Denoising Normalizing Flow

Anonymous Author(s)
Affiliation
Address
email

Abstract

Normalizing flows (NF) are expressive as well as tractable density estimation methods whenever the support of the density is diffeomorphic to the entire data-space. However, real-world data sets typically live (close) on low-dimensional manifolds thereby challenging the applicability of standard NF on real-world problems. Here we propose a novel method - called Denoising Normalizing Flow (DNF) - that estimates the density on the low-dimensional manifold while learning the manifold as well. The DNF works in 3 steps. First, it inflates the manifold - making it diffeomorphic to the entire data-space. Secondly, it learns a NF on the inflated manifold and finally it learns a denoising mapping - similarly to denoising autoencoders. The DNF relies on a single cost function and does not require to alternate between a density estimation phase and a manifold learning phase - as it is the case with other recent methods. Furthermore, we show that the DNF can learn meaningful low-dimensional representations from naturalistic images as well as generate high-quality samples.

1 Introduction

Given samples from the data-density $p(x)$, key objectives in probabilistic Machine Learning are 1. estimating $p(x)$ *(density estimation)*, 2. generating new data points from $p(x)$ *(sampling)*, and 3. finding low-dimensional representations of the data *(inference)*. The three main methods used to perform these tasks are *Normalizing Flows* (NFs) ([32]), *Generative Adversarial Networks* (GANs) ([15]), and *Variational autoencoders* (VAEs) ([24]). Among those methods, only the VAE does check all desired objectives by default. However, it does so by sacrificing the sample quality (compared to GANs), and only learning a lower bound on $p(x)$ (rather than the exact value as NFs do). Finding new ways to meet these key objectives, based on either the known main methods or new ones, is an active and important research area.

How can we infer low-dimensional representations using NFs? The standard NF requires the data-density $p(x)$ to have a support diffeomorphic to the entire data-space $\mathbb{R}^D$. However, the manifold hypothesis conjectures that the data-manifold $\mathcal{M}$ lies close to a d-dimensional (Riemannian) manifold embedded in $\mathbb{R}^D$, $d < D$. Thus, unfortunately, a standard NF cannot be used to infer the latent space. Recently, [9] proposed to overcome this limitation by first projecting the data into a d-dimensional space using the first $d-$components of a standard NF, and then use another NF to learn the latent distribution $\pi(u)$. For their method to work they propose different learning schemes, separating the manifold learning from the density estimation.

In this paper, we propose an easy and new way to learn low-dimensional representations with NFs. Our idea is based on the theoretical work derived in [18] where it was shown that by inflating the data-manifold with Gaussian noise $\varepsilon \sim \mathcal{N}(0, \sigma^2 I_D)$, i.e. $\tilde{x} := x + \varepsilon$, the data-density $p(x)$ can be well approximated by learning the inflated distribution $q_\sigma(\tilde{x})$. More concretely, the main result in [18] states sufficient conditions on the choice of noise $q_\sigma(\tilde{x}|x)$ and type of manifold such that $q_\sigma(\tilde{x}) = p(x)q_\sigma(\tilde{x}|x)$. By adding a penalty term to the usual KL-divergence used to learn $q_\sigma(\tilde{x})$, we

Submitted to 35th Conference on Neural Information Processing Systems (NeurIPS 2021). Do not distribute.

ensure that the first d-components of the corresponding flow are noise insensitive and thus encode the manifold. This penalty term is essentially the objective function of a *Denoising autoencoder* (DAE), and thus we call our method *Denoising Normalizing Flow* (DNF).

In summary, our contributions are the following:

1. We propose a new method, the Denoising Normalizing Flow, which combines two previously well-known methods (DAE and NF), and is able to
   - approximate $p(x)$,
   - sample new data $x \sim p(x)$ as a non-linear transformation of $u \sim \mathcal{N}(u; 0, I_d)$,
   - infer low-dimensional latent variables $u \sim p(u|x)$ given $x \sim p(x)$.
2. We demonstrate on naturalistic data that our method learns meaningful latent representations without sacrificing the sample quality.

Notations: We adapt the notation used in [9]. To further simplify it and avoid clutter, we denote the determinant of the Gram matrix of g evaluated at $g^{-1}(x)$ as

$$\det G_g(x) := \det \left(J_g(g^{-1}(x))^T J_g(g^{-1}(x)) \right) \tag{1}$$

where $J_g(g^{-1}(x))^T$ is the transpose of the Jacobian of $g : \mathbb{R}^d \to \mathbb{R}^D$, $d \leq D$, evaluated at $g^{-1}(x)$.

2 Problem statement

In the following, we are going to show why standard NFs are not suited to infer low-dimensional representations of the given data. We end the section with the research question we are going to study in the subsequent.

General setting: In generative modeling, it is assumed that the data $x \in \mathcal{M} \subset \mathbb{R}^D$ is generated by a non-linear, injective transformation g of some latent variables $u \in \mathcal{U} \subset \mathbb{R}^d$, i.e. $x = g(u)$ where $u \sim \pi(u)$. Typically, $d < D$ and the latent distribution $\pi(u)$ is assumed to be Gaussian, $\pi(u) = \mathcal{N}(u; 0, I_d)$. Hence, the latent random variable u generates x, and the data-density $p(x)$ evaluated at $x \in \mathbb{R}^D$ is given by

$$p(x) = \int_{\mathcal{U}} \pi(u) \delta(x - g(u)) du, \tag{2}$$

where δ denotes the Dirac function, see [3]. Whenever g is a diffeomorphism, we have for $x \in \mathcal{M}$ that

