# **Super-resolution Variational Auto-Encoders**

Ioannis Gatopoulos<sup>123</sup> Maarten Stol<sup>3</sup> Jakub M. Tomczak<sup>2</sup>

# Abstract

The framework of variational autoencoders (VAEs) provides a principled method for jointly learning latent-variable models and corresponding inference models. However, the main drawback of this approach is the blurriness of the generated images. Some studies link this effect to the objective function, namely, the (negative) log-likelihood (nll). Here, we propose to enhance VAEs by adding a random variable that is a downscaled version of the original image and still use the log-likelihood function as the learning objective. Further, by providing the downscaled image as an input to the decoder, it can be used in a manner similar to the superresolution. We present empirically that the proposed approach performs comparably to VAEs in terms of the nll, but it obtains a better Fréchet Inception Distance (FID) score in data synthesis.

# 1. Introduction

Unlike many other sensory systems, the human visual system (i.e. the components from the eye to neural circuits) develops largely after birth, especially in the first few years of life (Banks & Salapatek, 1978). In the beginning, even though the visual structures are fully present, they are still immature in their potentials. As the neural circuits adapt to natural light, they learn to enhance the already known signals with the new information that is becoming available (Figure 1). Furthermore, (Ayzenberg, 2019) suggested that the human visual system for object recognition tasks initially starts with the skeletal structure of the object and then maps other properties, such as textures and colors, onto it before it is classified. It seems that humans reinforce their capabilities by sequentially applying new content of



Figure 1: Infant vision development during the first six months. The previous processed information is enriched with new signals in order for the sensory systems to be able to analyze high fidelity signals (Haskett, 2019).

information over time and some specific processes, like object detection, are divided into two or more simpler tasks.

Inspired by this learning procedure, we formulate a generative model that mimics, to some extent, the human visual process. Specifically, we enhance the framework of Variational Auto-Encoder (VAEs) by introducing a downscaled representation of the image as a random variable, and utilize it in a super-resolution manner (Chang et al., 2004; Dong et al., 2015; Freeman et al., 2002) to generate high quality images. As a result, we obtain a two-level VAE with three latent variables, where one is the downscaled version of the original image.

In summary, our contributions are as follows:

- We present a powerful Variational Auto-Encoder that consists of a novel DenseNet-based encoder, a DenseNet-based decoder, and a flow-based prior. It achieves SOTA in terms of the log-likelihood function among singe-leveled VAEs.
- We propose a new class of VAEs that contain a superresolution part for generating crisp images, and is still trained using the log-likelihood objective.
- We present empirical results on CIFAR-10 and ImageNet32 where our approach achieves descent scores in terms of the bits per dimension (bpd) on CIFAR-10 and ImageNet32, and impressive FID scores.

#### 2. Variational Auto-Encoders

Let  $\mathbf{X} = {\mathbf{x}_1, ..., \mathbf{x}_N}$  with  $\mathbf{x}_n \in \mathbb{R}^D$  be the observable data that we wish to model. Further, we consider a latent variable model with latent (unobserved) variables  $\mathbf{z} \in \mathbb{R}^M$ ,

<sup>&</sup>lt;sup>1</sup>Department of Computer Science, University of Amsterdam, Amsterdam, The Netherlands <sup>2</sup>Department of Computer Science, Vrije Universiteit Amsterdam, Amsterdam, The Netherlands <sup>3</sup>BrainCreators B.V., Amsterdam, The Netherlands. Correspondence to: Ioannis Gatopoulos <ioannis.gatopoulos@gmail.com>.

Second workshop on *Invertible Neural Networks, Normalizing Flows, and Explicit Likelihood Models* (ICML 2020), Virtual Conference. Copyright 2020 by the author(s).

 $p_{\vartheta}(\mathbf{x}) = \int p_{\vartheta}(\mathbf{x}, \mathbf{z}) d\mathbf{z}$ , where  $\vartheta$  denotes parameters. We consider the optimization through maximum likelihood estimation (MLE) of  $p_{\vartheta}(\mathbf{x})$ , however, it becomes infeasible due to the intractability of the integration at hand. One possible way of overcoming this issue and obtaining a highly scalable framework is by introducing an amortized variational family Q in order to identify its member  $q_{\phi}(\mathbf{z}|\mathbf{x})$  that minimizes the Kullback-Leibler divergence to the *real* posterior  $p(\mathbf{z}|\mathbf{x})$ . In consequence, we derive a tractable objective function, namely the *evidence lower bound* (ELBO) (Jordan et al., 1999):

$$\log p_{\vartheta}(\mathbf{x}) \geq \mathbb{E}_{q_{\phi}(\mathbf{z}|\mathbf{x})} \log \left[ \frac{p_{\theta}(\mathbf{x}, \mathbf{z})}{q_{\phi}(\mathbf{z}|\mathbf{x})} \right]$$
$$= \mathbb{E}_{q_{\phi}(\mathbf{z}|\mathbf{x})} \left[ \log p_{\theta}(\mathbf{x}|\mathbf{z}) - \log q_{\phi}(\mathbf{z}|\mathbf{x}) + \log p_{\lambda}(\mathbf{z}) \right]$$
$$\equiv \mathcal{L}(\theta, \phi, \lambda), \tag{1}$$

