{original}}^{(i)}$ of $|f'_{\text{original}}|$ are modified to $\lambda_\alpha^{(i)} = \alpha\lambda_{\text{original}}^{(i)} + (1 - \alpha)$. This moves all eigenvalues closer to 1. Choose $\alpha > 0$ such that $\min_i \lambda_\alpha^{(i)} \geq L^{-1}$ and $\max_i \lambda_\alpha^{(i)} \leq L$ to achieve L -bi-Lipschitzness. \square

C.1.3. PROOF OF THEOREM 5.2

The following theorem based on Lemma 5.3 shows that $\Delta_{\text{affine}}(p_\theta(z))$ is a useful measure of convergence to the standard normal distribution:

Theorem 5.2. *With the definitions from Definition 5.1:*

$$p_\theta(z) = \mathcal{N}(z; 0, I) \iff \Delta_{\text{affine}}(p_\theta(z)) = 0. \quad (125)$$

Proof. The forward direction is trivial: $p(z) = \mathcal{N}(0, I)$ and therefore $\mathcal{D}_{\text{KL}}(p(z) \|\mathcal{N}(0, I)) = 0$. As adding a identity layer is a viable solution to Equation (23), there is a φ with $\mathcal{D}_{\text{KL}}(p_\varphi(z) \|\mathcal{N}(0, I)) = 0$, and thus $\Delta_{\text{affine}}(p_\theta(z)) = 0$.

For the reverse direction, start with $\Delta_{\text{affine}}(p_\theta(z)) = 0$. Then, by Lemma 5.3, also $\Delta_{\text{affine}}^*(p_\theta(z)) = 0$.

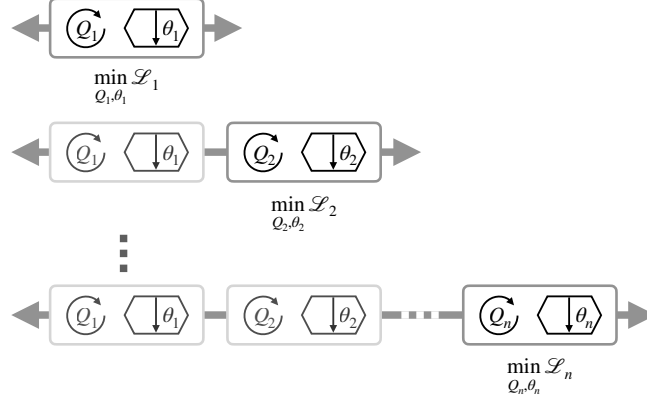


Figure 4. The normalizing flow we construct in our proof is remarkably simple: We iteratively add coupling blocks, optimizing the parameters of the new block while keeping previous parameters fixed. Theorem 5.2 shows that if adding another blocks shows no improvement in the loss, the flow has converged to a standard normal distribution in the latent space. Since the total loss that can be removed is finite, the flow converges.

The maximally achievable loss improvement for any rotation Q is then given by Equation (31):

$$\Delta_{\text{affine}}^*(p_\theta(z)) = \max_Q \frac{1}{2} \sum_{i=1}^{D/2} \mathbb{E}_b [m_i(b)^2 + \sigma_i(b)^2 - 1 - \log \sigma_i(b)^2] = 0. \quad (126)$$

It holds that both $x^2 \geq 0$ and $x^2 - 1 - \log x^2 \geq 0$. Thus, the following two summands are zero:

$$0 = \frac{1}{2} \mathbb{E}_b [m_i(b)^2], \quad (127)$$

$$0 = \frac{1}{2} \mathbb{E}_b [\sigma_i(b)^2 - 1 - \log \sigma_i(b)^2]. \quad (128)$$

This holds for all Q since the maximum over Q is zero.

By continuity of $p(b)$ and $m_1(b)$ in p , this implies for all b :

$$\mathbb{E}_{a_1|b} [a_1] = 0. \quad (129)$$

Fix b_1 and marginalize out the remaining dimensions $b_2, \dots, D/2$ to compute the mean of a_1 conditioned on b_1 :

$$m_{a_1|b} = \mathbb{E}_{a_1|b_1} [a_1] = \mathbb{E}_{b_1, \dots, D/2} [\mathbb{E}_{a_1|b} [a_1]] = \mathbb{E}_{b_1, \dots, D/2} [0] = 0. \quad (130)$$

As a_1 and b_1 are arbitrary orthogonal directions since the above is valid for any Q , we can employ Theorem C.1 to follow that $p(x)$ is spherically symmetric.