$$p(x) = |\det G_g(x)|^{-\frac{1}{2}} \pi(g^{-1}(x)). \tag{3}$$

If $d = D$, the target density $p(x)$ can be learned, in principle, exactly using a NF ([19]). In general, a NF is a diffeomorphism mapping $\mathcal{U} \subset \mathbb{R}^d$ to $\mathcal{M} \subset \mathbb{R}^D$, see [30] or [25] for some recent reviews. Denoting this diffeomorphism as $g_\theta : \mathcal{U} \to \mathcal{M}$ and its parameters as θ, the induced density on $\mathcal{M}$ is given by

$$p_\theta(x) = |\det G_{g_\theta}(x)|^{-\frac{1}{2}} p_u(g_\theta^{-1}(x)), \tag{4}$$

where $p_u(u)$ is a known reference density (usually set to be standard Gaussian). The parameters θ are updated such that the KL-divergence between $p(x)$ and $p_\theta(x)$,

$$D_{\mathrm{KL}}(p(x)||p_\theta(x)) = -\mathbb{E}_{x \sim p(x)}[\log p_\theta(x)] + const. \tag{5}$$

is minimized. Thus, to learn θ efficiently, one needs to evaluate $T_1(x) := \log p_u(g_\theta^{-1}(x))$ and $T_2(x) := \frac{1}{2} \log |\det G_{g_\theta}(x)|$ efficiently.

Topological constraints: The evaluation of $T_1(x)$ is efficient since we are free to choose the reference measure p_u, and $g_\theta^{-1}(x)$ is the forward pass of a neural network constructed to be a diffeomorphism ([11],[19]). Therefore, the majority of the NF literature focuses on designing clever flow-architectures to be able to calculate $T_2(x)$ efficiently without sacrificing the flow's expressiveness (i.e. the size of the space of diffeomorphisms able to learn). However, so far these architectures are constructed for $d = D$ since in this case, the Jacobian of g_θ is a square matrix, and thus $T_2(x)$ becomes

$$\frac{1}{2} \log |\det G_{g_\theta}(x)| = \log |\det J_{g_\theta}(g_\theta^{-1}(x))|. \tag{6}$$

Hence, to calculate the Gram determinant of g_θ efficiently is equivalent to calculate the determinant of the Jacobian of g_θ efficiently. Popular choices are to construct g_θ such that J_{g_θ} is lower triangular, as in this case, the determinant is simply the product of diagonal elements. Unfortunately, for $d < D$, J_{g_θ} is not a square matrix, and the full Gram determinant needs to be calculated. This makes NFs unsuitable for finding low-dimensional representations u of high-dimensional data points for large d as the computational complexity to calculate G_{g_θ} is $\mathcal{O}(d^2 D) + \mathcal{O}(d^3)$.
Research question: As mentioned by [9], finding ways to design g_θ such that $T_2(x)$ can be efficiently calculated for $d < D$ is an interesting research question. Here, we address it from a different angle.

Question 1 *Let $\mathcal{M}$ be a $d-$dimensional Riemannian manifold embedded in $\mathbb{R}^D$ through g. Can we construct a diffeomorphism $g_\theta : \mathbb{R}^d \to \mathcal{M}$ to*

1. *generate $x \sim p(x)$ in 2 steps: (a) generate $u \sim \mathcal{N}(u; 0, I_d)$ and (b) set $x = g_\theta(u)$?*
2. *infer $u \in \mathbb{R}^d$ such that $x = g_\theta(u)$?*
3. *approximate $\det G_g(x)$ efficiently?*

3 Denoising Normalizing Flow

We answer Question 1 based on the theoretical work developed in [18]. First, we briefly review this work, and then introduce the DNF.

Preliminaries: In Section 2, we discussed why classical NFs are not suited to infer low-dimensional representations of the given data. Also, if one is only interested in the value of $p(x)$, standard NF cannot be used. Intuitively, a standard NF f_ψ with parameters ψ simply squeezes or expands a volume element where the net change is given by its Jacobian determinant. A volume element of a $d-$dimensional manifold is $d-$dimensional and thus has $D-$dimensional Lebesgue measure 0. Thus, we are asking f_ψ to expand a $d-$dimensional volume to a $D-$dimensional one which will lead to a degeneration of $|\det G_{f_\psi}(x)|$ and manifest in numerical instabilities. Therefore, [18] proposed to inflate the manifold by adding Gaussian noise to the data points. This inflated manifold is $D-$dimensional and thus a usual flow can be used to learn the corresponding density. Their main result states sufficient conditions on the choice of noise and type of manifold $\mathcal{M}$, such that the learned inflated distribution can be deflated, and $p(x)$ is exactly retrieved.
More precisely, given a random variable $x \sim p(x)$ with probability measure $\mathbb{P}_X$ and taking values in a $d-$dimensional manifold $\mathcal{M}$ embedded in $\mathbb{R}^D$, if we add some noise $\varepsilon \sim \mathcal{N}(\varepsilon; 0\sigma^2 I_D)$ to it, the resulting new random variable $\tilde{x} = x + \varepsilon$ has the following density:

$$q_\sigma(\tilde{x}) = \int_{\mathcal{X}} q_\sigma(\tilde{x}|x) d\mathbb{P}_X(x). \tag{7}$$

In [18], it was shown that if (a) the noise is only added in the $(D-d)-$dimensional normal space N_x in x, i.e. $q_\sigma(\tilde{x}|x) = q_\sigma(\tilde{x} \cap N_x|x)$, (b) the noise magnitude σ is sufficiently small, and (c) the manifold $\mathcal{M}$ is sufficiently smooth and disentangled, the resulting inflated distribution takes the following product form:

$$q_\sigma(\tilde{x}) = p(x) q_\sigma(\tilde{x}|x), \tag{8}$$

where x is almost surely uniquely determined by $\tilde{x}$ as the orthogonal projection of $\tilde{x}$ on $\mathcal{M}$. For this projection to be well-defined, a sufficient condition is that the manifolds reach number is finite ([5]). [1] A manifold where almost every point $\tilde{x} \sim q(\tilde{x})$ in the inflated set has a unique projection on $\mathcal{M}$ was called $Q-$normally reachable in [18], where Q denotes the collection of noise distributions $q(\tilde{x}|x)$. Their main Theorem proves that for any $Q-$normally reachable manifold equation (8) holds.