where  $q_{\phi}(\mathbf{z}|\mathbf{x})$  is the variational posterior (or the *encoder*),  $p_{\theta}(\mathbf{x}|\mathbf{z})$  is the likelihood function (or the *decoder*) and  $p_{\lambda}(\mathbf{z})$  is the *prior* over the latent variables, parameterized by and  $\phi$ ,  $\theta$  and  $\lambda$  respectively. The optimization is done efficiently by computing the expectation by Monte Carlo integration while exploiting the *reparameterization trick* in order to obtain an unbiased estimator of the gradients. This generative model framework is known as *Variational Auto-Encoder* (VAE) (Kingma & Welling, 2013; Rezende et al., 2014).

**VAE with a bijective prior** Even though the lowerbound suggests that the prior plays a crucial role in improving the variational bounds, usually it is modelled by a fixed distribution (i.e., a standard multivariate Gaussian). While being relatively simple and computationally cheap, a fixed prior is known to result in over-regularized models that tend to ignore more of the latent dimensions (Burda et al., 2015; Tomczak & Welling, 2017). Moreover, as the objective function is optimised to match the variational posterior with the prior, (Rosca et al., 2018) argued that even if the former becomes the optimal one, namely the *aggregated posterior*, it may still not match a unit Gaussian distribution.

However, it is possible to obtain a rich, multi-modal prior distribution  $p(\mathbf{z})$  by using a *bijective model*. Formally, given a latent code  $\mathbf{z} \sim q_Z(\mathbf{z}|\mathbf{x})$ , a base distribution  $p_V(\mathbf{v})$ on a latent variable  $\mathbf{v} \in V$ , and  $f: V \to Z$  consisting of a sequence of L diffeomorphic transformations<sup>1</sup>, where  $f_i(\mathbf{v}_{i-1}) = \mathbf{v}_i$ ,  $\mathbf{v}_0 = \mathbf{v}$  and  $\mathbf{v}_L = \mathbf{z}$ , the sequential use of the *change of variable* can be used to express the distribution of  $\mathbf{z}$  as a function of  $\mathbf{v}$  as follows:

$$\log p_Z(\mathbf{z}) = \log p_V(\mathbf{v}) - \sum_{i=1}^L \log \left| \frac{\partial f_i(\mathbf{v}_{i-1})}{\partial \mathbf{v}_{i-1}} \right|, \quad (2)$$

where  $\left|\frac{\partial f_i(\mathbf{v}_{i-1})}{\partial \mathbf{v}_{i-1}}\right|$  is the Jacobian-determinant of the  $i^{th}$  transformation.

Thus, using the transformed prior we end up with the following training objective function:

$$\mathcal{L}(\theta, \phi, \lambda) = \mathbb{E}_{q_{\phi}(\mathbf{z}|\mathbf{x})} \Big[ \log p_{\theta}(\mathbf{x}|\mathbf{z}) - \log q_{\phi}(\mathbf{z}|\mathbf{x}) + \\ + \log p_{V}(\mathbf{v}_{0}) + \sum_{i=1}^{L} \log \left| \frac{\partial f_{i}^{-1}(\mathbf{v}_{i})}{\partial \mathbf{v}_{i}} \right| \Big].$$
(3)

In this paper, we utilize RealNVP (Dinh et al., 2016) as the prior, however, any other flow-based model could be used (van den Berg et al., 2018; Kingma & Dhariwal, 2018).

#### 3. Our method

#### 3.1. Model formulation

Let us introduce an additional variable  $\mathbf{y} \in \mathbb{R}^{C}$  that is a compressed representation of  $\mathbf{x}^{2}$ , where  $C \leq D$ . Further, let  $\mathbf{u} \in \mathbb{R}^{K}$  and  $\mathbf{z} \in \mathbb{R}^{M}$  be two stochastic latent variables that interact with the above observed ones in a way that is presented in Figure 2.

From the dependencies of the considered probabilistic graphical model, we can write the joint probability as follows:

$$p(\mathbf{x}, \mathbf{y}, \mathbf{z}, \mathbf{u}) = p(\mathbf{x} | \mathbf{y}, \mathbf{z}) p(\mathbf{z} | \mathbf{y}, \mathbf{u}) p(\mathbf{y} | \mathbf{u}) p(\mathbf{u})$$