We are left with showing that for a spherical $p(x)$, if for all Q there is no improvement $\Delta_{\text{affine}}(Q)$, then $p(x) = \mathcal{N}(0, I)$.

Without loss of generality, we can fix $Q = I$, as $(Q_\# p)(x) = p(x)$ for all Q . We write $x = (p; a)$.

As $\Delta_{\text{affine}} = 0$, we can follow $\sigma_i(b) = 1$ like above. This implies that:

$$\mathbb{E}_{a|b} [\|a\|^2] = \sum_{i=1}^{D/2} (m_i(b)^2 + \sigma_i(b)^2) = D/2. \quad (131)$$

In particular, this is independent of b and we can thus apply Theorem C.2 with $\alpha = 2$.

Finally, $m(b) = 0$ and $\sigma_i(b) = 1$ for all Q imply that $p(x) = \mathcal{N}(0, I)$. \square

C.2. Proof of Theorem 5.4

Theorem 5.4. For every continuous $p(x)$ with finite first and second moment with infinite support, there is a sequence of normalizing flows $f_n(x)$ consisting of n affine coupling blocks such that:

$$p_n(z) \xrightarrow{n \rightarrow \infty} \mathcal{N}(z; 0, I), \quad (132)$$

in the sense that $\Delta_{\text{affine}}(p_n(z)) \xrightarrow{n \rightarrow \infty} 0$.

Proof. The proof idea of iteratively adding new layers which are trained without changing previous layers is visualized in Figure 4.

Let us consider a coupling-based normalizing flow of depth n and call the corresponding latent distributions $p_n(z)$, where $n = 0$ corresponds to the initial data distribution $p(x)$. Denote by $\mathcal{L}_n = \mathcal{D}_{\text{KL}}(p_n(z) \| p(z))$ the corresponding loss. Then, if we add another layer to the flow, we achieve a difference in loss of: $\Delta_{\text{affine}}(p_n(z)) = \mathcal{L}_n - \mathcal{L}_{n+1}$.

Without loss of generality, we may assume that the rotation layer Q of each block can be chosen freely. Otherwise, add 48 coupling blocks with fixed rotations that together exactly represent the Q we want, as shown by Koehler et al. (2021, Theorem 2).

We construct the blocks of the flow iteratively: Choose the rotation and subnetwork parameters φ_{n+1} of each additional block such that the block maximally reduces the loss, keeping the parameters of the previous blocks $\varphi_{1,\dots,n}$ fixed. Then, $\Delta_{\text{affine}}(p_n(z))$ attains the value given in Equation (24) (Equation (114)):

$$\Delta_{\text{affine}}(p_n(z)) = \mathcal{L}_n - \mathcal{L}_{n+1} = \mathcal{D}_{\text{KL}}(p_{\varphi_{1,\dots,n}}(z) \| p(z)) - \min_{\varphi_{n+1}} \mathcal{D}_{\text{KL}}(p_{\varphi_{1,\dots,n} \cup \varphi_{n+1}}(z) \| p(z)) \geq 0, \quad (133)$$

Each layer contributes a non-negative improvement in the loss which can at most sum up to the initial loss:

$$\sum_{k=0}^{n-1} \Delta_{\text{affine}}(p_k(z)) = \mathcal{L}_0 - \mathcal{L}_n \leq \mathcal{L}_0 \quad \text{for all } n \geq 1, \quad (134)$$

and the inequality is due to $\mathcal{L}_n \geq 0$. For a non-negative series that is bounded above, the terms of the series must converge to zero (Rudin, 1976, Theorems 3.14 and 3.23), which shows convergence in terms of Section 5.2:

$$\sum_{n=0}^{\infty} \Delta_{\text{affine}}(p_n(z)) \leq \mathcal{L}_0 < \infty \Rightarrow \Delta_{\text{affine}}(p_n(z)) \rightarrow 0. \quad (135)$$

