

VFlow: More Expressive Generative Flows with Variational Data Augmentation

Jianfei Chen^{1,2} Cheng Lu¹ Biqi Chenli^{1,2} Jun Zhu^{1,2} Tian Tian^{1,2}

Abstract

Generative flows are promising tractable models for density modeling that define probabilistic distributions with invertible transformations. However, tractability imposes architectural constraints on generative flows. In this work, we study a previously overlooked constraint that all the intermediate representations must have the same dimensionality with the data due to invertibility, limiting the width of the network. We propose VFlow to tackle this constraint on dimensionality. VFlow augments the data with extra dimensions and defines a maximum evidence lower bound (ELBO) objective for estimating the distribution of augmented data jointly with the variational data augmentation distribution. Under mild assumptions, we show that the maximum ELBO solution of VFlow is always better than the original maximum likelihood solution. For image density modeling on the CIFAR-10 dataset, VFlow achieves a new state-of-the-art 2.98 bits per dimension.

1. Introduction

Generative flows (Dinh et al., 2014; 2017; Kingma & Dhariwal, 2018; Ho et al., 2019) are a promising class of generative models. They define a probability distribution $p(\mathbf{x})$ by applying an invertible transformation $\mathbf{x} = \mathbf{f}^{-1}(\epsilon)$ to some simple and known distribution $p(\epsilon)$. Stacking a sequence $\mathbf{f}_1 \dots, \mathbf{f}_L$ of deep neural networks as the transformation, generative flows can model complicated high-dimensional data. Comparing with generative adversarial networks (GANs) (Goodfellow et al., 2014) and variational autoencoders (VAEs) (Kingma & Welling, 2014), generative flows are particularly attractive because their sampling process and density estimation are tractable. Due to these

¹Department of Computer Science and Technology, Institute for AI, BNRist Center, Tsinghua University ²RealAI. Correspondence to: Jun Zhu <dczsj@mail.tsinghua.edu.cn>.

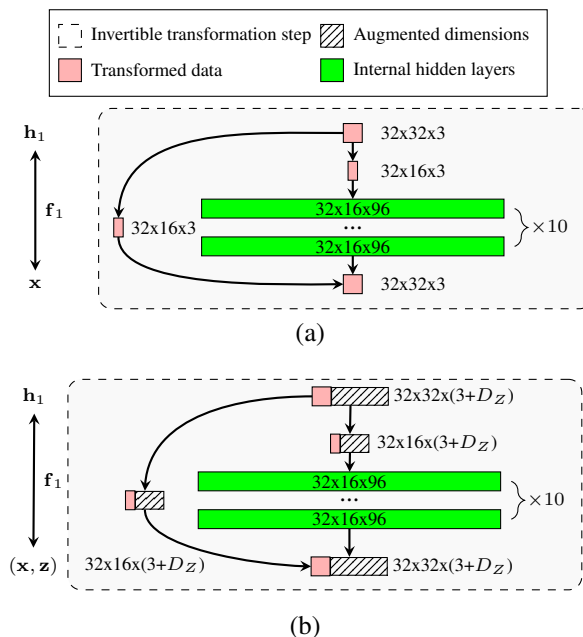


Figure 1. (a) Bottleneck problem in a Flow++ (Ho et al., 2019) for CIFAR-10. Dimensionality of the transformed data (red) limits the model capacity. (b) Our solution VFlow, where D_Z is the dimensionality of the augmented random variable. Only the transformation step f_1 is shown due to space constraint.

advantages, generative flows have been applied to a wide range of problems including image generation (Kingma & Dhariwal, 2018), speech synthesis (Prenger et al., 2019), 3D point cloud generation (Yang et al., 2019), semi-supervised learning (Nalisnick et al., 2019), anomaly detection (Choi et al., 2018), and ray tracing (Müller et al., 2019).

However, tractability comes with a cost of model expressiveness. To be tractable, generative flows have more architectural constraints compared with other non-invertible models, such as GANs and VAEs. One particular constraint is that the determinant of the Jacobian of \mathbf{f} must be efficient to compute. While previous work typically adopts transformations with diagonal (Dinh et al., 2014; 2017; Kingma & Dhariwal, 2018) or triangular Jacobian (Papamakarios et al., 2017), there has been lots of recent work developing transformations with free-form Jacobians, including invertible 1x1

convolution (Kingma & Dhariwal, 2018), continuous time flows (Chen et al., 2018; Grathwohl et al., 2019), invertible residual blocks (Behrmann et al., 2019; Chen et al., 2019), and emerging convolutions (Hoogeboom et al., 2019).

In this paper, we study another orthogonal architectural constraint, *the bottleneck problem*. To be invertible, all the transformation steps $\mathbf{f}_1, \dots, \mathbf{f}_L$ must output the same dimensionality with the input data \mathbf{x} , although each transformation (*i.e.*, neural network) can have internal hidden layers of higher dimensionality. This contradicts with the commonly adopted wisdom of deep learning to learn over-complete features, *i.e.*, higher dimensional features than the data. As an example, Fig. 1(a) presents a state-of-the-art Flow++ (Ho et al., 2019) architecture. Although each transformation \mathbf{f}_l has internal higher-dimensional hidden layers (green), its input and output (red) still lie on the lower-dimensional data space. This makes the generative flow highly inefficient because the high-dimensional features extracted within a transformation step cannot be reused by subsequent steps.

We propose VFlow as a solution to the bottleneck problem. VFlow augments the data \mathbf{x} by extra dimensions \mathbf{z} , which are interpreted as latent variables. We develop a variational inference framework to learn a generative flow $p(\mathbf{x}, \mathbf{z})$ in the augmented data space jointly with the augmented data distribution $q(\mathbf{z}|\mathbf{x})$. We show that VFlow is a generalization of the vanilla generative flows, so the augmented dimensions always help. VFlow improves existing generative flows, and achieves a state-of-the-art 2.98 bits per dimension likelihood on the CIFAR-10 dataset.

On the efficiency side, the additional $q(\mathbf{z}|\mathbf{x})$ network and higher data dimensionality of VFlow only add marginal overhead to the vanilla generative flows. Meanwhile, VFlow can be more compact, since more information can be shared between individual transformation steps. Thus, each transformation step can be simpler by avoiding extracting high-dimensional features from scratch. We show that VFlow can be 2.6 times more compact than vanilla generative flows, while achieving similar model quality. Our code is open-sourced at <https://github.com/thu-ml/vflow>.

2. Backgrounds

In this section, we review the basics of generative flows and formally define the bottleneck problem.

2.1. Generative Flows

Given a distribution of D_X -dimensional data \mathbf{x} on the space \mathbb{R}^{D_X} , the task of generative modeling aims to learn a model distribution $p(\mathbf{x}; \boldsymbol{\theta})$ parameterized by $\boldsymbol{\theta}$ that approximates the data distribution. The model can be learned with the

maximum likelihood principle

$$\max_{\boldsymbol{\theta}} \mathbb{E}_{\hat{p}(\mathbf{x})} [\log p(\mathbf{x}; \boldsymbol{\theta})], \quad (1)$$

where $\hat{p}(\mathbf{x})$ is the empirical data distribution.