It was also shown that for the case where $d \ll D$ (as it is generally assumed for high-resolution images), full Gaussian noise is an excellent approximation for a Gaussian in the normal space.

Main idea: We use a standard flow f_ψ to learn $q_\sigma(\tilde{x})$ such that the corresponding density has the product form of equation (8), and the first $d-$components u of the flow's output $f^{-1}(\tilde{x})$ are noise-insensitive whereas the remaining $(D-d)-$components v remain noise-sensitive. Thus, intuitively, we want the first $d-$components to *denoise* the inflated data.

[1] Informally, this reach condition ensures that a manifold is learnable through samples.

DNF: Let $f_\psi : \mathbb{R}^D \to \mathbb{R}^D$ be a standard flow with reference measure $p_z(z)$. We denote the first $d-$components of the flows output as u, and the remaining ones as v, i.e. $(u, v)^T = f_\psi^{-1}(\tilde{x})$. More formally,

$$u = u(\tilde{x}) = \text{Proj}_u(f_\psi^{-1}(\tilde{x})), \quad v = v(\tilde{x}) = \text{Proj}_v(f_\psi^{-1}(\tilde{x})) \tag{9}$$

with $\text{Proj}_u(z) = (z_1, \ldots, z_d)^T$ and $\text{Proj}_v(z) = (z_{d+1}, \ldots, z_D)^T$ for $z \in \mathbb{R}^D$. Note that $q_\sigma(\tilde{x}|x)$ is a $(D-d)-$dimensional Gaussian distribution with covariance $\sigma^2 I_{D-d}$. Therefore, as reference measure p_z, we choose $p_z((u, v)) = p_u(u)p_v(v)$, with $p_v(v) = \mathcal{N}(v; 0, \sigma^2 I_{D-d})$ modelling the noise-sensitive part of equation (8), and $p_u(u)$ the noise-insensitive part. For that, we want the image $f_\psi(u, 0)$ to be in the manifold. Therefore, we embed u back in $\mathbb{R}^D$ by padding the missing coordinates with 0,

$$\text{Pad}(u) = (u, \underbrace{0, \ldots, 0}_{(D-d)\text{-times}})^T, \tag{10}$$

such that the operation $\text{Pad}(\text{Proj}_u((u, v))) = (u, 0)^T$ ignores the noise-sensitive part in the latent variable. This operator allows us to define a denoising function $r_\psi(\tilde{x})$ as

$$r_\psi(\tilde{x}) := (f_\psi \circ \text{Pad} \circ \text{Proj}_u \circ f_\psi^{-1})(\tilde{x}) \tag{11}$$

and we regularize ψ by minimizing

$$\mathcal{C}(\psi) := \mathbb{E}_{x \sim p(x)} \mathbb{E}_{\tilde{x} \sim q_\sigma(\tilde{x}|x)} ||x - r_\psi(\tilde{x})||^2 \tag{12}$$

where $|| \cdot ||$ denotes the L_2 norm.

We have not specified the reference measure $p_u(u)$. To facilitate the disentanglement of noise-insensitivity and noise-sensitivity in the u and v variables, we transform u with yet another flow h_ϕ with paramters ϕ and reference measure $p_{\tilde{u}}$ (e.g. standard Normal). Now, our sampling procedure looks as follows: 1. sample $\tilde{u} \sim p_{\tilde{u}}(\tilde{u})$, 2. apply h_ϕ to obtain a sample u from $p_u(u)$, i.e. $u = h_\phi(\tilde{u})$, and 3. set $x = f_\psi(\text{Pad}(u))$. Denoting $\theta = [\psi, \phi]$, our final model is

$$\begin{aligned} q_\theta(\tilde{x}) =& |\det G_{f_\psi}(\tilde{x})|^{-\frac{1}{2}} p_u(u(\tilde{x})) p_v(v(\tilde{x})) \\ =& |\det G_{f_\psi}(\tilde{x})|^{-\frac{1}{2}} |\det G_{h_\phi}(u(\tilde{x}))|^{-\frac{1}{2}} p_{\tilde{u}}(h_\phi^{-1}(u(\tilde{x}))) p_v(v(\tilde{x})). \end{aligned} \tag{13}$$

Note that the reconstruction loss, equation (12), is essentially the objective function for the *denoising autoencoder* introduced in [1]. Therefore, we call our method *denoising Normalizing flow* (DNF) and it is trained on

$$\mathcal{L}_{\text{DNF}}(\theta) := D_{\text{KL}}(q_\sigma(\tilde{x}) || q_\theta(\tilde{x})) + \lambda \mathcal{C}(\psi) \tag{14}$$

where $\lambda > 0$ is the penalty hyperparameter and is trading the density estimation with the manifold learning. A graphical description of the DNF model is given in Figure 1 (a), and an algorithmic description in the DNF Algorithm below.

