

This CVPR workshop paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

Generative Flows as a General Purpose Solution for Inverse Problems

José A. Chávez San Pablo Catholic University Peru

jose.chavez.alvarez@ucsp.edu.pe

Abstract

Due to the success of generative flows to model data distributions, they have been explored in inverse problems. Given a pre-trained generative flow, previous work proposed to minimize the 2-norm of the latent variables as a regularization term. The intuition behind it was to ensure high likelihood latent variables that produce the closest restoration. However, high-likelihood latent variables may generate unrealistic samples as we show in our experiments. We therefore propose a solver to directly produce high-likelihood reconstructions. We hypothesize that our approach could make generative flows a general purpose solver for inverse problems. Furthermore, we propose 1×1 coupling functions to introduce permutations in a generative flow. It has the advantage that its inverse does not require to be calculated in the generation process. Finally, we evaluate our method for denoising, deblurring, inpainting, and colorization. We observe a compelling improvement of our method over prior works.

1. Introduction

With the success of generative models in the generation of synthetic data [10, 15], prior works [2–4] explored them for solving inverse problems in image processing. In this approach, a generative model can be pre-trained with some data such as medical images, faces, animals, and more [12, 20, 21]. Then, this trained model is used to find the best restoration for an inverse problem.

In an inverse problem we aim to recover $x^* \in \mathbb{R}^n$ from the possible-noisy linear measurements $y = Ax^* + \eta$, where $A \in \mathbb{R}^{m \times n}$ is the matrix measurement and $\eta \in \mathbb{R}^m$ is the noise. For instance, y would be a noisy or blurred image, and x^* would be the restoration of this image. So given a generative model $G : z \to x$, researchers [4] proposed the following objective function

$$\hat{z} = \underset{z \in \mathbb{R}^k}{\arg\min} \|AG(z) - y\|^2 + \gamma \cdot \|z\|^2,$$
 (1)

where $z \in \mathbb{R}^k$ and $x \in \mathbb{R}^n$. And the most typical choice for $\|\cdot\|$ is the 2-norm [3,4]. Note that the network G was trained with similar samples than x^* , thus the intuition is to find \hat{z} such that the synthetic image $\hat{x} = G(\hat{z})$ is close to x^* .

Generative Adversarial Networks (GANs) have shown great success in image synthesis [12, 18]; however they are not invertible. Experiments in image manipulation with GANs [1] have shown that *projected images* [22] are not identical to the original ones. Thus, even if we set the identity matrix as A, and $\eta = 0$ in Equation 1; we will probably never achieve the zero-error reconstruction ($x^* = \hat{x}$). On the other hand, Variational Auto-Encoders (VAEs) perform a stochastic mapping from samples to latent variables. Although VAEs have an encoder, we can infer only approximately the corresponding latent variable for a given sample.

Generative flows [8] have a particular advantage among generative models; they are invertible networks. Hence, these models can perform the re-generation of any imagethere exists a latent variable for any image x^* in an inverse problem, and is unique. A sequence of bijective functions allows the exact reconstruction of a sample with its corresponding latent variable. Therefore, zero-error reconstruction is possible and easy to achieve. Authors in [2] proposed to solve a unique inverse problem with generative flows. In this approach, the aim is to encode the corresponding measurements y from a given image x^* . Then, once the model is trained with synthetic pairs of $\{x^*, y\}$, it may estimate the corresponding original image from some measurements y. Note that this method recovers the original image from its measurements once it is trained on a specific inverse problem. One must train different pairs of $\{x^*, y\}$ for each task.

In contrast, GlowIP [3] explored pre-trained generative flows for compressed sensing, denoising and inpainting. In GlowIP, the Glow architecture [14] is trained on the CelebA dataset [16], thus encoding a latent variable for each image as $G^{-1}: x \to z$. Once the model is trained, GlowIP performs Equation 1 to solve an inverse problem. This model does not require the knowledge of any specific inverse problem in the training process—it is only trained with images



Figure 1. Samples from our generative flow trained on the CelebA dataset. We generate images from an isotropic Gaussian distribution $\mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$. Each row represents a batch of 16 samples by setting a specific value for σ . Note that $\sigma = 0$ represents $z = \mathbf{0}$.

from CelebA.