Generative flows define a sequence of invertible transformation steps $\mathbf{f}_1, \dots, \mathbf{f}_L$, that transform a datum \mathbf{x} to some random variable ϵ ,

$$\mathbf{x} \xleftarrow{\mathbf{f}_1} \mathbf{h}_1 \xleftarrow{\mathbf{f}_2} \mathbf{h}_2 \cdots \xleftarrow{\mathbf{f}_L} \epsilon,$$

where ϵ follows a simple factorized distribution that $p_\epsilon(\epsilon) = \prod_i p_\epsilon(\epsilon_i)$, such as the standard normal distribution. For notational simplicity, we define $\mathbf{h}_0 = \mathbf{x}$ and $\mathbf{h}_L = \epsilon$. Let \mathbf{f} be the composition of all the L transformations, such that $\epsilon = \mathbf{f}(\mathbf{x}; \boldsymbol{\theta})$, a generative flow defines the model distribution with the change-of-variables formula

$$\log p(\mathbf{x}; \boldsymbol{\theta}) = \log p_\epsilon(\epsilon) + \log \left| \frac{\partial \epsilon}{\partial \mathbf{x}} \right|,$$

where $\log \left| \frac{\partial \epsilon}{\partial \mathbf{x}} \right|$ is the log-absolute-determinant of the Jacobian of \mathbf{f} . Samples from $p(\mathbf{x}; \boldsymbol{\theta})$ can be obtained by taking the inverse transformation from p_ϵ :

$$\epsilon \sim p_\epsilon, \quad \mathbf{x} = \mathbf{f}^{-1}(\epsilon; \boldsymbol{\theta}).$$

One popular invertible transformation is the affine coupling layer (Dinh et al., 2017), where each transformation $\mathbf{h}_l = \mathbf{f}_l(\mathbf{h}_{l-1}; \boldsymbol{\theta})$ is defined as

$$\begin{aligned} \mathbf{x}_1, \mathbf{x}_2 &= \text{split}(\mathbf{h}_{l-1}), \\ \mathbf{y}_1 &= \mathbf{x}_1, \quad \mathbf{y}_2 = \boldsymbol{\mu}(\mathbf{x}_1; \boldsymbol{\theta}) + \exp(\mathbf{s}(\mathbf{x}_1; \boldsymbol{\theta})) \circ \mathbf{x}_2, \\ \mathbf{h}_l &= \mathbf{f}_l(\mathbf{h}_{l-1}; \boldsymbol{\theta}) = \text{concat}(\mathbf{y}_1, \mathbf{y}_2), \end{aligned} \quad (2)$$

where $\text{split}(\cdot)$ is any operation that splits the input into two disjoint parts, $\text{concat}(\cdot)$ is its inverse operation, and $\boldsymbol{\mu}, \mathbf{s}$ are neural networks with B hidden layers and D_H units per layer.

2.2. The Bottleneck Problem

Starting from the work on universal approximation theorems (Gybenko, 1989; Mhaskar, 1993) of multi-layer perceptrons, it is well known that network width plays an important role on the model capacity. The impact of network width is also verified empirically by recent works such as Wide ResNet (Zagoruyko & Komodakis, 2016) and EfficientNet (Tan & Le, 2019). Almost all existing non-invertible deep models, such as residual networks (He et al., 2016) and generative adversarial networks (Goodfellow et al., 2014) have features in a higher-dimensional space than the original data space.

However, for generative flows, all the transformed data $\mathbf{h}_0, \dots, \mathbf{h}_L$ must have the same dimensionality D_X with the input \mathbf{x} due to invertibility, as illustrated in Fig. 1. This architecture is ineffective for three reasons:

1. Few features (green in Fig. 1) extracted within each transformation step can pass through the bottleneck (red in Fig. 1), so subsequent transformation steps must extract their own features from scratch;
2. For fixed dimensional \mathbf{h}_l , the benefit of increasing the hidden layer size D_H is limited. Unlike non-invertible deep networks, which can approximate arbitrary functions with large D_H , the capacity of a single transformation step is intrinsically limited by architectural constraints, even with infinite D_H . For example, an affine coupling layer (Dinh et al., 2017) can not alter all the dimensions at once, while an invertible residual block (Behrmann et al., 2019) has a bounded Lipschitz constant;
3. Due to the limited capacity of a single transformation step, a sufficiently powerful generative flow needs to have *many* transformation steps, which is expensive.

We refer to this issue as the *bottleneck problem*. To reflect the impact of the bottleneck width on model capacity, we denote a generative flow with D -dimensional bottleneck as a *D-dimensional flow*. Ideally, a D_H -dimensional flow completely eliminates the bottleneck.

3. VFlow

We present VFlow, a variational data augmentation framework and compare it with the vanilla generative flows.

3.1. Variational Data Augmentation

The bottleneck problem can be tackled by increasing the dimensionality of the original data, so that the dimensionality of the flow is also increased. To achieve this, we augment the data \mathbf{x} with an additional D_Z -dimensional random variable $\mathbf{z} \in \mathbb{R}^{D_Z}$, and model the augmented data distribution $p(\mathbf{x}, \mathbf{z}; \boldsymbol{\theta})$ with a $(D_X + D_Z)$ -dimensional flow. The new flow $p(\mathbf{x}, \mathbf{z}; \boldsymbol{\theta})$ is more powerful since its dimensionality can be adjusted freely by setting D_Z . The underlying invertible transformation becomes $\boldsymbol{\epsilon} = \mathbf{f}(\mathbf{x}, \mathbf{z}; \boldsymbol{\theta})$ where $\boldsymbol{\epsilon} \in \mathbb{R}^{D_X + D_Z}$.