Answer to Question 1: If $\mathcal{L}_{\text{DNF}}(\theta) = 0$, the reconstruction error expressed in equation (12) is 0. Thus, the generative story of the DNF is exactly the one described in point 1. of question 1 with $g_\theta := f_\psi \circ \text{Pad} \circ h_\phi$. The inverse, $h_\phi^{-1} \circ \text{Proj}_{1:d} \circ f_\psi^{-1}$, can be used to infer $\tilde{u}$ s.t. point 2. holds. Finally, to show the third claim, we additionally assume that $\mathcal{U}$ (which is the domain of the manifold generating function g) is diffeomorphic to $\mathbb{R}^d$, s.t. without loss of generality we set $\pi(u) = \mathcal{N}(u; 0, I_d)$. Then, we exploit equation (8) and calculate $|\det G_g(x)|$ efficiently with the help of f_ψ and h_ϕ, see Proposition 1 and its proof in the supplementary.

Proposition 1 *Let $\mathcal{M}$ be a $d-$dimensional Riemannian manifold embedded in $\mathbb{R}^D$ through $g : \mathbb{R}^d \to \mathcal{M}$. Let $x \sim p(x)$ be generated by $g(u)$, where $u \sim \mathcal{N}(u; 0, I_d)$, i.e. $x = g(u)$. Assume that equation (8) holds. If $\mathcal{L}_{DNF}(\theta) = 0$, then*

$$|\det G_g(x)| = |\det G_{f_\psi}(x)| |\det G_{h_\phi}(x)|. \tag{15}$$

4 Related work

Autoencoders: An autoencoder learns to compress the data x using an encoder $f_\psi^{-1}(x)$, and then to reconstruct x using a decoder $g_\theta(z)$. As the AE is only trained on the reconstruction error

DNF Algorithm: Training of denoising normalizing flow. For simplicity, we show a stochastic gradient descent with a constant learning rate. Alternative optimization methods and learning rate schedules can be easily adapted.

Require: Learning rate α, penalty parameter λ, inflation variance σ^2, batch size n, number of epochs E.
Initialize: Parameters ψ and ϕ for flows f_ψ and h_ϕ.

```
while θ = [ψ, φ] has not converged do
  for e = 1 to E do
    for i = 1 to n do
      Sample: x_i ~ p(x)                                          # sample data
      Inflate: x̃_i = x_i + ε_i, where ε_i ~ N(0, σ² I_D)          # add noise
      (u_i, v_i) ← f_ψ^{-1}(x̃_i)                                  # project on u and v
      x'_i ← f_ψ(u_i, 0)                                          # reconstruct x
      ũ_i ← h_φ^{-1}(u_i)                                         # transform u
    end for
    L_DNF ← (1/n) Σ_i^n log p_ũ(ũ_i) − log|det J_{h_φ}(ũ_i)| + p_v(v_i) − log|det J_{f_ψ}(u_i, v_i)|
            +λ||x_i − x'_i||²        # calculate log q_θ(x̃) and add reconstruction error
    θ ← θ − α∇_θ L_DNF                                            # update model parameters
  end for
end while
```

$||x - g_\phi(f_\psi^{-1}(x))||^2$, it can't be used to generate new data or estimate $p(x)$. However, if the input is corrupted by Gaussian noise ε, an optimal AE that can reconstruct the uncorrupted input depends on the gradient of the data-loglikelihood ([1]). This can be exploited to estimate $p(x)$ ([6]). A variational autoencoder (VAE) ([24]) is a stochastic version of an AE. More concretely, a lower bound on the data log-likelihood, known as ELBO or free energy, is maximized to learn parameters ϕ and θ such that the true conditional distributions $p(x|z)$ and $p(z|x)$ are approximated by $p_\theta(x|z) = \mathcal{N}(x; g_\theta(z), I_D)$ and $q_\phi(z|x) = \mathcal{N}(x; f_\psi(x), \sigma_\theta(x) I_D)$, respectively. If one is not interested in $p(x)$ but still wants to generate new data using an AE, regularization terms can be exploited ([26],[35]). Alternatively, a NF can be used to normalize the learned latent variables $u = f_\psi(x)$. Thus, after training a usual AE, the latent density associated with the latent variable u is learned using a standard NF. This probabilistic autoencoder (PAE) was recently introduced by [8].

NFs based models: Recently, a few attempts have been made to overcome the topological constraint of NFs in terms of density estimation ([32], [33], [9],[29],[18]), sampling ([4],[10],[22],[9]), and inference ([9]). If the manifold is known, i.e. its atlas is given, a usual flow in $\mathbb{R}^d$ is sufficient to learn the data-density $p(x)$ as we can use these charts to push-back $p(x)$ to $\mathbb{R}^d$. This was first done in [14] for a manifold consisting of a single chart g. However, if g is not known, only recently some methods were proposed to learn it ([9],[4]). We review these methods in the following, highlight their differences to the DNF and illustrate them in Figure 1.

Pseudo invertible encoder (PIE): The idea of [4] is to define the manifold as a level set of a usual NF f. For that, they propose to treat the first $d-$latent variables u differently than the remaining $D-d$ variables v, by using different reference measures for u and v, respectively, i.e. $f^{-1}(x) = z = (u, v)^T \in \mathbb{R}^d \times \mathbb{R}^{D-d}$ and $p_z(z) = p_u(u)p_v(v)$. The gist is very similar to the DNF. However, there is no incentive for f to encode the manifold in u which manifests in poor sample quality (see [9]). In addition, this approach does not work for data living exactly on a low-dimensional manifold as the Jacobian determinant of f degenerates. In Figure 1, we depict a version of the PIE model where u is additionally transformed with a flow h. Note that a PIE model is a DNF with $\lambda = \sigma = 0$.