□

C.3. Relation to Convergence in KL

Corollary C.3. *Given a series of probability distributions $p_n(z)$. Then, convergence in KL divergence*

$$\mathcal{D}_{\text{KL}}(p_n(z) \| \mathcal{N}(0, 1)) \xrightarrow{n \rightarrow \infty} 0 \quad (136)$$

implies convergence in the loss improvement by a single affine coupling as in Definition 5.1:

$$\Delta_{\text{affine}}(p_n(z)) \xrightarrow{n \rightarrow \infty} 0. \quad (137)$$

Proof. By assumption, for every $\epsilon > 0$ there exists $N \in \mathbb{N}$ such that:

$$\mathcal{D}_{\text{KL}}(p_n(z) \| \mathcal{N}(0, 1)) < \epsilon \quad \forall n > N. \quad (138)$$

This implies convergence of $\Delta_{\text{affine}}(p_n(z))$, by the following upper bound via the sum of all possible future improvements which is bounded from above by the total loss:

$$\Delta_{\text{affine}}(p_n(z)) \leq \sum_{m=n}^{\infty} \Delta_{\text{affine}}(p_m(z)) \leq \mathcal{D}_{\text{KL}}(p_n(z) \| \mathcal{N}(0, 1)) < \epsilon \quad \forall n > N. \quad (139)$$

□

C.4. Convergence in Wasserstein but not in KL Divergence

In Section 5.1 we argued that convergence under Wasserstein distance $W_2(p, p_n) \xrightarrow{n \rightarrow \infty} 0$ does not imply convergence under KL divergence $\mathcal{D}_{\text{KL}}(p||p_n) \xrightarrow{n \rightarrow \infty} 0$. We illustrate that via an example:

Take a standard normal target $p(x) = \mathcal{N}(0, 1)$ in 1D and approximate by a mixture of δ -distributions:

$$p_n(x) = \sum_{i=-\infty}^{\infty} p([i a_n, (i+1)a_n])\delta(x - a_n i). \quad (140)$$

This mixture splits the input space into bins of width a_n , and positions a δ -peak at the left of each bin, weighted by the amount of mass in the bin in the target distribution.

The optimal transport plan underlying the Wasserstein distance redistributes the weight from the left edge of each bin over the entire bin. This means that the total distance any point has to travel under the optimal transport plan is a . It thus holds that $W_2(p, p_n) \leq a_n$. If we choose $a_n \rightarrow 0$, so does $W_2(p, p_n) \rightarrow 0$.

The KL divergence can be lower bounded by the total variation via Pinsker's inequality (see also Equation (79)):

$$\delta_{\text{TV}}(p, p_n) = \sup_{A \text{ measurable}} |P(A) - P_n(A)| \leq \sqrt{\frac{1}{2} D_{\text{KL}}(p||p_n)} \quad (141)$$

The set of all bin interiors $I = \cup_{i=-\infty}^{\infty} (i a_n, (i+1)a_n)$ is measurable. It holds that $P(I) = 1$ (since we only exclude the zero-set of bin edges to get I from \mathbb{R}). Also, $P_n(I) = 0$ since all the mass is concentrated at the bin edges in p_n , and so $1 \leq \delta_{\text{TV}}(p, p_n) \leq \sqrt{\frac{1}{2} D_{\text{KL}}(p||p_n)}$ regardless of a_n .

Thus $\mathcal{D}_{\text{KL}}(p||p_n) \geq 2 > 0$ regardless of n , but $W_2(p, p_n) \rightarrow 0$.

Intuitively, the construction in Koehler et al. (2021) is related to the mixture above. The vanishing scaling terms from latent to data space in their universality proof squeeze the distribution. The translation terms ensure that this squeezed distribution is distributed over the space such that the error in terms of Wasserstein distance is bounded by the grid length ϵ .