Note that Equation 1 has a regularization term $||z||^2$. As authors suggest in [3,4], the intuition behind it is to encourage high likelihood latent variables, because G is typically trained with the latent distribution $\mathcal{N}(\mathbf{0}, \mathbf{I})$. However, a high log-likelihood in the latent representation may generate a low likelihood sample x. Indeed, the sample $x_0 = G(\mathbf{0})$ is a blurry and unlikely image in a generative flow (see Figure 1). Our hypothesis is that z must produce realistic samples instead. Recently, researchers [19] have attempted to generalize Equation 1, they proposed a maximum a posteriori (MAP) objective for image denoising as follows

$$\hat{z} = \underset{z \in \mathbb{R}^k}{\operatorname{arg\,min}} \frac{1}{2\sigma^2} \|G(z) - y\|^2 - \beta \cdot \log p_G(G(z)), \quad (2)$$

where $\sqrt{\mathbb{E}[\|\eta\|^2]} = \sigma$, and $p_G(G(z))$ is the density estimation by G. However, one must be careful when using density estimates from a generative flow. In the training process, G only sees samples from its corresponding dataset. Nevertheless, one may obtain the density estimates from out-of-distribution samples. In fact, a generative flow may assign a high likelihood to constant or random images [17]. Additionally, estimate $p_G(G(z))$ involves using the generative flow in the form of G and then G^{-1} . As the authors suggest in [19], the hyperparameter β should be configured to control this problem.

On the other hand, permutations represent a key process in generative flows to handle the expressiveness of bijective functions. The Glow architecture [14] requires invertible 1×1 convolutions to permute the entries of hidden feature maps. Note that this architecture requires computing the inverse of the invertible 1×1 convolutions in order to solve Equation 1 or Equation 2. Invertible 1×1 convolutions have shown impressive performance in the generation of synthetic high-resolution images. Nonetheless, the inverse of this layer has a significant cost in limited resources, which could be avoided. In this work, we aim to explore pre-trained generative flows as a general purpose solution for inverse problems. The contributions are summarized in the following points:

- A solver to produce realist restorations in inverse problems through a pre-trained generative flow. Our approach demonstrates quantitive and qualitative improvement among the state-of-the-art methods in many inverse problems.
- We introduce 1 × 1 coupling functions to permute the channels of feature maps in a generative flow. This permutation does not require the computation of its inverse during the generation of samples.

Code to implement our solver and reproduce results is publicly available¹.

2. Preliminaries

In this section, we give an overview on the generative flow model as well as existing methods to introduce permutations.

2.1. Generative flows

Generative flows can be trained by the maximum likelihood using the change of variable formula

$$\log(p_X(x)) = \log(p_Z(z)) + \log(|\det \mathbf{J}_F(x)|), \quad (3)$$

where $F: x \to z$ is an invertible network that maps from a sample $x \sim p_X$ to the latent variable $z \sim p_Z$; and $\mathbf{J}_F(x) = \frac{\partial F(x)}{\partial x}$ is the Jacobian of F evaluated at x. The network F is composed of a sequence of bijections as follows

$$x \xrightarrow[F_1]{F_1} h_1 \xrightarrow[F_2^{-1}]{F_2} h_2 \cdots h_{l-1} \xrightarrow[F_l^{-1}]{F_l} z, \qquad (4)$$

where F must be a transformation such that its Jacobian determinant is computationally tractable [8]. The idea is to choose transformations F_i whose Jacobian $\frac{\partial h_i}{\partial h_{i-1}}$ represents a triangular matrix—the determinant of a triangular matrix is the product of its diagonal elements. Training F in Equation 3 with maximum likelihood is also known as *Gaussianization* [6], where the dimensions in z becomes independent and therefore

$$p_Z(z) = \prod_{i=1}^k p_{Z_i}(z_i),$$
(5)

where k is number of entries of z. Note that Equation 5 simplifies the maximization of the right-hand side in Equation 3. Moreover, each bijection F_i is also called *coupling*

layer [8], and it is a reversible triangular transformation where its determinant is easy to compute.