Manifold flow: The idea of [9] is to learn the embedding g directly by first mapping the data into the latent space using a usual NF f, projecting it to the first $d-$coordinates u, and finally transform these variables using yet another flow h in $\mathbb{R}^d$ with reference measure $p_{\tilde{u}}$, see Figure 1 for a depiction of their $\mathcal{M}-$flow. The resulting density,

$$p_{\mathcal{M}}(x) = p_{\tilde{u}}(h^{-1}(u(x)))|\det G_h(u(x))|^{-\frac{1}{2}}|\det G_g(u(x))|^{-\frac{1}{2}}, \tag{16}$$

is indeed a density defined only on the manifold. As noted in [9], calculating the Gram determinant of g may be very slow for large d. Therefore, they propose to train the parameters of g using the

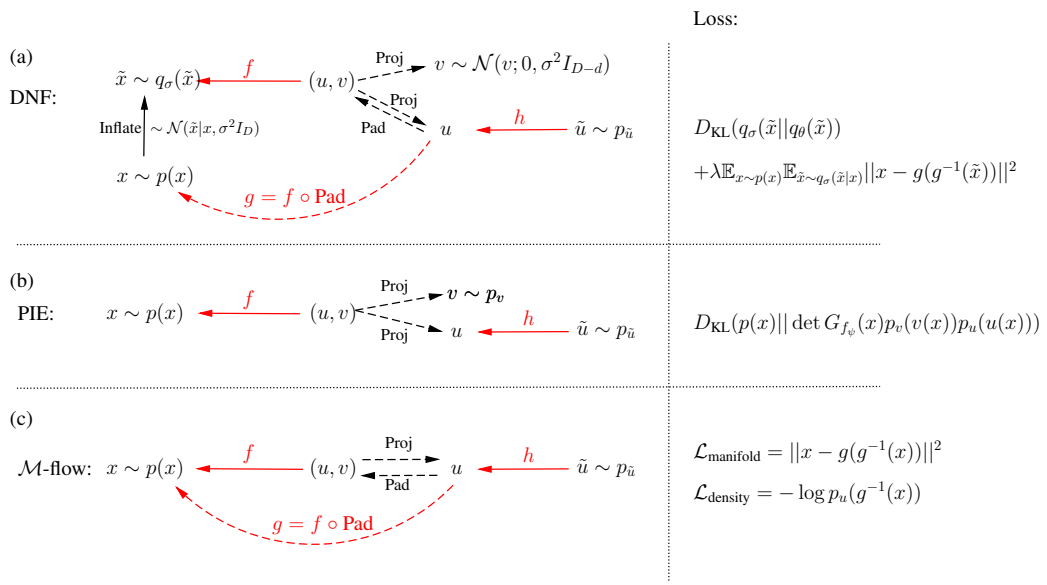

Figure 1: Schematic overview of different NF-based methods (PIE, $\mathcal{M}$-flow, DNF) (adopted from [9]). **Left:** Graphical model. Red solid lines are NFs, red dashed lines the corresponding generator. Black Dashed lines describe the projection or padding operation, respectively, as described in the main text. **Right:** The losses used to train the model.

reconstruction error $||x - g^{-1}(g(x))||_2^2$ only, and the parameters of h using the usual D_{KL} objective for NFs. They propose several training strategies to ensure that g encodes the manifold (manifold phase), and h learns the density (density phase). One strategy alternates between the manifold- and density phase (alternate training) for every training epoch. Another strategy first trains g on the reconstruction error and then learns the density through h (sequential training). Note that a PAE is also trained using a sequential training scheme, as opposed to the DNF which combines the manifold- and density learning phase into a single cost function, equation (14). Another difference to the DNF is that the Gram determinant of f is neither used to train the $\mathcal{M}-$flow nor to estimate $p(x)$.

5 Results

In this chapter, we confirm experimentally our claims made at the end of Section 1, formalized in question 1, and answered at the end of Section 3. First, we show that the DNF circumvents the degeneracy problem of a standard NF for manifold-supported densities. Afterward, we show that the DNF can learn meaningful latent representations for naturalistic images, without sacrificing the sample quality compared to the $\mathcal{M}-$flow (see Section 4), PAE and VAE. Our method depends on two hyperparameters, the noise magnitude σ^2 and the penalty coefficient λ. σ^2 needs to be sufficiently small such that equation (8) approximately holds (see [18] for more details on the choice of σ). If $\lambda = 0$, only the density is learned, and if $\lambda \gg 1$ the manifold learning dominates.

We closely follow the experimental setting of [9]. For the $\mathcal{M}-$flow and DNF, we use the same network architectures and training protocols. These architectures are based on affine-coupling layers ([11]), neural splines ([12]), and trainable permutations ([23]). For the PAE, we use convolutional neuronal networks for the encoder and decoder, respectively. For images, we use a recently developed variant of the VAE, the InfoMax-VAE ([31]), which uses insights from information theory to learn meaningful latent representations. We refer to the supplementary for more training details. [2]

5.1 Density estimation

[2] Our main code is based on the original $\mathcal{M}-$flow implementation made public by the authors of [9] under the MIT license. For the InfoMax-VAE, we use the official implementation made public by the authors of [31] under the Apache License 2.0

We consider a $1-$dimensional manifold embedded in $\mathbb{R}^2$, a thin spiral. We draw the latent variable u from an exponential distribution with scale 0.3 and generate the spiral through $g(u) = u(\cos(u), \sin(u))^T$ (Figure 2 top left).