By modeling the augmented data distribution, the log marginal likelihood $\log p(\mathbf{x}; \boldsymbol{\theta}) = \log \int p(\mathbf{x}, \mathbf{z}; \boldsymbol{\theta}) d\mathbf{z}$ and optimization problem (1) become intractable in general. Thus, we resort to the variational methods and establish a lower bound of the marginal likelihood with a variational distribution of the augmented data $q(\mathbf{z}|\mathbf{x}; \boldsymbol{\phi})$:

$$\log p(\mathbf{x}; \boldsymbol{\theta}) \geq \mathbb{E}_{q(\mathbf{z}|\mathbf{x}; \boldsymbol{\phi})} [\log p(\mathbf{x}, \mathbf{z}; \boldsymbol{\theta}) - \log q(\mathbf{z}|\mathbf{x}; \boldsymbol{\phi})], \quad (3)$$

which is known as evidence lower bound (ELBO) in variational inference literature. VFlow optimizes the following maximum ELBO objective as a surrogate of the maximum

likelihood objective Eq. (1):

$$\max_{\boldsymbol{\theta}, \boldsymbol{\phi}} \mathbb{E}_{\hat{p}(\mathbf{x})q(\mathbf{z}|\mathbf{x}; \boldsymbol{\phi})} [\log p(\mathbf{x}, \mathbf{z}; \boldsymbol{\theta}) - \log q(\mathbf{z}|\mathbf{x}; \boldsymbol{\phi})]. \quad (4)$$

After training, density estimation can be achieved with importance sampling

$$\log p(\mathbf{x}; \boldsymbol{\theta}) \approx \log \left(\frac{1}{S} \sum_{i=1}^S \frac{p(\mathbf{x}, \mathbf{z}_i; \boldsymbol{\theta})}{q(\mathbf{z}_i|\mathbf{x}; \boldsymbol{\phi})} \right), \quad (5)$$

where $\mathbf{z}_1, \dots, \mathbf{z}_S \sim q(\mathbf{z}|\mathbf{x}; \boldsymbol{\phi})$ are the S samples.

The augmented data distribution $q(\mathbf{z}|\mathbf{x}; \boldsymbol{\phi})$ is modeled with another conditional flow defined with an invertible transformation $\mathbf{z} = \mathbf{g}^{-1}(\boldsymbol{\epsilon}_q; \mathbf{x}, \boldsymbol{\phi})$:

$$\log q(\mathbf{z}|\mathbf{x}; \boldsymbol{\phi}) = \log p_{\boldsymbol{\epsilon}}(\boldsymbol{\epsilon}_q) - \log \left| \frac{\partial \mathbf{z}}{\partial \boldsymbol{\epsilon}_q} \right|,$$

where $\boldsymbol{\epsilon}_q$ follows the same distribution $p_{\boldsymbol{\epsilon}}$ with $\boldsymbol{\epsilon}$. Given that $\mathbf{z} = \mathbf{g}^{-1}(\boldsymbol{\epsilon}_q; \mathbf{x}, \boldsymbol{\phi})$ is a differentiable reparameterization of $\boldsymbol{\epsilon}_q$, the ELBO in Eq. (3) can be optimized with the reparameterization trick (Kingma & Welling, 2014). VFlow is illustrated in Fig. 1(b). By choosing different architectures for $p(\mathbf{x}, \mathbf{z}; \boldsymbol{\theta})$ and $q(\mathbf{z}|\mathbf{x}; \boldsymbol{\phi})$, VFlow can be combined with various existing generative flows (Kingma & Dhariwal, 2018; Ho et al., 2019; Chen et al., 2019), and improve their expressiveness and efficiency.

3.2. Connection to Vanilla Generative Flows

While VFlow tackles the bottleneck problem, it only maximizes a lower bound of the likelihood. It is thus worth studying whether the gain from increased dimensionality of the flow surpasses the gap between the marginal likelihood and the ELBO.

We now show that VFlow is indeed better even it only optimizes a lower bound. Before presenting the theoretical results, we need to clarify the parameter space of different flow models.

- A vanilla generative flow defines $p_x(\mathbf{x}; \boldsymbol{\theta}_x)$, where $\boldsymbol{\theta}_x \in \Theta_x$, and Θ_x is the parameter space.
- For any $D_Z > 0$, a VFlow defines $p_a(\mathbf{x}, \mathbf{z}; \boldsymbol{\theta}_a)$, where $\mathbf{z} \in \mathbb{R}^{D_Z}$, $\boldsymbol{\theta}_a \in \Theta_a$, and Θ_a is the parameter space. Marginalizing \mathbf{z} yields $p_a(\mathbf{x}; \boldsymbol{\theta}_a)$.
- For any $D_Z > 0$, the variational distribution is $q(\mathbf{z}|\mathbf{x}; \boldsymbol{\phi})$, where $\mathbf{z} \in \mathbb{R}^{D_Z}$, $\boldsymbol{\phi} \in \Phi$, and Φ is the parameter space.

With these notations, the maximum likelihood solution of vanilla generative flows (Eq. 1) can be written as $\max_{\boldsymbol{\theta}_x} \mathbb{E}_{\hat{p}(\mathbf{x})} [\log p_x(\mathbf{x}; \boldsymbol{\theta}_x)]$, and the maximum ELBO solution of VFlow can be written as $\max_{\boldsymbol{\theta}_a, \boldsymbol{\phi}} \mathbb{E}_{\hat{p}(\mathbf{x})q(\mathbf{z}|\mathbf{x}; \boldsymbol{\phi})} [\log p_a(\mathbf{x}, \mathbf{z}; \boldsymbol{\theta}_a) - \log q(\mathbf{z}|\mathbf{x}; \boldsymbol{\phi})]$.

Our analysis is based on the following assumptions:

A1 (*high-dimensional flow can emulate low-dimensional flow*) For all $\theta_x \in \Theta_x$ and $D_Z > 0$, there exists $\theta_a \in \Theta_a$, such that for all \mathbf{x} and \mathbf{z} ,

$$p_a(\mathbf{x}, \mathbf{z}; \theta_a) = p_x(\mathbf{x}; \theta_x) p_\epsilon(\mathbf{z}).$$

A2 (*the variational family has an identity transformation*) For all $D_Z > 0$, there exists $\phi \in \Phi$, such that for all \mathbf{x} and \mathbf{z} , $q(\mathbf{z}|\mathbf{x}; \phi) = p_\epsilon(\mathbf{z})$, where $p_\epsilon(\mathbf{z})$ is the simple factorized distribution defined in Sec. 2.1.

Assumptions A1 and A2 can be verified for most existing invertible transformation steps (Dinh et al., 2017; Kingma & Dhariwal, 2018; Chen et al., 2019). Consider the simplest case of a linear flow $\epsilon = \mathbf{x}\theta_x$, where $\theta_x \in \Theta_x$ is an orthonormal matrix. Taking $\theta_a = \begin{bmatrix} \theta_x & \mathbf{0} \\ \mathbf{0} & \mathbf{I} \end{bmatrix}$ yields $p_a(\mathbf{x}, \mathbf{z}; \theta_a) = p_\epsilon \left(\begin{bmatrix} \mathbf{x} & \mathbf{z} \\ \mathbf{0} & \mathbf{I} \end{bmatrix} \begin{bmatrix} \theta_x & \mathbf{0} \\ \mathbf{0} & \mathbf{I} \end{bmatrix} \right) = p_\epsilon(\mathbf{x}\theta_x) p_\epsilon(\mathbf{z})$, satisfying Assumption A1. Moreover, $q(\mathbf{z}|\mathbf{x}; \mathbf{I}) = p_\epsilon(\mathbf{z}) = p_\epsilon(\mathbf{z})$, satisfying Assumption A2. We leave the detailed verification for Glow (Kingma & Dhariwal, 2018) and Residual Flow (Chen et al., 2019) in Appendix A.