Then, we define the amortized variational posterior of  $p(\mathbf{y}, \mathbf{z}, \mathbf{u} | \mathbf{x})$  as follows:

$$q(\mathbf{y}, \mathbf{z}, \mathbf{u} | \mathbf{x}) = q(\mathbf{z} | \mathbf{y}, \mathbf{x}) q(\mathbf{u} | \mathbf{y}) q(\mathbf{y} | \mathbf{x}) \equiv q(\mathbf{w} | \mathbf{x})$$

where  $\mathbf{w} = {\mathbf{y}, \mathbf{z}, \mathbf{u}}$ , and derive the corresponding lower bound of the likelihood function in the following manner:

$$\log p(\mathbf{x}) \geq \mathbb{E}_{q(\mathbf{w})} \log \frac{p(\mathbf{x}, \mathbf{w})}{q(\mathbf{w})}$$
$$= \mathbb{E}_{q(\mathbf{w})} \Big[ \log p(\mathbf{x}, \mathbf{w}) \Big] - \mathbb{E}_{q(\mathbf{w})} \Big[ \log q(\mathbf{w}) \Big]$$
(4)
$$\equiv \mathcal{L}(\mathbf{x}).$$

After expanding and rearranging the above objective function (please see A.1 for full derivation), we obtain:

$$\mathcal{L}(\mathbf{x}) = \mathbb{E}_{q(\mathbf{z}|\mathbf{x},\mathbf{y}) q(\mathbf{y}|\mathbf{x})} \log p_{\theta}(\mathbf{x}|\mathbf{y},\mathbf{z}) + \\ - \mathbb{E}_{q(\mathbf{u}|\mathbf{y})} \mathcal{K} \mathcal{L}(q(\mathbf{y}|\mathbf{x})) || p_{\theta}(\mathbf{y}|\mathbf{u})) + \\ - \mathbb{E}_{q(\mathbf{u}|\mathbf{y}) q(\mathbf{y}|\mathbf{x})} \mathcal{K} \mathcal{L}(q(\mathbf{z}|\mathbf{x},\mathbf{y})) || p(\mathbf{z}|\mathbf{y},\mathbf{u})) + \\ - \mathbb{E}_{q(\mathbf{y}|\mathbf{x})} \mathcal{K} \mathcal{L}(q(\mathbf{u}|\mathbf{y})) || p(\mathbf{u})),$$
(5)

where  $\mathcal{KL}(\cdot || \cdot)$  denotes the Kullback-Leibler divergence.

<sup>&</sup>lt;sup>1</sup>That is, invertible and differentiable transformations.

<sup>&</sup>lt;sup>2</sup>Here, we consider a downscaled  $\mathbf{x}$  however, our framework allows for any compressed transformation of the original data.



Figure 2: Stochastic dependencies of the proposed model. Our approach takes advantage of a compressed representation y of the data in the variational part, that is then utilized in the super-resolution in the generative part.

## 3.2. Properties

There are two main properties that we are going to take advantage of.

A. IF  $q(\mathbf{y}|\mathbf{x})$  is both **deterministic** and **discrete**, then  $\mathbb{E}_{q(\mathbf{y}|\mathbf{x})} \left[ \log q(\mathbf{y}|\mathbf{x}) \right] = 0.$ 

Dependence between two random variables can take a variety of forms, of which *stochastic independence* and *functional dependence* can be argued to be most opposite in character. In the former case, neither of the variables provide any information about each other, whereas in the latter, there is a full determination. Even though the proposed framework allows to model  $q(\mathbf{y}|\mathbf{x})$  as a stochastic dependency, the choice of a deterministic relationship is more attractive, as the transformations to a compressed representation are usually available (e.g., a downscaled image), the optimization process would be faster, easier and the model overall will require less trainable parameters.

We will define this deterministic transformation as a *degenerate* probability distribution which provides a way to deal with constant values in a probabilistic framework. It trivially gives rise to a probability mass function satisfying  $P(\Omega) = 1$  and has an expectation of a constant value  $c \in \mathbb{R}$ , a variance of 0 and most importantly, its entropy is also equal to 0. Thus, modelling the distribution  $q(\mathbf{y}|\mathbf{x})$  as a discrete degenerate distribution yields:

$$\mathbb{H}[q(\mathbf{y}|\mathbf{x})] = 0 \Leftrightarrow \mathbb{E}_{q(\mathbf{y}|\mathbf{x})}[-\log q(\mathbf{y}|\mathbf{x})] = 0$$

which simplifies the derived objective function in (5).

# B. The distribution $q(\mathbf{z}|\mathbf{y}, \mathbf{x})$ can be simplified to $q(\mathbf{z}|\mathbf{x})$ .

One of the core motivations behind the architecture of the two staged approach is that the latent variable z will be able to capture the missing information between x and y. While u would allow to produce the global structure of the data (e.g., a shape of a horse), the variation of z will alter high-level features (e.g., varying z will result into a different color of a horse). Thus, since y is a compressed representation of x, it does not introduce any additional information about z that is not already in x. This intuitively allows to model z only using x, and, in essence, to replace  $q(\mathbf{z}|\mathbf{y}, \mathbf{x})$  with  $q(\mathbf{z}|\mathbf{x})$ .