2.2. Permutations in generative flows

A coupling layer leaves half of its input unchanged while the other is modified. This expressiveness of coupling layers caused by the half unchanged can be handled with permutations. In NICE [8], a swap operation exchanges the halves of the coupling layer input, so in two consecutive coupling layers, the entire input is modified. The Glow model [14] implements learned invertible 1×1 convolutions, which performs impressive results in the CelebA HQ dataset [11]. However, as the authors said the cost of computing its determinant is $\mathcal{O}(c^3)$, where c is the number of input channels in the invertible 1×1 convolution. Furthermore, its inverse has a cost of $\mathcal{O}(c^3)$. On the other hand, Glow resizes the feature maps with squeeze layers [9]. Due to a squeeze layer should produce the same number of elements for a given input, it increments the number of channels c that will be permuted by invertible 1×1 convolutions. This process increases the number of parameters of invertible 1×1 convolution, and consequently the cost of computing its Jacobian determinant and its inverse.

The Glow architecture [14] is used in GlowIP [3] for inverse problems. Note the process of solving inverse problems in Equation 1 involves using the inverse of the network. It means that GlowIP should compute the inverse of each invertible 1×1 convolution in the sample generation. Moreover, training the Glow architecture to infer latent variables from images also involves computing the determinant of this permutation layer.

3. Proposed approach

In this section, we first describe our inverse problem solver to generate realist reconstructions, and then our proposed 1×1 coupling functions to introduce permutations.

3.1. Inverse problem solver

Once we fix our generative flow F_{ϕ} (parametrized by ϕ) with the CelebA dataset [16], we proceed to solve inverse problems. Let assume that we aim to recover $x^* \in \mathbb{R}^n$ from possible-noisy linear measurements $y = Ax^* + \eta$, where $A \in \mathbb{R}^{m \times n}$ and $\eta \in R^m$. Then, given a pre-trained generative flow $F_{\phi} : x \to z$, where $x, z \in \mathbb{R}^n$, we may rewrite Equation 3 as

$$\log(p_{\phi}(x)) = \log(p_Z(z)) - \log(|\det \mathbf{J}_{F_{-}^{-1}}(z)|), \quad (6)$$

where $\log(p_{\phi}(x))$ is the log-likelikehood. Note that we replace the term $|\det \mathbf{J}_{F_{\phi}}(x)|$ by $|\det \mathbf{J}_{F_{\phi}^{-1}}(z)|^{-1}$ in Equation 3. Now, $\mathbf{J}_{F_{\phi}^{-1}}(z)$ is the Jacobian of F_{ϕ}^{-1} evaluated at z.

 $^{^{\}rm l} The \ code \ is \ available \ at \ https://github.com/jarchv/ solverip.git$

Note that the maximization of $\log(p_Z(z))$ is equivalent to the minimization of $||z||^2$, which is proposed in [3, 4]. However, it does not ensure the maximization of the loglikelihood $\log(p_{\phi}(x))$ —we must maximize the Jacobian term as well. The maximization of the right-hand side in Equation 6 directly ensures the generation of high loglikelihood samples x, which is what we pretend. Then, our objective function to find the latent variable that recovers x^* from y in an inverse problem is

$$\hat{z} = \operatorname*{arg\,min}_{z \in \mathbb{R}^k} \left\{ \|AG(z) - y\|_1 + \alpha \cdot \mathcal{L}(z) \right\}, \quad (7)$$

$$\mathcal{L}(z) = -\log(p_Z(z)) + \log(|\det \mathbf{J}_{F_{\phi}^{-1}}(z)|), \quad (8)$$

where $\alpha > 0$ is a hyperparameter. We find empirically that our regularization term $\mathcal{L}(z)$ performs better with the 1-norm in Equation 7.