The following theorem compares the maximum ELBO solution Eq. (4) of VFlow with the maximum likelihood solution Eq. (1) of vanilla generative flows.

Theorem 1. *Under Assumptions A1 and A2, for any $D_Z > 0$, we have*

$$\begin{aligned} & \max_{\theta_x \in \Theta_x} \mathbb{E}_{\hat{p}(\mathbf{x})} [\log p_x(\mathbf{x}; \theta_x)] \\ & \leq \max_{\theta_a \in \Theta_a, \phi \in \Phi} \mathbb{E}_{\hat{p}(\mathbf{x})} [\log p_a(\mathbf{x}, \mathbf{z}; \theta_a) - \log q(\mathbf{z}|\mathbf{x}; \phi)]. \end{aligned}$$

Proof. Our proof is based on a simple construction. Given any vanilla flow model $p_x(\mathbf{x}; \theta_x)$, according to Assumptions A1 and A2, for any $D_Z > 0$, we can construct

- $\theta(\theta_x) \in \Theta_a$, such that $p_a(\mathbf{x}, \mathbf{z}; \theta(\theta_x))$ is a factorized distribution $p_a(\mathbf{x}, \mathbf{z}; \theta(\theta_x)) = p_x(\mathbf{x}; \theta_x) p_\epsilon(\mathbf{z})$. This is a very weak model that does not utilize \mathbf{z} at all.
- $\phi \in \Phi$, such that the variational distribution is trivial $q(\mathbf{z}|\mathbf{x}; \phi) = p_\epsilon(\mathbf{z})$.

Even using these special models, we have

$$\log p_a(\mathbf{x}, \mathbf{z}; \theta(\theta_x)) - \log q(\mathbf{z}|\mathbf{x}; \phi) = \log p_x(\mathbf{x}; \theta_x).$$

Now, starting from

$$\begin{aligned} & \max_{\theta_x \in \Theta_x} \mathbb{E}_{\hat{p}(\mathbf{x})} [\log p_x(\mathbf{x}; \theta_x)] \\ & = \max_{\theta_a \in \Theta_a, \phi \in \Phi} \mathbb{E}_{\hat{p}(\mathbf{x})} [\log p_x(\mathbf{x}; \theta_x) + \log p_\epsilon(\mathbf{z}) - \log p_\epsilon(\mathbf{z})], \end{aligned}$$

considering the special $\theta(\theta_x)$, we have

$$= \max_{\theta_x \in \Theta_x} \mathbb{E}_{\hat{p}(\mathbf{x})} [\log p_a(\mathbf{x}, \mathbf{z}; \theta(\theta_x)) - \log p_\epsilon(\mathbf{z})],$$

allowing the parameter of p_a to be chosen freely from Θ_a , not just $\theta(\theta_x) \in \Theta_a$, we have

$$\leq \max_{\theta_a \in \Theta_a} \mathbb{E}_{\hat{p}(\mathbf{x})} [\log p_a(\mathbf{x}, \mathbf{z}; \theta_a) - \log p_\epsilon(\mathbf{z})],$$

replacing $p_\epsilon(\mathbf{z})$ by the trivial variational distribution, we have

$$= \max_{\theta_a \in \Theta_a} \mathbb{E}_{\hat{p}(\mathbf{x})} [\log p_a(\mathbf{x}, \mathbf{z}; \theta_a) - \log q(\mathbf{z}|\mathbf{x}; \phi)], \quad (6)$$

allowing ϕ to be chosen freely from Φ , we have

$$\leq \max_{\theta_a \in \Theta_a, \phi \in \Phi} \mathbb{E}_{\hat{p}(\mathbf{x})} [\log p_a(\mathbf{x}, \mathbf{z}; \theta_a) - \log q(\mathbf{z}|\mathbf{x}; \phi)].$$

□

Remark 1: Theorem 1 does not consider optimization issues, such as convergence speed. However, as we shall see in Appendix A, it is rather simple for a VFlow to mimic a vanilla generative flow by setting some parameters to zero, due to the residual structure of transformation steps. Therefore, we hypothesize that VFlow should still be better than vanilla generative flows under the same number of optimizer iterations. This is empirically verified in Fig. 6.

Remark 2: Variational inference-based models such as VAEs rely heavily on the quality of the variational posterior $q(\mathbf{z}|\mathbf{x}; \phi)$ to work well. Unlike VAE, VFlow is better than vanilla generative flows even with a trivial variational distribution $q(\mathbf{z}|\mathbf{x}) = p_\epsilon(\mathbf{z})$. This can be seen from Eq. (6).

Finally, combining Theorem 1 with the variational bound Eq. (3), we have

Corollary 1. *Under Assumptions A1 and A2, for any $D_Z > 0$, we have*

$$\max_{\theta_x \in \Theta_x} \mathbb{E}_{\hat{p}(\mathbf{x})} [\log p_x(\mathbf{x}; \theta_x)] \leq \max_{\theta_a \in \Theta_a} \mathbb{E}_{\hat{p}(\mathbf{x})} [\log p_a(\mathbf{x}; \theta_a)]$$

3.3. Efficiency

While VFlow makes the model more expressive, its overhead is only marginal. To see this, note that the overhead of VFlow arises from two parts: (1) the cost of computing $q(\mathbf{z}|\mathbf{x}; \phi)$ and (2) the increase of the cost for computing $p(\mathbf{x}, \mathbf{z}; \theta)$ due to the increase of the dimensionality. The first cost can be small by using a much smaller network for $q(\mathbf{z}|\mathbf{x}; \phi)$ than $p(\mathbf{x}, \mathbf{z}; \theta)$. As an extreme case, one can eliminate the cost by adopting $q(\mathbf{z}|\mathbf{x}; \phi) = p_\epsilon(\mathbf{z})$, according to Remark 2. The second cost is small because most computation of $p(\mathbf{x}, \mathbf{z}; \theta)$ is spent on the internal hidden layers (green layers in Fig. 1), whose time complexity is only related to the hidden layer size D_H instead of the flow dimensionality $D_X + D_Z$.

On the other hand, VFlow can be more compact and efficient than a vanilla generative flow to achieve similar modeling quality, as it alleviates the ineffectiveness listed in Sec. 2.2 caused by the bottleneck problem.