#### FINAL ELBO

With these two properties in mind, the final lower bound of the marginal likelihood of  $\mathbf{x}$  is the following:

$$\mathcal{L}(\mathbf{x}) = \underbrace{\mathbb{E}_{q(\mathbf{z}|\mathbf{x}) q(\mathbf{y}|\mathbf{x})} \log p_{\theta}(\mathbf{x}|\mathbf{y}, \mathbf{z})}_{\text{RE}_{x}} + \underbrace{\mathbb{E}_{q(\mathbf{u}|\mathbf{y})q(\mathbf{y}|\mathbf{x})} \log p_{\theta}(\mathbf{y}|\mathbf{u})}_{\text{RE}_{y}} + \underbrace{\mathbb{E}_{q(\mathbf{u}|\mathbf{y})q(\mathbf{y}|\mathbf{x})} \log p_{\theta}(\mathbf{y}|\mathbf{u})}_{\text{RE}_{y}} + \underbrace{\mathbb{E}_{q(\mathbf{u}|\mathbf{y})q(\mathbf{y}|\mathbf{x})} \mathcal{K}\mathcal{L}(q(\mathbf{z}|\mathbf{x}))||p(\mathbf{z}|\mathbf{y},\mathbf{u}))}_{\text{KL}_{z}} + \underbrace{\mathbb{E}_{q(\mathbf{y}|\mathbf{x})} \mathcal{K}\mathcal{L}(q(\mathbf{u}|\mathbf{y}))||p(\mathbf{u}))}_{\text{KL}_{u}} + \underbrace{\mathbb{E}_{q(\mathbf{y}|\mathbf{x})} \mathcal{K}\mathcal{L}(q(\mathbf{u}|\mathbf{y}))||p(\mathbf{u})}_{\text{KL}_{u}} + \underbrace{\mathbb{E}_{q(\mathbf{y}|\mathbf{x})} \mathcal{K}\mathcal{L}(q(\mathbf{u}|\mathbf{y}))||p(\mathbf{u})}_{\text{KL}_{u}} + \underbrace{\mathbb{E}_{q(\mathbf{y}|\mathbf{x})} \mathcal{K}\mathcal{L}(q(\mathbf{u}|\mathbf{y}))||p(\mathbf{u})}_{\text{KL}_{u}} + \underbrace{\mathbb{E}_{q(\mathbf{y}|\mathbf{x})} \mathcal{K}\mathcal{L}(q(\mathbf{u}|\mathbf{y}))||p(\mathbf{u})}_{\text{KL}_{u}} + \underbrace{\mathbb{E}_{q(\mathbf{y}|\mathbf{x})} \mathcal{K}\mathcal{L}(\mathbf{y}|\mathbf{x})}_{\text{KL}_{u}} + \underbrace{\mathbb{E}_{q(\mathbf{y}|\mathbf{x})} + \underbrace{\mathbb{E}_{q(\mathbf{y}|\mathbf{x})} \mathcal{K}\mathcal{L}(\mathbf{y}|\mathbf{x})}_{\text{KL}_{u}} + \underbrace{\mathbb{E}_{q(\mathbf{y}|\mathbf{x})} + \underbrace{\mathbb{E}_{q(\mathbf{x$$

#### 3.3. Super-resolution VAE (srVAE)

We choose the following distributions in our model:

$$\begin{aligned} q_{\phi_1}\left(\mathbf{u}|\mathbf{y}\right) &= \mathcal{N}\left(\mathbf{u}|\boldsymbol{\mu}_{\phi_1}(\mathbf{y}), \operatorname{diag}\left(\boldsymbol{\sigma}_{\phi_1}(\mathbf{y})\right)\right) \\ q(\mathbf{y}|\mathbf{x}) &= \delta(\mathbf{y} = d(\mathbf{x})) \\ q_{\phi_2}\left(\mathbf{z}|\mathbf{x}\right) &= \mathcal{N}\left(\mathbf{z}|\boldsymbol{\mu}_{\phi_2}(\mathbf{x}), \operatorname{diag}\left(\boldsymbol{\sigma}_{\phi_2}(\mathbf{x})\right)\right) \\ p_{\lambda}\left(\mathbf{u}\right) &= p(\mathbf{v})\prod_{i=1}^{L} \left| \operatorname{det} \frac{\partial f_i(\mathbf{v}_{i-1})}{\partial \mathbf{v}_{i-1}} \right|^{-1} \\ p(\mathbf{v}) &= \mathcal{N}\left(\mathbf{v}|\mathbf{0},\mathbf{1}\right) \\ p_{\theta_1}\left(\mathbf{y}|\mathbf{u}\right) &= \sum_{i=1}^{K} \pi_i^{(\mathbf{u})} \operatorname{Dlogistic}\left(\boldsymbol{\mu}_i^{(\mathbf{u})}, s_i^{(\mathbf{u})}\right) \\ p_{\theta_2}\left(\mathbf{z}|\mathbf{y},\mathbf{u}\right) &= \mathcal{N}\left(\mathbf{z}|\boldsymbol{\mu}_{\theta_2}(\mathbf{y},\mathbf{u}), \operatorname{diag}\left(\boldsymbol{\sigma}_{\theta_2}(\mathbf{y},\mathbf{u})\right)\right) \\ p_{\theta_3}\left(\mathbf{x}|\mathbf{z},\mathbf{y}\right) &= \sum_{i=1}^{K} \pi_i^{(\mathbf{z},\mathbf{y})} \operatorname{Dlogistic}\left(\boldsymbol{\mu}_i^{(\mathbf{z},\mathbf{y})}, s_i^{(\mathbf{z},\mathbf{y})}\right). \end{aligned}$$