3.2. 1×1 Coupling functions

Figure 2 shows our generative flow architecture. A *flow step* represents a sequence of an activation normalization [14] and a coupling layer. Inside a coupling layer, we implement swap operations. It introduce permutations as follow

$$h_1, h_2 = \mathsf{swap}(\mathsf{split}(h)), \tag{9}$$

where split separates the input tensor $h \in \mathbb{R}^n$ into two chunks $h_1 \in \mathbb{R}^{n/2}$ and $h_2 \in \mathbb{R}^{n/2}$. Note that a swap operation has constant time complexity $\mathcal{O}(1)$. Then, we set a convolutional neural network CNN as the *coupling function*, which process the first chunk as

$$s, t = \operatorname{CNN}(h_1), \tag{10}$$

where s and t have the dimension as h_1 . Inspired on Glow [14], we perform the following affine transformation on h_2

$$h_2' = h_2 \odot \operatorname{sigmoid}(s+2) + t, \tag{11}$$

where sigmoid is the sigmoid function. Finally, we apply a concat operation to obtain the output tensor h', which has the same dimension as the input h

$$h' = \operatorname{concat}(\operatorname{swap}(h_1, h_2')). \tag{12}$$

Note that if we apply a swap operation in Equation 9, then we should apply another one in Equation 12 to hold the order of the entries in h.



Figure 2. Our generative flow architecture. Each *flow step* is composed by an activation normalization [14] followed by our proposed coupling layer. These flow steps are combined with squeeze layers. We also split low-scale feature maps in two chunks, and force one of the halves (h_z) to directly follow p_Z . The other half continues its path until produce the 1×1 latent variable z.

An additional permutation is performed in our generative flow architecture. Our coupling function applies 1×1 convolutions on 1×1 feature maps h. Hence, each 1×1 coupling function permutes all the elements of 1×1 features maps. However, each of these coupling functions does not require the computation of its inverse during sample generation. Finally, in the process of training our generative flow, the Jacobian of 1×1 coupling function is not required to be computed.

4. Experiments

This section illustrates the experimental details and presents the comparison of our proposed solver to other state-of-the-art methods. The performance evaluation is split into four parts: denoising, deblurring, inpainting, and colorization. Then, we show how realistic are samples with different standard deviations in the latent distribution. Finally, we conduct an ablation study to compare the time speed and memory usage when we use 1×1 coupling functions and invertible 1×1 convolution.

5. Training details and parameters settings

First, we resize to 32×32 resolution all images of the CelebA dataset [16]. Then, we fix a generative flow with the training samples. The Adam optimizer [13] is used during the training with parameters $\alpha = 0.0001$, $\beta_1 = 0.5$ and $\beta_2 = 0.999$. Inspired on [9], we force half of the low-scale feature maps to follow the latent distribution as we can see in Figure 2. We perform *squeeze layers* [9] to reduce and increase the scale of feature maps. Additionally, we set L = 5 as the number of scales, and K = 2 as the number of flow steps per scale. The value of K is duplicated for scales 2×2 and 1×1 —we use 14 flow steps in total. We set the batch size as 32 and stop the training process after 100 epochs. Finally, we use 512 hidden channels for each network CNN.

Once the model is trained, we optimize Equation 7 and



Figure 3. Image denoising on images from the CelebA dataset. The first row is the target image x^* and the second row is the noisy image y, we aim to recover x^* from y. From the 3rd-7th rows, we show results for CSGM, GlowIP, MAP, BM3D and our proposed solver. We synthetically create noisy images with a noise level of $\sqrt{\mathbb{E}[|\eta||^2]} = 0.1$.