3.4. Modeling Discrete Data

The discussion so far is limited to continuous data \mathbf{x} . If the data follow a discrete distribution $P(\mathbf{x})$, an additional dequantization step is needed to convert the data from discrete to continuous. Ho et al. (2019) propose to bound the discrete density with a variational dequantization distribution $r(\mathbf{u}|\mathbf{x})$:

$$\log P(\mathbf{x}) \geq \mathbb{E}_{r(\mathbf{u}|\mathbf{x})}[\log p(\mathbf{x} + \mathbf{u}) - \log r(\mathbf{u}|\mathbf{x})],$$

where \mathbf{u} is continuous and $p(\mathbf{x} + \mathbf{u})$ is a generative flow for continuous data. Combining with the ELBO Eq. (3), we obtain a lower bound for discrete data

$$\log P(\mathbf{x}) \geq \mathbb{E}_{r(\mathbf{u}|\mathbf{x}), q(\mathbf{z}|\mathbf{x}+\mathbf{u})}[\log p(\mathbf{x} + \mathbf{u}, \mathbf{z}) - \log r(\mathbf{u}|\mathbf{x}) - \log q(\mathbf{z}|\mathbf{x} + \mathbf{u})]. \quad (7)$$

Estimating the marginal density $\log P(\mathbf{x})$ involves similar importance sampling procedure with Eq. (5), but the samples are drawn from the joint distribution $r(\mathbf{u}|\mathbf{x})q(\mathbf{z}|\mathbf{x} + \mathbf{u})$ of dequantization noise and augmented data.

Although both variational dequantization and VFlow introduce variational distributions, their purposes are different. Variational dequantization aims to reduce the gap between the discrete data distribution and continuous model distribution, while VFlow aims to increase the dimensionality of the flow. These approaches are orthogonal to each other.

4. Related Works

There exists a large bulk of works on developing more flexible transformation steps, such as transformations with free-form Jacobians (Grathwohl et al., 2019; Behrmann et al., 2019; Chen et al., 2019), fast fourier transformation-based invertible convolutions (Hooeboom et al., 2019), flexible coupling functions (Ho et al., 2019; Durkan et al., 2019; Müller et al., 2019), and masked convolutional layers (Hooeboom et al., 2019; Song et al., 2019). VFlow is orthogonal with these approaches since it tackles a different bottleneck of dimensionality, and can be combined with these works to create better models.

The bottleneck problem is studied for *discriminative* invertible models including neural ODEs (Dupont et al., 2019) and i-RevNets (Jacobsen et al., 2018), where zeros are padded to the input data to increase the number of dimensions. In contrast, VFlow studies the much more challenging *generative* modeling problem. For generative modeling, zero padding does not work because the padded data $(\mathbf{x}, \mathbf{0})$ still lies on a D_X -dimensional manifold, while the distribution $p_\epsilon(\epsilon)$ is defined on a $D_X + D_Z$ dimensional space. Therefore, an invertible transformation does not exist. Similarly, simply replicating the data does not help. Another possible solution is reducing the number of transformations L to one. While

this does eliminate the bottleneck problem, the capacity of a single transformation is limited, as discussed in Sec. 2.2.

Variational autoencoders (Kingma & Welling, 2014) can be understood as VFlows where both $p(\mathbf{x}, \mathbf{z}; \theta)$ and $q(\mathbf{z}|\mathbf{x}; \phi)$ are generative flows with a single affine coupling layer. Particularly, a Gaussian VAE $p(\mathbf{x}, \mathbf{z}; \theta) = \mathcal{N}(\mathbf{z}; \mathbf{0}, \mathbf{I})\mathcal{N}(\mathbf{x}; \mu(\mathbf{z}), \exp(\mathbf{s}(\mathbf{z}))^2)$ is equivalent with

$$\begin{aligned} \epsilon_Z &\sim \mathcal{N}(\mathbf{0}, \mathbf{I}), & \epsilon_X &\sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \\ \mathbf{z} &= \epsilon_Z, & \mathbf{x} &= \mu(\epsilon_Z) + \exp(\mathbf{s}(\epsilon_Z)) \circ \epsilon_X, \end{aligned}$$

which shares the same form with the affine coupling layer defined in Eq. (2), despite in the opposite direction. VFlows are more general than VAEs by not assuming the hierarchical structure $p(\mathbf{x}, \mathbf{z}) = p(\mathbf{z})p(\mathbf{x}|\mathbf{z})$. Though it is possible for VAEs to implement both $p(\mathbf{z})$ and $p(\mathbf{x}|\mathbf{z})$ with generative flows (Morrow & Chiu, 2020; Chen et al., 2017), the flow $p(\mathbf{x}|\mathbf{z})$ is still D_X -dimensional, so the bottleneck problem persists. Another line of work implement $q(\mathbf{z}|\mathbf{x})$ with generative flows (Kingma et al., 2016; Rezende & Mohamed, 2015) but leaves $p(\mathbf{x}, \mathbf{z})$ unchanged. VFlow has identical $q(\mathbf{z}|\mathbf{x})$ but more powerful $p(\mathbf{x}, \mathbf{z})$ than these works. There are also a number of works combining VAEs with autoregressive models (Chen et al., 2017; Gulrajani et al., 2017). However they suffer from slow sampling due to the sequential nature of autoregressive models. Finally, while a powerful $q(\mathbf{z}|\mathbf{x})$ is critical for VAEs, it is less important for VFlows, since $p(\mathbf{x}, \mathbf{z})$ itself is powerful even with $q(\mathbf{z}|\mathbf{x}) = p_\epsilon(\mathbf{z})$, as discussed in Sec. 3.2.

Augmented Normalizing Flow (ANF) (Huang et al., 2020) is an independent parallel work of VFlow. Both ANF and VFlow combine generative flows and variational inference. The main difference is the theoretical guarantee. We view VFlow as a general improvement of flow models, so our theory compares the modeling quality of original and augmented data. In contrast, ANF is more closely related to variational autoencoders, and it rather focuses on the universal approximation of probability distributions in the asymptotic case.

5. Toy Data Experiments

We first evaluate VFlow on a toy $D_X = 2$ *Checkerboard* dataset (Behrmann et al., 2019), which is multimodal and its density is shown in Fig. 2(a). The baseline model is Glow (Kingma & Dhariwal, 2018), where each transformation step consists of an affine coupling layer with 2 hidden layers and $D_H = 50$ hidden units per layer. VFlow further augments Glow with a conditional Glow $q(\mathbf{z}|\mathbf{x}; \phi)$ and various number of extra dimensions D_Z . All the models are trained for 100,000 iterations with Adam (Kingma & Ba, 2015) and a batch size of 64, and each experiment is repeated 5 times with different random seeds. Model quality is measured with the log-likelihood $\log p(\mathbf{x})$ on a 1,000-

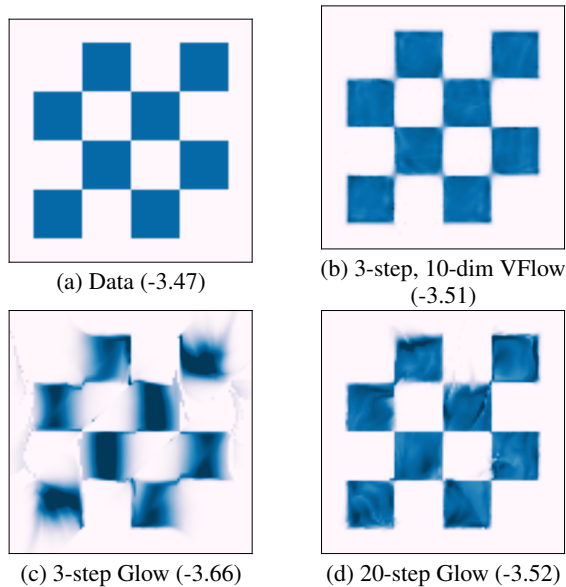


Figure 2. Data and model density on toy data, log-likelihood is shown in parenthesis.

sample test set. For VFlow, likelihood is evaluated with 100-sample importance sampling by Eq. (5).