where Dlogistic is defined as the discretized logistic distribution (Salimans et al., 2017),  $\delta(\cdot)$  is the Dirac's delta, and  $d(\mathbf{x})$  denotes the downscaling transformation that returns a discrete values.

Table 1: Negative log-likelihood for CIFAR-10 and ImageNet32 test set.	. The dataset marked with the symbol $^{\dagger}$ denotes a
model trained only on CIFAR-10 and evaluated on ImageNet32. For FID,	we provide values obtained on the test set and the
training set (in brackets).	

Dataset	Model	nll	reconstruction loss		regularization loss		FID
		(bits/dim)	$RE_x$	$\operatorname{RE}_y$	$\mathrm{KL}_z$	$KL_u$	TID
Cifar10	VAE	3.51	5540	-	1966	-	41.36 (37.25)
	srVAE	3.65	5107	1241	619	819	34.71 (29.95)
ImageNet32 <sup>†</sup>	VAE	3.80	6386	-	1707	-	51.82 (N/A)
	srVAE	4.00	5907	1257	597	805	45.37 (N/A)

In VAEs, it is possible to use the following functionality:

- Generation: The model is able to generate new images through the following process:  $\mathbf{z} \sim p(\mathbf{z}) \rightarrow \mathbf{x} \sim p(\mathbf{x}|\mathbf{z})$ .
- Reconstruction: The model allows to reconstruct x by using the following scheme: x → z ~ q(z|x) → x ~ p(x|z).

Interestingly, our approach allows four operations:

- Generation: The model allows to generate novel content by applying the following hierarchical sampling process:  $\mathbf{u} \sim p(\mathbf{u}) \rightarrow \mathbf{y} \sim p(\mathbf{y}|\mathbf{u}) \rightarrow \mathbf{z} \sim p(\mathbf{z}|\mathbf{u}, \mathbf{y}) \rightarrow \mathbf{x} \sim p(\mathbf{x}|\mathbf{z}, \mathbf{y}).$
- Conditional Generation (or Super-Resolution Generation): Given y, we can sample the latent codes: u ~ q(u|y) → z ~ p(z|y, u), → x ~ p(x|z, y).
- **Reconstruction:** Similarly to standard VAE, we can reconstruct  $\mathbf{x}$ :  $\mathbf{y} \sim q(\mathbf{y}|\mathbf{x}) \rightarrow \mathbf{z} \sim q(\mathbf{z}|\mathbf{x}) \rightarrow \mathbf{x} \sim p(\mathbf{x}|\mathbf{z}, \mathbf{y})$ .
- Generative Reconstruction: Additionally, we can reconstruct  $\mathbf{x}$  by combining the generation and the reconstruction:  $\mathbf{y}^* \sim q(\mathbf{y}^*|\mathbf{x}) \rightarrow \mathbf{u} \sim q(\mathbf{u}|\mathbf{y}^*) \rightarrow$  $\mathbf{y} \sim p(\mathbf{y}|\mathbf{u}) \rightarrow \mathbf{z} \sim p(\mathbf{z}|\mathbf{y}, \mathbf{u}), \rightarrow \mathbf{x} \sim p(\mathbf{x}|\mathbf{z}, \mathbf{y}).$

In order to highlight the super-resolution part in our model, we refer to it as the *super-resolution VAE* (srVAE).

#### 4. Experiments

#### 4.1. Setup

We evaluated the following two models for density estimation; (i) a VAE (ii) the proposed two-level VAE (sr-VAE). In both models we employed RealNVP as a bijective prior (Dinh et al., 2016). Specifically, for our model, even though the compressed image y can be given by any deterministic and discrete transformation of the input data (i.e., the image label, grey-scale transformation, Fourier transform, sketch representation), we provide results with a  $2\times$  downscaled image. The downscaled images still preserve the global structure of the samples while they disregard the high resolution details. Moreover, it will allow us to evaluate the model for its ability to perform super-resolution tasks. We set  $K = M = 16 \times 8 \times 8$  as latent dimentions in our experiments. The building blocks of the neural network implementation details are described in the Appendix, see Figure 4. We used a composition of DenseNets (Huang et al., 2016) and channel attention (Zhang et al., 2018) with ELUs (Clevert et al., 2015) as activation functions.