Figure 2: Density estimation

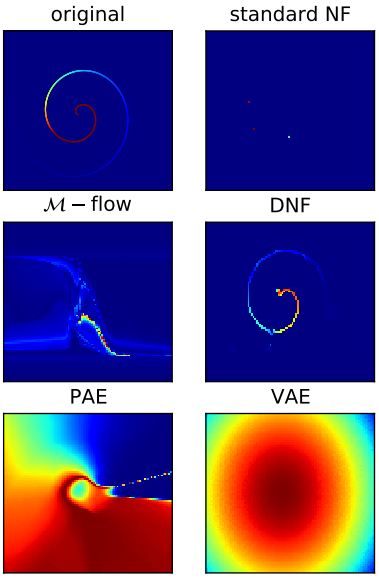

We train a standard NF *(upper right)*, $\mathcal{M}-$flow *(middle left)*, DNF *(middle right)*, PAE *(lower left)*, VAE *(lower right)* with similar architectures (if possible) on 100 epochs with a batch size of 100. As we can see in Figure 2, the Jacobian determinant of the standard NF degenerates, and the learned density collapses into single points. Surprinsingly, the $\mathcal{M}-$flow fails to learn $p(x)$ as well, no matter which training schedule we use (we display the result of the sequential training). The PAE does not learn $p(x)$ as expected. The standard VAE simply equates $p(x)$ with the Gaussian prior on the latent variables (we take a point estimate of the ELBO to approximate $p(x)$).

For the DNF, we set $\sigma^2 = 0.01$ and $\lambda = 1$. Note that the DNF is trained on the inflated density $q_\sigma(\tilde{x})$, and encodes the noisy part in $v(\tilde{x})$. Thus, we approximate $p(x) \approx q_\sigma(x)/q_\sigma(x|x)$ with $q_\theta(x)/q_\sigma(x|x)\chi_{\{v(x)\leq\sigma/3\}}$, where χ denotes the characteristic function, and $q_\sigma(x|x) = \sqrt{\pi\sigma^2}$. The fact that the $\mathcal{M}-$flow fails to learn $p(x)$ and the DNF succeeds, suggests the regulatory importance of the Gram determinant of f_ψ (which is only used from the DNF, and reflects the main difference between those two methods next to the training schemes).

The densities are evaluated on a 100×100 grid. For that, the $\mathcal{M}-$flow calculates the Gram determinant of g (see equation (16) and Figure 1 (c)) for each point which takes 288 seconds on a CPU.[3] On the same device, the DNF needs only 5 seconds to evaluate the density on the grid. This illustrates the drawback of needing to calculate the full Gram determinant (see Section 2).

5.2 StyleGAN image manifold

One drawback of both the $\mathcal{M}-$flow and DNF is that the true manifold dimension d must be known beforehand. For real-world datasets, d is unknown. Therefore, [9] uses a StyleGAN2 model ([21]) trained on the FFHQ dataset ([20]) to generate an $n-$dimensional manifold by only varying the first n latent variables while keeping the remaining fixed.

$\mathbf{n = 2}$: We want to show that the DNF learns meaningful latent representations. For that, we first train a DNF on 10^4 images using 100 epochs with $\sigma^2 = 0.1$ and $\lambda = 1000$. Then, we generate an image grid by varying the $2-$dimensional latent variables $\tilde{u}$ and mapping them to the data space using the learned mapping g_θ. In Figure 5.2, we see on the top left the original grid obtained by the StyleGAN2 model, and on the top right the grid generated by the DNF. A smooth mapping from latent to data space is learned. The PAE (lower left), does not learn such a smooth mapping which may be due to the lack of learning $p(x)$, in contrast to the InfoMax-VAE (lower right). However, the latter lacks variability which shows that the learned posterior $p(z|x)$ does not match the prior $\mathcal{N}(z; 0, I_d)$.

To further evaluate the quality of generated test samples, we display the Fréchet inception distance (FID score) [4] ([17], [28]) in tabular 5.2 along with the mean reconstruction error which measures the manifold learning. We slightly outperform the $\mathcal{M}-$flow in terms of FID score, and significantly in terms of the mean reconstruction error. The PAE and InfoMax-VAE perform worse compared to the $\mathcal{M}-$flow and DNF indicating a suboptimal choice of latent dimension for these models.

[3] The Gram matrix can be calculated using automatic differentiation.

[4] The FID is the Wasserstein-2 distance between two Gaussians. The mean values and covariance matrices are obtained as sample estimates of the original and model data, respectively. However, instead of using the pixel values, one uses the outcomes of the next to the last layer of the Inception v3 ([34]) image classifier trained on the corresponding data set.

Figure 3: StyleGAN images (top left) are used to train a DNF (top right), PAE (bottom left), and InfoMax-VAE (bottom right).

	StyleGAN $n = 2$		StyleGAN $n = 64$	
Model	FID	Mean reconstr. error	FID	Mean reconstr. error
$\mathcal{M}-$flow	5.32 ± 0.19	349.48 ± 12.73	19.41	1020.78
DNF	$\mathbf{4.07} \pm 0.1$	$\mathbf{215.1} \pm 4.6$	**17.97**	**887.72**
InfoMax-VAE	38.35 ± 3	11878.3 ± 1256.44	53.2	16397.54
PAE	31.74 ± 0.93	1106.64 ± 688.55	–	–

Table 1: FID and mean reconstruction error of the $\mathcal{M}-$flow and DNF on the StyleGAN image manifold for $n = 2$ and $n = 64$. For $n = 2$, we train 3 models with different initialization and report the mean and standard deviation.