Figure 4. Image deblurring on images from the CelebA dataset. The first row is the target image x^* and the second row is the blurry image y, we aim to recover x^* from y. From the 3rd-5th rows, we show results for CSGM, GlowIP, and our proposed solver. We create synthetic blurry images by applying a 3×3 average pooling with stride 1 on targets images.

initialize the latent variable at $z_0 \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$ for $\sigma = 0.1$. In this process, we set $\alpha = 0.005$, $\beta_1 = 0.9$ and $\beta_2 = 0.999$. We configure the batch size as 32, and after 1500 iterations we return the restored image for the respective inverse problem. All experiments are made in an NVIDIA GeForce GTX 750 Ti.

To compare with other methods such as CSGM [4] and GlowIP [3], we use their initialization and objective func-



Figure 5. Image inpainting on images from the CelebA dataset. The first row is the target image x^* and the second row is the masked image y, we aim to recover x^* from y. From the 3rd-5th rows, we show results for CSGM, GlowIP, and our proposed solver. We crop a region of images to create the masked images.

tion in our pre-trained generative flow. Note that CSGM and GlowIP perform the same objective function (Equation 1), the unique difference is the initialization: random initialization for CSGM and zero initialization for GlowIP. We therefore test many values for their penalization hyperparameter γ . Each value of this hyperparameter in our experiments is the one that achieves the best performance for each inverse problem. Additionally, in image denoising, we also compare with BM3D [7] and MAP [19]. We use the PSNR and SSIM metrics to measure the quality of restorations on each inverse problem.

5.1. Denoising

For image denoising we use a noise level of $\sqrt{\mathbb{E}[||\eta||^2]} = 0.1$. Figure 3 shows that CSGM ($\gamma = 0.1$) smooths the noisy image but can not remove all distortions. These distortions are slightly more visible in GlowIP ($\gamma = 0.1$) than CSGM. On the other hand, MAP [19] has similar visual results to GlowIP. In MAP, we use the same initialization as ours, and 0.0015 as the learning rate with $\beta = 0.5$ —we find empirically the best performance for this method using these settings. Our method outperforms BM3D in the PSNR and SSIM metrics as you can see in Table 1. Indeed, our restorations look sharper and more realistic than those from BM3D (see Figure 3). Nonetheless, our method eliminates some details from the background in samples. In this case, we set $\alpha = 0.05$ in Equation 7.

5.2. Deblurring

We create synthetic blurry images by applying an average pooling in the target image x^* , and use it as the measurements y. There is no significant difference between CSGM ($\gamma = 0.01$) and GlowIP ($\gamma = 0$) in image deblurring (see Figure 4). Both produce visual artifacts that are exacerbated if we increase their penalization term—CSGM and GlowIP mitigate visual artifacts with $\gamma = 0.01$ and $\gamma = 0$ respectively. As you can see in Table 1, our method outperforms GlowIP and CSGM in the PSNR and SSIM metrics. Moreover, our results look more sharp and close to the target image. Our hyperparameter α in this case is 0.02.

5.3. Inpainting

For inpainting, we mask a square in the center of the target image and fill it with zeros. We obtain the best performance in CSGM for $\gamma = 0.01$. But even so, CSGM produces many visual artifacts. In contrast, GlowIP ($\gamma = 0$) produces a realistic reconstruction of the black square with few visual artifacts. As we can see in Table 1, our method slightly outperforms GlowIP, however our visual results look similar to those from GlowIP (see Figure 5). Apparently, the log-likelihood term does not significantly improve the performance. We obtain our best performance for $\alpha = 0.002$.

5.4. Colorization

In the case of the colorization problem, we average the three channels of the target image to produce y. CSGM ($\gamma = 0.01$) generates over-saturated images, which have different tones than target images as we can see in Figure 6. Instead, GlowIP ($\gamma = 0$) produces washed-out images, and similar to image deblurring, its regularization term does not improve its performance. Our method performs the colorization problem, however our results are similar to those

Table 1. Average PSNR(dB) and SSIM for CSGM, GlowIP, MAP, BM3D and our method. We evaluate on inverse problems such as denoising, deblurring, inpainting and colorization. Each numerical measurement is the average over 192 samples.