D. Benefits of More Expressive Coupling Blocks

To see what the best improvement for an infinite capacity coupling function can ever be, we make use of the following Pythagorean identity combined from variants in Chen & Gopinath (2000); Cardoso (2003); Draxler et al. (2022):

$$\mathcal{L} = \mathcal{D}_{\text{KL}}(p_\theta(z)||\mathcal{N}(0, I)) = B + \mathbb{E}_{b \sim p(a,b)} [D(b) + J(b) + S(b)]. \quad (142)$$

The symbols $B, D(b), J(b), S(b)$ all denote KL divergences:

The first two terms remain unchanged under a coupling layer: The KL divergence to the standard normal in the passive dimensions $B = \mathcal{D}_{\text{KL}}(p_\theta(b)||\mathcal{N}(0, I_{D/2}))$, which are left unchanged. The *dependence* between active dimensions $D(b) = \mathcal{D}_{\text{KL}}(p_\theta(a|b)||p_\theta(a_1|b) \cdots p_\theta(a_{D/2}|b))$ measures the multivariate mutual information between active dimensions. It is unchanged because each dimension a_i is treated conditionally independent of the others (Chen & Gopinath, 2000).

The remaining terms measure how far each dimension $p_\theta(a_i|b)$ differs from the standard normal: The negentropy measures the divergence to the Gaussian with the same first moments as $p_\theta(a_i|b)$ in each dimension, summing to $J(b) = \sum_{i=1}^{D/2} \mathcal{D}_{\text{KL}}(p_\theta(a_i|b)||\mathcal{N}(m_i(b), \sigma_i(b)))$. Finally, the non-Standardness $S(b) = \sum_{i=1}^{D/2} \mathcal{D}_{\text{KL}}(\mathcal{N}(m_i(b), \sigma_i(b))||\mathcal{N}(0, 1))$ measures how far these 1d Gaussian are away from the standard normal distribution.

Note that the total loss \mathcal{L} is invariant under a rotation of the data. The rotation does, however, affect how that loss is distributed into the different components in Equation (142).

If we restrict the coupling function to be affine-linear $c(a_i; \theta) = sa_i + t$ (i.e. a RealNVP coupling), then this means that also $J(b)$ is left unchanged, essentially because $p_\theta(a_i|b)$ and $\mathcal{N}(m_i(b), \sigma_i(b))$ undergo the same transformation (Draxler et al., 2022, Lemma 1). Only a nonlinear coupling function $c(a_i; \theta)$ can thus affect $J(b)$ and reduce it to $J(\tilde{b}) \leq J(b)$.

Taking the loss difference between two layers, we find Equation (36).

E. Experimental Details

We base our code on PyTorch (Paszke et al., 2019), Numpy (Harris et al., 2020), Matplotlib (Hunter, 2007) for plotting and Pandas (McKinney, 2010; The pandas development team, 2020) for data evaluation.

E.1. Layer-Wise Flow

In experiment on a toy dataset for Figure 1, we demonstrate that a coupling flow constructed layer by layer as in Equation (28) learns a target distribution. We proceed as follows:

We construct a data distribution on a circle as a Gaussian mixture of M Gaussians with means $m_i = (r \cos \varphi_i, r \sin \varphi_i)$, where $\varphi_i = 0, \frac{1}{M}2\pi, \dots, \frac{M-1}{M}2\pi$ are equally spaced, and $\sigma_i = 0.3$. The advantage of approximating the ring with this construction is that this yields a simple to evaluate data density, which we need for accurately plotting $p_\theta(z)$:

$$p(x) = \frac{1}{M} \sum_{i=1}^M \mathcal{N}(x; m_i, \sigma^2 I). \tag{143}$$

We then fit a total 100 layers in the following way: First, treat $p(x)$ as the initial guess for the latent distribution. Then, we build the affine coupling block that maximally reduces the loss using Equation (28). We therefore need to know the conditional mean $m(b)$ and standard deviation $\sigma(b)$ for each b . We approximate this from a finite number of samples N which are grouped by the passive coordinate b into B bins so that N/B samples are in each bin. We then compute the empirical mean m_i and standard deviation σ_i over the active dimension in each bin $i = 1, \dots, B$. According to Equation (28), we define $s_i = \frac{1}{\sigma_i}$ and $t_i = -\frac{1}{\sigma_i}m_i$ at the bin centers and interpolate between bins using a cubic spline. Outside the domain of the splines, we extrapolate with constants s, t with the value of the closest bin. We do not directly optimize over Q , but choose the Q that reduces the loss most out of N_Q random 2d rotation matrices.