We study the impact of the dimensionality of the flow $D_X + D_Z \in \{2, 4, 6, 8, 10\}$, where $D_X + D_Z = 2$ is the baseline Glow and $D_X + D_Z > 2$ is VFlow. To control the model size, we vary the total number of transformation steps $L \in \{2, 3, 4, 5, 10, 15, 20\}$. For baseline Glow, the p -network has all the $L_p = L$ transformation steps; and for VFlow, p -network has $L_p = L - 1$ transformation steps and q -network has one transformation step. The result is shown in Fig. 3, VFlow significantly outperforms Glow under similar model size. For example, a 3-step, 10-dimensional VFlow achieves -3.51 ± 0.01 log-likelihood (Fig. 2(b)), outperforming the baseline 3-step Glow with -3.67 ± 0.03 log-likelihood (Fig. 2(c)) by a large margin. The 3-layer, 10-dimensional VFlow even outperforms a much larger 20-step Glow, which achieves -3.54 ± 0.05 log-likelihood (Fig. 2(d)), showing that the model can be much more compact by solving the bottleneck problem.

To further understand why the dimensionality of is important, we visualize the learnt representation for a 2-step Glow and a 2-step VFlow, which has a single transformation step for both p and q . To make visualization possible, z is only one-dimensional, so $D_X + D_Z = 3$. Note that having odd number of dimensions is suboptimal because the affine coupling layer cannot split the data into two parts with equal number of dimensions. Moreover, affine coupling layer cannot represent the one-dimensional distribution $q(z|\mathbf{x})$, so we replace it with a Gaussian layer $\mathcal{N}(z; \boldsymbol{\mu}(\mathbf{x}), \boldsymbol{\sigma}(\mathbf{x}))$ without changing the architecture of $\boldsymbol{\mu}(\mathbf{x})$ and $\boldsymbol{\sigma}(\mathbf{x})$. The

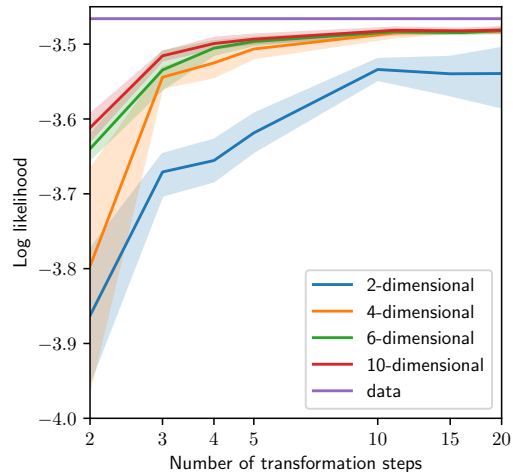


Figure 3. Impact of the dimensionality on the toy dataset.

learnt transformations are visualized in Fig. 4. While Glow struggles to map different modes to the compact space of ϵ , VFlow does a much better job. VFlow learns a pile of “pies” in the ϵ space, where each mode is a pie, and different modes are directly distinguished by the extra dimension z .¹ By shifting the pies to different positions on the x -plane based on z , VFlow can easily handle the multi-modality of the data. Comparing with Glow which needs to map an irregular shape in the ϵ space to a square in the x space, VFlow requires a much simpler transformation, because each mode is already a regular pie in the ϵ space.

6. Density Estimation of Images

In this section, we evaluate VFlow on CIFAR-10² and ImageNet (Russakovsky et al., 2015) for density estimation of images. VFlow augments a state-of-the-art generative flow, Flow++ (Ho et al., 2019) by introducing extra dimensions and another variational distribution $q(\mathbf{z}|\mathbf{x})$. More specifically, the $p(\mathbf{x}, \mathbf{z})$ network is similar with Flow++ shown in Fig. 1, and the main difference is the dimensionality of the flow. Variational dequantization is deployed according to Sec. 3.4. We choose the network architecture for $q(\mathbf{z}|\mathbf{x})$ to be similar with the variational dequantization network $r(\mathbf{u}|\mathbf{x})$ of Flow++. A detailed description of the model architecture is in Appendix B. While we only consider Flow++ for this section due to its impressive density estimation result, our variational data augmentation framework is general and can be combined with future advances of the model architecture.

¹The three axes x_0 , x_1 and z do not distinguish from each other in the ϵ space because the invertible 1×1 convolution in Glow can rotate them arbitrarily.