We applied the proposed model to CIFAR-10 for quantitative and qualitative evaluation of natural images. Additionally, we applied the model trained on CIFAR-10 to ImageNet32, without any additional fine-tuning in order to illustrate its adaption performance to a similar dataset. We evaluate the density estimation performance by *bits per dimension* (bits/dim),  $\mathcal{L}/(hwc \cdot \log(2))$ , where *h*, *w* and *c* denote the height, width, and channels, respectively and we use the *Fréchet Inception Distance* (FID) (Heusel et al., 2017) as a metric for image generation quality. The negative log-likelihood value (*nll*) was estimated using 500 weighted samples (Burda et al., 2015), and for computing the FID scored we used 10k generated images and 10k real images from the test set, but also 50k generated images and 50k real images from the train set.

The code for this paper is available at https://github.com/ioangatop/srVAE.

#### 4.2. Evaluation on Natural Images

**Quantitative results** The density estimation and the image generation scores on CIFAR-10 and ImageNet32 are presented in Table 1. Even though the VAE with the RealNVP prior follows an architecture without the use of any auto-regressive components and a single stochastic la-

Super-resolution Variational Auto-Encoders



CM: Compressed, OG: Original, SR: Super-Resolution, RS: Reconstruction, GR: Generative Reconstruction

Figure 3: Qualitative results of the VAE (a) and the srVAE (b) on CIFAR-10.

tent variable, it achieves a very competitive log-likelihood score (see Table 2 in the Appendix). The importance of the data-driven prior is further supported by the low FID, as it manages to outperform various flow-based generative models (see the Appendix, Table 3). Finally, analyzing results in Table 1, we make two observations. First, the  $RE_x$  is better in case of our approach. This result is to be expected since our model contains a super-resolution part. However, we pay a price for that, namely, we have an extra error coming from  $RE_y$ . Second, both  $KL_z$  and  $KL_u$  are relatively large, and, thus, we claim the model does not suffer from the posterior collapse. Interestingly, the KL part of the VAE is 1966 and 1707 for CIFAR-10 and ImageNet32, respectively, while our model achieves the sum of  $KL_z$  and  $KL_u$  around 1500 on both datasets. This result suggests that introducing an additional random variable y helps to match the variational posteriors and the (conditional) priors, and the model does not bypass the latent variable z, verifying its importance.

Even though our model outperforms the VAE in the reconstruction loss of the original image ( $RE_x$ ), due to the summation with the value of  $RE_y$  it results in a poorer likelihood. However, we see that the srVAE significantly improves the FID score, as it produces more coherent and visually pleasing generations.

**Qualitative results** We test the performance of the two models on image generation, reconstruction, and in the case of our model, additionally for conditional generation (super-resolution) and generative reconstruction tasks on CIFAR-10. The results are illustrated in Figure 3. The VAE with the bijective prior showcases an excellent performance on the natural image reconstruction task, which

is contrary to the performance that is often provided in the literature. This maybe be associated with the effectiveness of a powerful, invertible, data-driven prior like RealNVP and its ability to boost the performance significantly with negligible sacrifice on generation speed, and none on inference. The provided unconditional generations, instead of being characterised as blurry, manage to output images with a coherent global structure.

In contrast with the VAE, our approach breaks the generation of an image into a two-step process. It first generates a compressed sample through the latent variable u, and then adds local structure with the help of the stochastic variable z. The provided results illustrate that indeed the generations of the first step outputs an outline as a general concept which is then enriched with additional components, resulting in a sharp image. While a proportion of the generations of VAE tend to be noisy and abstract, the two-staged approach seems to generate smoother, higher fidelity results. Moreover, due to our choice of the model, i.e. the compressed image as a  $2 \times$  downscaled representation, the unconditional generation functionality is essentially a super-resolution task. The model manages to perform accurate reconstructions of the original images, providing a novel generative approach.

More detailed results and analysis of the conducted experiments are provided in the Appendix A.3.

# 5. Conclusion

We propose a new type of generative model which is able to perform both conditional and unconditional sampling which demonstrate improved quantitative performance in terms of FID of the generating sample on standard image modelling benchmark. In addition, we demonstrate that VAEs employed with a RealNVP prior can result in a competitive density estimation performance, despite its non-autoregressive architecture form and a single stochastic latent variable. Our approach opens new directions in the VAE framework. First, it allows usage of the log-likelihood-based objective to generate crisp images. Second, the introduction of a downscaled image in the framework alleviates common issues in learning latent variables. Third, it introduces the super-resolution into the VAE framework. All these aspects could be further studied and developed to obtain better quality of generated images.