Method	Denoising		Deblurring		Inpainting		Colorization	
	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
CSGM [4]	37.41	0.897	37.17	0.883	36.39	0.839	34.71	0.897
GlowIP [3]	37.21	0.894	35.82	0.867	37.81	0.903	35.79	0.937
MAP [19]	37.35	0.898	_	_	_	_	_	_
BM3D [7]	37.65	0.905	_	_	_	_	_	_
Ours	37.72	0.918	39.72	0.963	37.97	0.913	35.72	0.940



Figure 6. Image colorization on images from the CelebA dataset. The first row is the target image x^* and the second row is the grayscale image y, we aim to recover x^* from y. From the 3rd-5th rows, we show results for CSGM, GlowIP, and our proposed solver. We create grayscale images by averaging the channels of the target images.

from GlowIP. In this case, our solver does not achieve a significant improvement over state-of-the-art methods. Table 1 shows our results for $\alpha = 0.02$.

5.5. Additional experiments

Due to our generative flow is trained with a latent distribution $\mathcal{N}(\mathbf{0}, \mathbf{I})$, we must use this distribution to produce synthetic samples. Hence, we experiment with different standard deviations of the latent distribution. Given a latent variables $z \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$, we sample for many values of σ as you can see in Figure 1. Our experiments show that if we set $z_0 = \mathbf{0}$, the sample $F^{-1}(z_0) = x_0$ is a blurry and unrealistic image. As we increase the standard deviation in the latent variables, we generate high likelihood samples, however we also introduce visual artifacts. Our experiments show that latent variables with a σ in [0.8, 1.2] produce the most realistic samples.

As you can see in Figure 1, high likelihood latent variables do not ensure realistic samples. Both CSGM and GlowIP have problems in image denoising and image deblurring. In the case of inpainting, all methods (including ours) can not perform the reconstruction of non-frontal images. Maybe, a deeper generative flow could generate more realistic restorations—we perform a shallow generative flow for our experiments. An analysis of continuous normalizing flows [5] as a pre-trained model is beyond the scope of this work.

5.6. Ablation study

We conduct an ablation study to verify the computational cost on the generation process. As we mentioned before, solving inverse problems implicates using the inverse of our network. Thus, we compare two configurations of our architecture. The first one introduces permutations with invertible 1×1 convolutions; we follow the configuration of Glow [14] with three different scales (L = 3). In this configuration, each flow step is a sequence of activation normalization, invertible 1×1 convolution and coupling layer. The second architecture employs our 1×1 coupling functions for permutations. Both configurations set 14 as the

Table 2. We compare two architectures with respect to the number of parameters and time speed measured using milliseconds. We configure the first one with invertible 1×1 convolutions, and the second one with our proposed 1×1 coupling functions. The numerical measurements for time speed show the mean and standard deviation over 1500 iterations and 128 samples.

Method	Time (ms)	# of parameters
Inv. 1 × 1 Conv. [14]	10.86 ± 0.18	33.8M
Ours	4.46 ± 0.03	32.7M

total number of flow steps.

Table 2 proves that our method is faster than invertible 1×1 convolutions in the generation process. It is worth mentioning that, on average, our generation time is about 6.4ms less than the other configuration. Note that both architectures have a similar amount of parameters. Nevertheless, invertible 1×1 convolutions require the computation of its inverse in Equation 7. This process increases the time needed to generate a sample.

6. Discussion

We demonstrate that our solver improves the performance at denoising, deblurring and inpainting in lowresolution images from the CelebA dataset [16]. Indeed, our reconstructions are the sharpest and closest to the original image. Furthermore, we introduce 1×1 coupling functions to save computational time in the generation process of generative flows. Our method ensures the generation of high likelihood samples, which results in sharp and realistic restorations in inverse problems. Our experiments suggest that generative flows can be a general purpose solution for inverse problems. Note that we use the same network for each inverse problem, we do not need to re-train a new network for each task. A future direction is to evaluate our method on high-resolution images and real-world examples.