We limit the step size of each layer to avoid artifacts from finite training data, by mapping:

$$\tilde{x} = \alpha x + (1 - \alpha) f_{\text{blk}}(x). \tag{144}$$

In addition, we resample the training data from the ground truth distribution after every step to avoid overfitting. We do not explicitly control for the bi-Lipschitz constant of our coupling blocks because we do not encounter any numerical problems.

We choose $N = 2^{26}$, $B = 64$, $M = 20$, $\alpha = 0.5$, $N_Q = 10$. The resulting flow has $64 \cdot 2 \cdot 100 = 12,800$ learnable parameters. Figure 5 shows how the KL divergence vanishes for our layer-wise training, together with Δ_{affine} .

E.2. Volume-Preserving Normalizing Flows

The target distribution is a two-dimensional Gaussian Mixture Model with two modes. The two modes have the same relative weight but different covariance matrices ($\Sigma_1 = I \cdot 0.2$, $\Sigma_2 = I \cdot 0.1$) and means ($m_1 = [-0.5, -0.5]$, $m_2 = [0.5, 0.5]$).

The normalizing flow with a constant Jacobian determinant consists of 15 GIN coupling blocks as introduced in Sorrenson et al. (2019). This type of coupling blocks has a Jacobian determinant of one. To allow for a global volume change, a layer with a learnable global scaling is added after the final coupling block. This learnable weight is initialized as one. For the normalizing flow with variable Jacobian determinant, the GIN coupling is modified by removing the normalization of the scaling factors in the affine couplings. This allows the normalizing flow to have variable Jacobian determinants. In this case, the global scaling block is omitted. To implement the normalizing flow, we use the FrEIA package (Ardizzone et al., 2018a) implementation of the GIN coupling blocks.

In both normalizing flows, the two subnetworks used to compute the parameters of the affine couplings are fully connected neural networks with two hidden layers and a hidden dimensionality of 128. ReLU activations are used. The weights of the linear layers of the subnetworks are initialized by applying the PyTorch implementation of the Xavier initialization (Glorot & Bengio, 2010). In addition, the weights and biases of the final layer of each subnetworks are set to zero.

The networks are trained using the Adam (Kingma & Ba, 2017) with PyTorch’s default settings and an initial learning rate of $1 \cdot 10^{-3}$ which is reduced by a factor of ten after 5000, 10000 and 15000 training iterations. In total, the training ran for 25000 iterations. In each iteration, a batch of size 128 was drawn from the target distribution to compute the negative log likelihood objective. We use a standard normal distribution as the latent distribution.




Figure 5. Empirically, the KL divergence decreases as more coupling blocks are added for the toy distribution considered in Figure 1 (left). At the same time, there is a strong correlation between the loss improvement by a single coupling block $\Delta_{\text{affine}}(p_{\theta}(z))$ and the KL divergence (right). Crucially, the KL divergence does not saturate as the loss improvements become smaller. The coupling flow considered is trained according to the greedy layer-wise training in our proof construction.

For obtaining the optimal distribution $p^*(x)$, we follow the grid procedure in Section 4 and compute the probabilities on a regular 400×400 grid of grid spacing 0.01. The covariance of $p^*(z)$ is computed for the latent scaling layer by sampling 4096 points from the mixture model, moving them according to the volume-preserving flow learned using the grid and computing their empirical covariance matrix. This yields essentially the same scaling as obtained from training the volume-preserving flow.

In order to learn $p^*(r)$ as in Appendix B.2, we sample $5 \cdot 10^8$ samples from ground truth data distribution. These samples are passed through the volume-preserving flow and afterwards the L_2 norm is applied to the latent codes to obtain the latent radii r . We subtract the smallest observed radius from all radii to ensure that the distribution we construct is supported for all $r \geq 0$ and fit a histogram with 4200 bins to the radii. To obtain a smoother distribution, we use the left bin edges of the histogram and the corresponding density values to fit a cubic spline using SciPy’s `interpolate` package (Virtanen et al., 2020). We choose the partition function of the distribution $p_r(r)$ defined by the spline such that it integrates to one. For a given latent code z the latent density can be computed by evaluating p_r at $r = \|z\|_2$ and correcting for the volume at the given radius (see Equation (145)).

$$p(z) = \frac{1}{2\pi\|z\|_2} \cdot p_r(r = \|z\|_2) \quad (145)$$