#### References

- Ayzenberg, V., L. S. Skeletal descriptions of shape provide unique perceptual information for object recognition. *Scientific Reports*, 9, 06 2019. ISSN 1552-5783.
- Banks, M. S. and Salapatek, P. Acuity and contrast sensitivity in 1-, 2-, and 3-month-old human infants. *Investigative Ophthalmology and Visual Science*, 17(4): 361–365, 04 1978. ISSN 1552-5783.
- Burda, Y., Grosse, R., and Salakhutdinov, R. Importance weighted autoencoders, 2015.
- Chang, H., Yeung, D.-Y., and Xiong, Y. Super-resolution through neighbor embedding. In *Proceedings of the* 2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2004. CVPR 2004., volume 1, pp. I–I. IEEE, 2004.
- Chen, R. T. Q., Behrmann, J., Duvenaud, D., and Jacobsen, J.-H. Residual flows for invertible generative modeling, 2019.
- Chen, X., Mishra, N., Rohaninejad, M., and Abbeel, P. Pixelsnail: An improved autoregressive generative model, 2017.
- Clevert, D.-A., Unterthiner, T., and Hochreiter, S. Fast and accurate deep network learning by exponential linear units (elus), 2015.
- Dinh, L., Sohl-Dickstein, J., and Bengio, S. Density estimation using real nvp, 2016.
- Dong, C., Loy, C. C., He, K., and Tang, X. Image superresolution using deep convolutional networks. *IEEE transactions on pattern analysis and machine intelligence*, 38(2):295–307, 2015.
- Freeman, W. T., Jones, T. R., and Pasztor, E. C. Examplebased super-resolution. *IEEE Computer graphics and Applications*, 22(2):56–65, 2002.

- Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., and Courville, A. Improved training of wasserstein gans, 2017.
- Haskett, M. Human Versus Computer Vision., 2019. URL https://blinkidentity.com/ human-versus-computer-vision-2/.
- Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., and Hochreiter, S. Gans trained by a two time-scale update rule converge to a local nash equilibrium, 2017.
- Ho, J., Chen, X., Srinivas, A., Duan, Y., and Abbeel, P. Flow++: Improving flow-based generative models with variational dequantization and architecture design, 2019.
- Hoogeboom, E., Peters, J. W. T., van den Berg, R., and Welling, M. Integer discrete flows and lossless compression, 2019.
- Huang, G., Liu, Z., van der Maaten, L., and Weinberger, K. Q. Densely connected convolutional networks, 2016.
- Jordan, M. I., Ghahramani, Z., and et al. An introduction to variational methods for graphical models. In *MACHINE LEARNING*, pp. 183–233. MIT Press, 1999.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization, 2014.
- Kingma, D. P. and Dhariwal, P. Glow: Generative flow with invertible 1x1 convolutions, 2018.
- Kingma, D. P. and Welling, M. Auto-encoding variational bayes, 2013.
- Kingma, D. P., Salimans, T., Jozefowicz, R., Chen, X., Sutskever, I., and Welling, M. Improved variational inference with inverse autoregressive flow. In Lee, D. D., Sugiyama, M., Luxburg, U. V., Guyon, I., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 29*, pp. 4743–4751. Curran Associates, Inc., 2016.
- Liang, Z., Feng, Y., Guo, Y., Liu, H., Chen, W., Qiao, L., Zhou, L., and Zhang, J. Learning for disparity estimation through feature constancy, 2017.
- Maaløe, L., Fraccaro, M., Liévin, V., and Winther, O. Biva: A very deep hierarchy of latent variables for generative modeling, 2019.
- Ostrovski, G., Dabney, W., and Munos, R. Autoregressive quantile networks for generative modeling, 2018.
- Parmar, N., Vaswani, A., Uszkoreit, J., Łukasz Kaiser, Shazeer, N., Ku, A., and Tran, D. Image transformer, 2018.

- Radford, A., Metz, L., and Chintala, S. Unsupervised representation learning with deep convolutional generative adversarial networks, 2015.
- Rezende, D. J., Mohamed, S., and Wierstra, D. Stochastic backpropagation and approximate inference in deep generative models, 2014.
- Rosca, M., Lakshminarayanan, B., and Mohamed, S. Distribution matching in variational inference, 2018.
- Salimans, T. and Kingma, D. P. Weight normalization: A simple reparameterization to accelerate training of deep neural networks, 2016.
- Salimans, T., Karpathy, A., Chen, X., and Kingma, D. P. Pixelcnn++: Improving the pixelcnn with discretized logistic mixture likelihood and other modifications, 2017.
- Tomczak, J. M. and Welling, M. Vae with a vampprior. *arXiv preprint arXiv:1705.07120*, 2017.
- Vahdat, A., Macready, W. G., Bian, Z., Khoshaman, A., and Andriyash, E. Dvae++: Discrete variational autoencoders with overlapping transformations, 2018.
- van den Berg, R., Hasenclever, L., Tomczak, J. M., and Welling, M. Sylvester normalizing flows for variational inference, 2018.
- van den Oord, A., Kalchbrenner, N., and Kavukcuoglu, K. Pixel recurrent neural networks, 2016a.
- van den Oord, A., Kalchbrenner, N., Vinyals, O., Espeholt, L., Graves, A., and Kavukcuoglu, K. Conditional image generation with pixelcnn decoders, 2016b.
- Zhang, Y., Li, K., Li, K., Wang, L., Zhong, B., and Fu, Y. Image super-resolution using very deep residual channel attention networks, 2018.