$\mathbf{n = 64}$: We train the models on $2 \cdot 10^4$ images for 200 epochs. In Figure 4, we show samples from the original dataset (top left), $\mathcal{M}-$flow (second row), DNF (third row), and InfoMax-VAE (last row).[5] We show how smoothly these models interpolate linearly in the latent space (last 5 images). For that, we sample two $64-$dimensional latent variables, linearly connect them and display the generated images from the 3 intermediate latent points. We compare the FID and mean reconstruction error of the models in tabular 5.2. The DNF outperforms the $\mathcal{M}-$flow and InfoMax-VAE in terms of FID and mean reconstruction error.

[5]The same PAE model used to learn the StyleGAN $2d$ dataset did not converge.

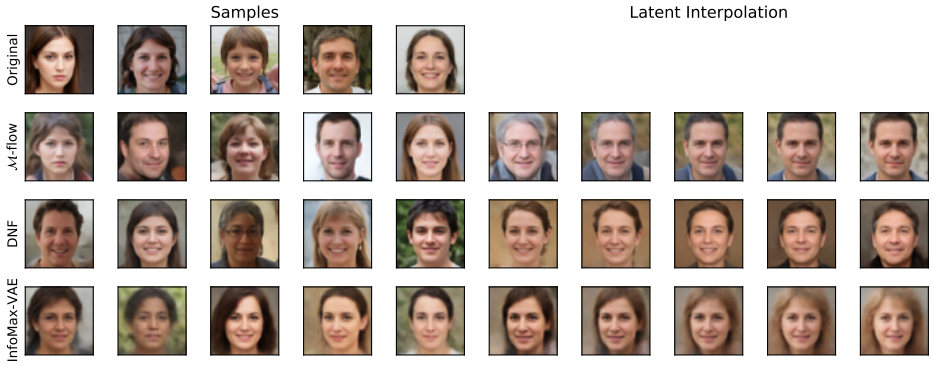

Figure 4: **Samples** from StyleGAN, $\mathcal{M}-$flow, DNF, and InfoMax-VAE (first 5 figures). **Latent interpolation** of $\mathcal{M}-$flow (top), DNF (middle), and InfoMax-VAE (remaining 5 figures).

6 Summary and discussion

Our model is based on the theoretical work established by [18] leading to a natural combination of NFs and DAE - the DNF. In contrast to similar methods, a DNF is trained on a single objective function combining manifold- and density learning. To learn a density supported on a low-dimensional manifold with a NF, one needs to compute the flows Gram determinant. The DNF circumvents this necessity and can be used to approximate it. We have pigeonholed the DNF into the literature, and compared its performance on naturalistic images with related methods ($\mathcal{M}-$ flow, PAE, VAE). Among those methods, we have seen that the DNF generates images with the highest quality (in terms of FID), and reconstructs a given input with the lowest L_2 distance.

It is well known that adding noise in the input increases the performance in supervised learning tasks ([7],[2],[27]). However, normally this comes with the price of lower sample quality or worse density estimation. The DNF has the potential to avoid this trade-off, and proves experimentally that adding noise to images can lead to better performance in terms of sampling quality and manifold learning.

Extensions: The theoretical foundation of the DNF is equation (8) which requires the manifold to be sufficiently smooth, and the noise to be added in the manifold's normal space. Although for the special case where $d \ll D$ standard Gaussian noise is an excellent approximation for a Gaussian in the normal space, more research on how to sample in a manifold's normal space is needed and could improve the performance of the DNF. Essentially, the DNF learns to encode the manifold in the u component and the off-manifold direction in the v component. Arguably, a pre-trained DNF could be used to improve the sampling ability in the normal space.

Another limitation of the DNF is that the manifold dimension d needs to be known. Nevertheless, if $p(x)$ is supported on a low-dimensional manifold, to the best of our knowledge the DNF together with the $\mathcal{M}-$flow are the only existing methods to learn $p(x)$ based on NF. Hence, in contrast to other methods which need to know the exact manifold, the DNF and $\mathcal{M}-$flow can still be used when d is estimated from samples. The latter is an active research area ([16],[13]).

Other recently developed methods (such as PAE or $\mathcal{M}-$flow) separate the manifold from the density learning. We have seen that this separation is not necessary, and have used different heuristics to evaluate the quality of learning. For the value of $p(x)$, we visualized the learned density (Figure 2). For the quality of latent representations, we generated an image grid (Figure 5.2), and for the smoothness in the latent space, we generated an image path (Figure 4). We measured the sampling quality using the FID, and the manifold-learning using the mean reconstruction error. Given these different heuristics, there is an aspiration for a unified performance criterion.

Broader Impact: To compress increasingly high-dimensional data with the least loss of information as possible is becoming increasingly important. The DNF ties in with the $\mathcal{M}-$flow and shows that such compression is possible, even for naturalistic images. Similar to all other generative models, a possible negative impact of DNFs would be the generation of fake data. On a more positive note, the DNF can be used to improve out-of-distribution detection or even increase the robustness to adversarial attacks which are both essential for reliable societal applications of Machine Learning.

References

[1] Guillaume Alain and Yoshua Bengio. What regularized auto-encoders learn from the data-generating distribution. *Journal of Machine Learning Research*, 15:3743–3773, 2014.

[2] Guozhong An. The effects of adding noise during backpropagation training on a generalization performance. *Neural Computation*, 8(3):643–674, 1996.

[3] Chi Au and Judy Tam. Transforming variables using the dirac generalized function. *The American Statistician*, 53(3):270–272, 1999.

[4] Jan Jetze Beitler, Ivan Sosnovik, and Arnold Smeulders. Pie: Pseudo-invertible encoder. 2018.

[5] Clément Berenfeld and Marc Hoffmann. Density estimation on an unknown submanifold. *arXiv preprint arXiv:1910.08477*, 2019.