²<https://www.cs.toronto.edu/~kriz/cifar.html>

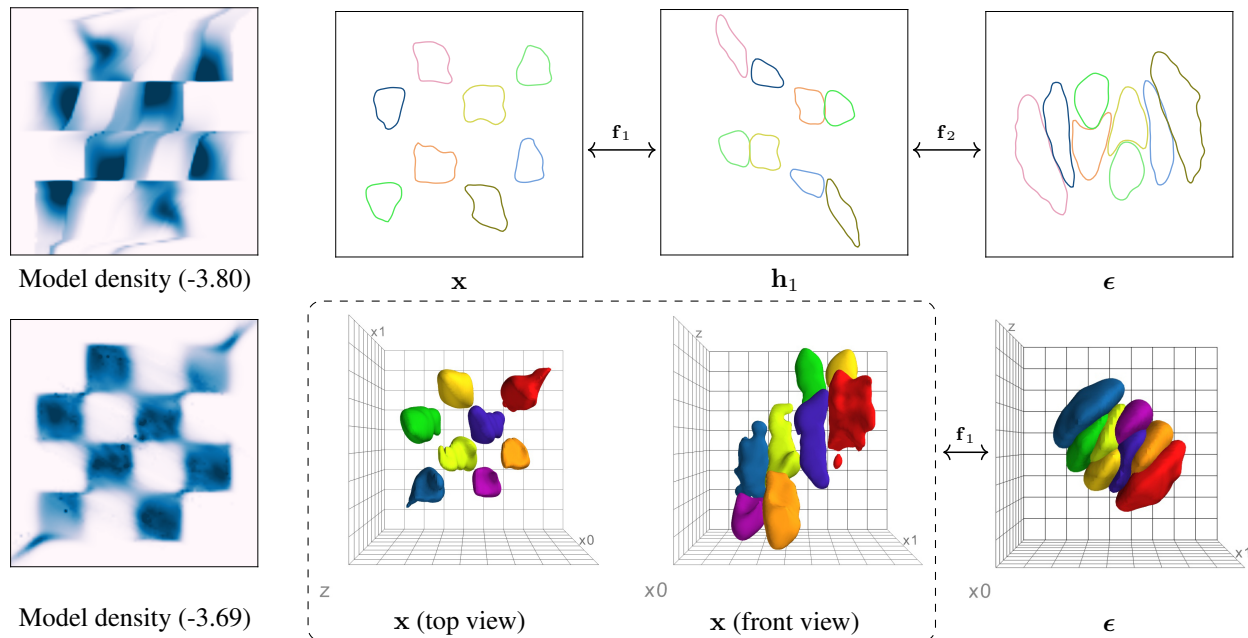


Figure 4. Visualization of learnt transformation on toy data. Top row: 2-step Glow. Bottom row: 2-step, 3-dimensional VFlow. Log-likelihood is shown in parenthesis. We sample ϵ and visualize the transformed density in x , h_1 and ϵ space. The density is estimated from samples by kernel density estimation, and we show the 50% probability contour / isosurface for each mode in different color.

Table 1. Density modeling results in bits/dim (bpd). We report testing bpd for CIFAR-10 and validation bpd for ImageNet.

Model	CIFAR-10	ImageNet 32x32	ImageNet 64x64
Glow (Kingma & Dhariwal, 2018)	3.35	4.09	3.81
FFJORD (Grathwohl et al., 2019)	3.40	-	-
Residual Flow (Chen et al., 2019)	3.28	4.01	3.76
MintNet (Song et al., 2019)	3.32	4.06	-
Flow++ (Ho et al., 2019)	3.08	3.86	3.69
VFlow	2.98	3.83	3.66

The model size is controlled by three main hyper-parameters, (1) the dimensionality of the flow; (2) the hidden layer size D_H ; and (3) the number of hidden layers B per transformation step. For brevity, we refer to a $32 \times 32 \times C$ -dimensional flow as a C -channel flow, where a 3-channel flow is the baseline Flow++.

The model is trained with an Adam optimizer (Kingma & Ba, 2015) with a batch size 64 for 2,000 epochs. Following (Ho et al., 2019), the learning rate linearly warms up to 0.0012 during the first 2,000 training steps, and exponentially decays at a rate of 0.99999 per step starting from the 50,000-th step until it reaches 0.0003. All the experiments are run on 16 RTX 2080Ti GPUs. The model quality is measured by bits per dimension (bpd) (Van Oord et al., 2016), where smaller bpd implies higher likelihood and better modeling quality. The likelihood $P(x; \theta)$ is evalu-


ated with importance sampling as Eq. (5) with $S = 4096$ samples for CIFAR-10 and $S = 1024$ for ImageNet.

6.1. Improving Existing Models

We compare a 6-channel VFlow with existing generative flows in Table 1, where the hyperparameters $D_H = 96$ and $B = 10$ are set identical to Flow++. By augmenting the number of channels from 3 to 6, For CIFAR-10, VFlow improves the bpd from 3.08 of Flow++ to 2.98. Samples from Flow++ and VFlow are shown in Fig. 5.

6.2. Ablation Study under Fixed Parameter Budget

To further investigate the impact of the dimensionality, we vary the number of channels under a fixed 4 million parameter budget on the CIFAR-10 dataset. In this set of experiments, we randomly hold out 10,000 samples from



(a) Flow++ (3.08 bpd)



(b) VFlow (2.98 bpd)

Figure 5. Random samples, where (a) is reprinted from (Ho et al., 2019).

the training set for validation. As the dimensionality grows, we reduce the number of hidden layers B to stay within the parameter budget. The training curve and final bpd are reported in Fig. 6 and Table 2 respectively, which clearly show that increased dimensionality is beneficial. Moreover, VFlow achieves better results at all stages of training. This supports Remark 1 in Sec. 3.2 that VFlow is still better even taken optimization issues into account. For this $D_H = 32$ network, a 6-channel VFlow, which has 24 channels on the 16×16 scale, is already sufficient to resolve the bottleneck problem. Going beyond 6 channels has marginal improvement on the model quality, while increasing the number of parameters.

Interestingly, Fig. 6 suggests that besides improved model capacity, the generalization gap of VFlow is also slightly

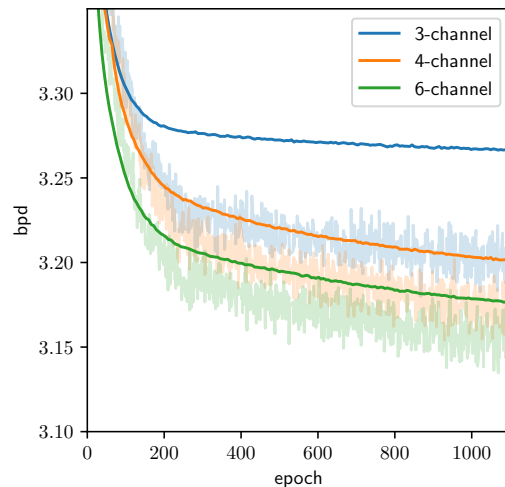


Figure 6. Bpd on training (light) and validation (dark) dataset of Flow++ and VFlow under a 4-million parameter budget (not fully converged). Here bpd is only an upper bound because we evaluate it with ELBO as Eq. (7) instead of the marginal likelihood.

Table 2. Impact of dimensionality on the CIFAR-10 dataset.

Model	bpd	Parameters	D_H	B
3-channel Flow++	3.23	4.02M	32	13
4-channel VFlow	3.16	4.03M	32	11
6-channel VFlow	3.13	4.01M	32	10

smaller than Flow++. We suspect the additional randomness introduced by \mathbf{z} acts as an implicit regularization.

6.3. Parameter Efficiency

The last set of experiments aims for more compact models with similar model capacity. As shown by Table 3, we can reduce the number of parameters of the baseline Flow++ from 31.4 million to 11.9 million, which is a 2.6 times reduction. This reduction of model size is achieved by reducing the hidden layer size D_H from 96 to 56. As we argued in Sec. 2.2, the excessive number of hidden units does not help much for a network with merely 3 channels. Increasing the dimensionality of the network (*i.e.*, the bottleneck width) is much more efficient than increasing D_H . Therefore, VFlow is not only more expressive but also more

Table 3. Parameter efficiency on CIFAR-10.

Model	bpd	Parameters	D_H	B
3-channel Flow++	3.08	31.4M	96	10
6-channel VFlow	2.98	37.8M	96	10
6-channel VFlow	3.03	16.5M	64	10
6-channel VFlow	3.08	11.9M	56	10

compact than vanilla low-dimensional flows. We emphasize that the reduction of model size come solely from resolving the bottleneck problem. Even smaller models can be forged by combining with potentially more compact architectures, such as MintNet (Song et al., 2019).