References

- Rameen Abdal, Peihao Zhu, Niloy J. Mitra, and Peter Wonka. Styleflow: Attribute-conditioned exploration of stylegan-generated images using conditional continuous normalizing flows. ACM Trans. Graph., 40(3), May 2021.
- [2] Lynton Ardizzone, Jakob Kruse, Carsten Rother, and Ullrich Köthe. Analyzing inverse problems with invertible neural networks. In *International Conference on Learning Representations*, 2019. 1
- [3] Muhammad Asim, Max Daniels, Oscar Leong, Ali Ahmed, and Paul Hand. Invertible generative models for inverse problems: mitigating representation error and dataset bias. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 399–409. PMLR, 13–18 Jul 2020. 1, 2, 3, 4, 5, 7

- [4] Ashish Bora, Ajil Jalal, Eric Price, and Alexandros G Dimakis. Compressed sensing using generative models. In *International Conference on Machine Learning*, pages 537– 546. PMLR, 2017. 1, 2, 4, 5, 7
- [5] Ricky T. Q. Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018. 7
- [6] Scott Chen and Ramesh Gopinath. Gaussianization. In T. Leen, T. Dietterich, and V. Tresp, editors, *Advances in Neural Information Processing Systems*, volume 13. MIT Press, 2000. 3
- [7] Kostadin Dabov, Alessandro Foi, Vladimir Katkovnik, and Karen Egiazarian. Image denoising by sparse 3-d transformdomain collaborative filtering. *IEEE Transactions on image processing*, 16(8):2080–2095, 2007. 6, 7
- [8] Laurent Dinh, David Krueger, and Yoshua Bengio. Nice: Non-linear independent components estimation. arXiv preprint arXiv:1410.8516, 2014. 1, 3
- [9] Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using real nvp. In *International Confer*ence on Learning Representations 2017 (Conference Track), 2017. 3, 4
- [10] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Q. Weinberger, editors, Advances in Neural Information Processing Systems, volume 27. Curran Associates, Inc., 2014. 1
- [11] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. arXiv preprint arXiv:1710.10196, 2017. 3
- [12] Tero Karras, Miika Aittala, Samuli Laine, Erik Härkönen, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Alias-free generative adversarial networks. *Advances in Neural Information Processing Systems*, 34, 2021. 1
- [13] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun, editors, 3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings, 2015. 4
- [14] Durk P Kingma and Prafulla Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018. 1, 2, 3, 4, 7, 8
- [15] Diederik P. Kingma and Max Welling. Auto-Encoding Variational Bayes. In 2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings, 2014. 1
- [16] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of the IEEE international conference on computer vision*, pages 3730–3738, 2015. 1, 3, 4, 8

- [17] Eric Nalisnick, Akihiro Matsukawa, Yee Whye Teh, Dilan Gorur, and Balaji Lakshminarayanan. Do deep generative models know what they don't know? arXiv preprint arXiv:1810.09136, 2018. 2
- [18] Axel Sauer, Katja Schwarz, and Andreas Geiger. Styleganxl: Scaling stylegan to large diverse datasets. arXiv preprint arXiv:2202.00273, 2022. 1
- [19] Jay Whang, Qi Lei, and Alex Dimakis. Solving inverse problems with a flow-based noise model. In *International Conference on Machine Learning*, pages 11146–11157. PMLR, 2021. 2, 6, 7
- [20] Yan Wu, Jeff Donahue, David Balduzzi, Karen Simonyan, and Timothy Lillicrap. Logan: Latent optimisation for generative adversarial networks. arXiv preprint arXiv:1912.00953, 2019. 1
- [21] Xin Yi, Ekta Walia, and Paul Babyn. Generative adversarial network in medical imaging: A review. *Medical Image Analysis*, 58:101552, 2019. 1
- [22] Jun-Yan Zhu, Philipp Krähenbühl, Eli Shechtman, and Alexei A. Efros. Generative visual manipulation on the natural image manifold. In Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling, editors, *Computer Vision – ECCV 2016*, pages 597–613, Cham, 2016. Springer International Publishing. 1