[6] Siavash A Bigdeli, Geng Lin, Tiziano Portenier, L Andrea Dunbar, and Matthias Zwicker. Learning generative models using denoising density estimators. *arXiv preprint arXiv:2001.02728*, 2020.

[7] Chris M Bishop. Training with noise is equivalent to tikhonov regularization. *Neural Computation*, 7(1):108–116, 1995.

[8] Vanessa Böhm and Uroš Seljak. Probabilistic auto-encoder. *arXiv preprint arXiv:2006.05479*, 2020.

[9] Johann Brehmer and Kyle Cranmer. Flows for simultaneous manifold learning and density estimation. *Advances in Neural Information Processing Systems*, 33, 2020.

[10] Edmond Cunningham, Renos Zabounidis, Abhinav Agrawal, Ina Fiterau, and Daniel Sheldon. Normalizing flows across dimensions. *arXiv preprint arXiv:2006.13070*, 2020.

[11] Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using real nvp. *arXiv preprint arXiv:1605.08803*, 2016.

[12] Conor Durkan, Artur Bekasov, Iain Murray, and George Papamakarios. Neural spline flows. In *Advances in Neural Information Processing Systems*, pages 7511–7522, 2019.

[13] Elena Facco, Maria d'Errico, Alex Rodriguez, and Alessandro Laio. Estimating the intrinsic dimension of datasets by a minimal neighborhood information. *Scientific reports*, 7(1):1–8, 2017.

[14] Mevlana C Gemici, Danilo Rezende, and Shakir Mohamed. Normalizing flows on riemannian manifolds. *arXiv preprint arXiv:1611.02304*, 2016.

[15] Ian J Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *arXiv preprint arXiv:1406.2661*, 2014.

[16] Matthias Hein and Jean-Yves Audibert. Intrinsic dimensionality estimation of submanifolds in rd. In *Proceedings of the 22nd international conference on Machine learning*, pages 289–296, 2005.

[17] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 6629–6640, 2017.

[18] Christian Horvat and Jean-Pascal Pfister. Density estimation: an inflation-deflation approach, 2021.

[19] Chin-Wei Huang, David Krueger, Alexandre Lacoste, and Aaron Courville. Neural autoregressive flows. volume 80 of *Proceedings of Machine Learning Research*, pages 2078–2087, Stockholmsmässan, Stockholm Sweden, 10–15 Jul 2018. PMLR.

[20] T Karras, S Laine, and T Aila. A style-based generator architecture for generative adversarial networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.

[21] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8107–8116. IEEE Computer Society, 2020.

[22] Hyeongju Kim, Hyeonseung Lee, Woo Hyun Kang, Joun Yeop Lee, and Nam Soo Kim. Softflow: Probabilistic framework for normalizing flow on manifolds. *Advances in Neural Information Processing Systems*, 33, 2020.

[23] Diederik P Kingma and Prafulla Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. *arXiv preprint arXiv:1807.03039*, 2018.

[24] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.

[25] Ivan Kobyzev, Simon Prince, and Marcus Brubaker. Normalizing flows: An introduction and review of current methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.

[26] Abhishek Kumar, Ben Poole, and Kevin Murphy. Regularized autoencoders via relaxed injective probability flow. In *International Conference on Artificial Intelligence and Statistics*, pages 4292–4301. PMLR, 2020.

[27] Bai Li, Changyou Chen, Wenlin Wang, and Lawrence Carin. Certified adversarial robustness with additive noise. *arXiv preprint arXiv:1809.03113*, 2018.

[28] Mario Lucic, Karol Kurach, Marcin Michalski, Sylvain Gelly, and Olivier Bousquet. Are gans created equal? a large-scale study. 2018.

[29] Emile Mathieu and Maximilian Nickel. Riemannian continuous normalizing flows. *arXiv preprint arXiv:2006.10605*, 2020.

[30] George Papamakarios, Eric Nalisnick, Danilo Jimenez Rezende, Shakir Mohamed, and Balaji Lakshminarayanan. Normalizing flows for probabilistic modeling and inference. *arXiv preprint arXiv:1912.02762*, 2019.

[31] Ali Lotfi Rezaabad and Sriram Vishwanath. Learning representations by maximizing mutual information in variational autoencoders. In *2020 IEEE International Symposium on Information Theory (ISIT)*, pages 2729–2734. IEEE, 2020.

[32] Danilo Rezende and Shakir Mohamed. Variational inference with normalizing flows. In *International Conference on Machine Learning*, pages 1530–1538. PMLR, 2015.

[33] Danilo Jimenez Rezende, George Papamakarios, Sébastien Racaniere, Michael Albergo, Gurtej Kanwar, Phiala Shanahan, and Kyle Cranmer. Normalizing flows on tori and spheres. In *International Conference on Machine Learning*, pages 8083–8092. PMLR, 2020.

[34] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015.

[35] Wenju Xu, Shawn Keshmiri, and Guanghui Wang. Adversarially approximated autoencoder for image generation and manipulation. *IEEE Transactions on Multimedia*, 21(9):2387–2396, 2019.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes] See Section 3.
   (b) Did you describe the limitations of your work? [Yes] See Section 6.
   (c) Did you discuss any potential negative societal impacts of your work? [Yes] See Section 6.
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes]
   (b) Did you include complete proofs of all theoretical results? [Yes] In the supplemental material.
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [No] But it is planned.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] In the supplemental material.
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] See Table 5.2.
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [No]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes] See footnote 2.
   (b) Did you mention the license of the assets? [Yes] See footnote 2.
   (c) Did you include any new assets either in the supplemental material or as a URL? [No]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [No]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [No]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]