7. Conclusions

We identify the bottleneck problem which limits the capacity of generative flows. To tackle this problem, we propose VFlow, a variational data augmentation framework that pads extra dimensions to the data with learnable variational distributions for the padded data. VFlow is a generalization of vanilla generative flows, and can be combined with existing generative flows to improve their expressiveness and compactness. In our experiments on the CIFAR-10 dataset, VFlow achieves a new state-of-the-art 2.98 bpd, while retaining the 3.08 bpd of vanilla Flow++ with 2.6 times less parameters.

Acknowledgements

We thank Chongxuan Li, Yucen Luo, Ziyu Wang, Kun Xu, Fan Bao, Shihong Song and Qian Fu for proofreading. This work was supported by the National Key Research and Development Program of China (No. 2017YFA0700904), NSFC Projects (Nos. 61620106010, U19B2034, U1811461), Beijing NSF Project (No. L172037), Beijing Academy of Artificial Intelligence (BAAI), Tsinghua-Huawei Joint Research Program, Huawei Hisilicon Kirin solution, the MindSpore team, a grant from Tsinghua Institute for Guo Qiang, Tiangong Institute for Intelligent Computing, and the NVIDIA NVAIL Program with GPU/DGX Acceleration.

References

- Behrmann, J., Grathwohl, W., Chen, R. T., Duvenaud, D., and Jacobsen, J.-H. Invertible residual networks. In *International Conference on Machine Learning*, pp. 573–582, 2019.
- Chen, T. Q., Rubanova, Y., Bettencourt, J., and Duvenaud, D. K. Neural ordinary differential equations. In *Advances in Neural Information Processing Systems*, pp. 6571–6583, 2018.
- Chen, T. Q., Behrmann, J., Duvenaud, D. K., and Jacobsen, J.-H. Residual flows for invertible generative modeling. In *Advances in Neural Information Processing Systems*, pp. 9913–9923, 2019.
- Chen, X., Kingma, D. P., Salimans, T., Duan, Y., Dhariwal, P., Schulman, J., Sutskever, I., and Abbeel, P. Variational lossy autoencoder. In *International Conference on Learning Representations*, 2017.
- Choi, H., Jang, E., and Alemi, A. A. Waic, but why? generative ensembles for robust anomaly detection. *arXiv preprint arXiv:1810.01392*, 2018.
- Dinh, L., Krueger, D., and Bengio, Y. Nice: Non-linear independent components estimation. In *International Conference on Learning Representations Workshop*, 2014.
- Dinh, L., Sohl-Dickstein, J., and Bengio, S. Density estimation using real nvp. In *International Conference on Learning Representations*, 2017.
- Dupont, E., Doucet, A., and Teh, Y. W. Augmented neural odes. In *Advances in Neural Information Processing Systems*, pp. 3134–3144, 2019.
- Durkan, C., Bekasov, A., Murray, I., and Papamakarios, G. Neural spline flows. In *Advances in Neural Information Processing Systems*, pp. 7509–7520, 2019.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, pp. 2672–2680, 2014.
- Grathwohl, W., Chen, R. T., Bettencourt, J., Sutskever, I., and Duvenaud, D. Ffjord: Free-form continuous dynamics for scalable reversible generative models. In *International Conference on Learning Representations*, 2019.
- Gulrajani, I., Kumar, K., Ahmed, F., Taiga, A. A., Visin, F., Vazquez, D., and Courville, A. Pixelvae: A latent variable model for natural images. In *International Conference on Learning Representations*, 2017.
- Gybenko, G. Approximation by superposition of sigmoidal functions. *Mathematics of Control, Signals and Systems*, 2(4):303–314, 1989.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 770–778, 2016.
- 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*, pp. 2722–2730, 2019.
- Hoogeboom, E., Van Den Berg, R., and Welling, M. Emerging convolutions for generative normalizing flows. In *International Conference on Machine Learning*, pp. 2771–2780, 2019.

- Huang, C.-W., Dinh, L., and Courville, A. Augmented normalizing flows: Bridging the gap between generative flows and latent variable models. *arXiv:2002.07101*, 2020.
- Jacobsen, J.-H., Smeulders, A., and Oyallon, E. i-revnet: Deep invertible networks. In *International Conference on Learning Representations*, 2018.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2015.
- Kingma, D. P. and Dhariwal, P. Glow: Generative flow with invertible 1x1 convolutions. In *Advances in Neural Information Processing Systems*, pp. 10215–10224, 2018.
- Kingma, D. P. and Welling, M. Auto-encoding variational bayes. In *International Conference on Learning Representations*, 2014.
- Kingma, D. P., Salimans, T., Jozefowicz, R., Chen, X., Sutskever, I., and Welling, M. Improved variational inference with inverse autoregressive flow. In *Advances in Neural Information Processing Systems*, pp. 4743–4751, 2016.
- Mhaskar, H. N. Approximation properties of a multilayered feedforward artificial neural network. *Advances in Computational Mathematics*, 1(1):61–80, 1993.
- Morrow, R. and Chiu, W.-C. Variational autoencoders with normalizing flow decoders, 2020. URL <https://openreview.net/forum?id=r1eh30NFwB>.
- Müller, T., McWilliams, B., Rousselle, F., Gross, M., and Novák, J. Neural importance sampling. *ACM Transactions on Graphics (TOG)*, 38(5):145, 2019.
- Nalisnick, E., Matsukawa, A., Teh, Y. W., Gorur, D., and Lakshminarayanan, B. Hybrid models with deep and invertible features. In *International Conference on Machine Learning*, pp. 4723–4732, 2019.
- Papamakarios, G., Pavlakou, T., and Murray, I. Masked autoregressive flow for density estimation. In *Advances in Neural Information Processing Systems*, pp. 2338–2347, 2017.
- Prenger, R., Valle, R., and Catanzaro, B. Waveglow: A flow-based generative network for speech synthesis. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 3617–3621. IEEE, 2019.
- Rezende, D. and Mohamed, S. Variational inference with normalizing flows. In *International Conference on Machine Learning*, pp. 1530–1538, 2015.
- Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A. C., and Fei-Fei, L. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015. doi: 10.1007/s11263-015-0816-y.
- Song, Y., Meng, C., and Ermon, S. Mintnet: Building invertible neural networks with masked convolutions. In *Advances in Neural Information Processing Systems*, pp. 11002–11012, 2019.
- Tan, M. and Le, Q. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International Conference on Machine Learning*, pp. 6105–6114, 2019.
- Van Oord, A., Kalchbrenner, N., and Kavukcuoglu, K. Pixel recurrent neural networks. In *International Conference on Machine Learning*, pp. 1747–1756, 2016.
- Yang, G., Huang, X., Hao, Z., Liu, M.-Y., Belongie, S., and Hariharan, B. Pointflow: 3d point cloud generation with continuous normalizing flows. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 4541–4550, 2019.
- Zagoruyko, S. and Komodakis, N. Wide residual networks. *arXiv preprint arXiv:1605.07146*, 2016.