# Appendix

# A.1. Derivation of the lower bound

Expanding the lower bound from (4), from the first part we will have

$$\mathbb{E}_{q(\mathbf{w})} \Big[ \log p(\mathbf{x}, \mathbf{w}) \Big] = \mathbb{E}_{q(\mathbf{w})} \Big[ \log p(\mathbf{x} | \mathbf{y}, \mathbf{z}) \, p(\mathbf{z} | \mathbf{y}, \mathbf{u}) \, p(\mathbf{y} | \mathbf{u}) \, p(\mathbf{u}) \Big] = \mathbb{E}_{q(\mathbf{z} | \mathbf{y}, \mathbf{x}) q(\mathbf{y} | \mathbf{x})} \Big[ \log p(\mathbf{x} | \mathbf{y}, \mathbf{z}) \Big] + \mathbb{E}_{q(\mathbf{z} | \mathbf{y}, \mathbf{x}) q(\mathbf{y} | \mathbf{x})} \Big[ \log p(\mathbf{z} | \mathbf{y}, \mathbf{u}) \Big] + \mathbb{E}_{q(\mathbf{u} | \mathbf{y}) q(\mathbf{y} | \mathbf{x})} \Big[ \log p(\mathbf{y} | \mathbf{u}) \Big] + \mathbb{E}_{q(\mathbf{u} | \mathbf{y}) q(\mathbf{y} | \mathbf{x})} \Big[ \log p(\mathbf{u}) \Big],$$

and for the second

$$\mathbb{E}_{q(\mathbf{w})} \Big[ \log q(\mathbf{w}) \Big] = \mathbb{E}_{q(\mathbf{z}|\mathbf{y},\mathbf{x})} \Big[ \log q(\mathbf{z}|\mathbf{y},\mathbf{x}) \Big] + \mathbb{E}_{q(\mathbf{y}|\mathbf{x})} \Big[ \log q(\mathbf{y}|\mathbf{x}) \Big] + \mathbb{E}_{q(\mathbf{u}|\mathbf{y}) q(\mathbf{y}|\mathbf{x})} \Big[ \log q(\mathbf{u}|\mathbf{y}) \Big]$$

Interestingly, plugging the above terms back to (4) and rearranging them, we will have

$$\mathcal{L}(\mathbf{x}) \stackrel{(4)}{=} \underbrace{\mathbb{E}_{q(\mathbf{z}|\mathbf{y},\mathbf{x})q(\mathbf{y}|\mathbf{x})} \Big[\log p(\mathbf{x}|\mathbf{y},\mathbf{z})\Big] + \mathbb{E}_{q(\mathbf{z}|\mathbf{y},\mathbf{x})q(\mathbf{u}|\mathbf{y})q(\mathbf{y}|\mathbf{x})} \Big[\log p(\mathbf{z}|\mathbf{y},\mathbf{u})\Big] - \mathbb{E}_{q(\mathbf{z}|\mathbf{y},\mathbf{x})} \Big[\log q(\mathbf{z}|\mathbf{y},\mathbf{x})\Big] + \underbrace{\mathbb{E}_{q(\mathbf{u}|\mathbf{y})q(\mathbf{y}|\mathbf{x})} \Big[\log p(\mathbf{y}|\mathbf{u})\Big] + \mathbb{E}_{q(\mathbf{u}|\mathbf{y})q(\mathbf{y}|\mathbf{x})} \Big[\log p(\mathbf{u})\Big] - \mathbb{E}_{q(\mathbf{y}|\mathbf{x})} \Big[\log q(\mathbf{y}|\mathbf{x})\Big] - \mathbb{E}_{q(\mathbf{u}|\mathbf{y})q(\mathbf{y}|\mathbf{x})} \Big[\log q(\mathbf{u}|\mathbf{y})\Big]}_{B}.$$

Working with term **B**, one can see that

$$\boldsymbol{B} = \mathbb{E}_{q(\mathbf{u}|\mathbf{y})q(\mathbf{y}|\mathbf{x})} \Big[ \log \frac{p(\mathbf{y}|\mathbf{u})p(\mathbf{u})}{q(\mathbf{u}|\mathbf{y})q(\mathbf{y}|\mathbf{x})} \Big],$$

which denotes a (hidden) lower bound on of the marginal  $\log p(\mathbf{y})$  with variational posterior  $q(\mathbf{u}|\mathbf{y})q(\mathbf{y}|\mathbf{x})$ . Thus, the resulted lower bound of the marginal likelihood of  $\mathbf{x}$  would be

$$\mathcal{L}(\mathbf{x}) = \mathbb{E}_{q(\mathbf{z}|\mathbf{x},\mathbf{y}) q(\mathbf{y}|\mathbf{x})} \log p_{\theta}(\mathbf{x}|\mathbf{y},\mathbf{z}) - \mathcal{K}\mathcal{L}(q(\mathbf{z}|\mathbf{x},\mathbf{y})||p(\mathbf{z}|\mathbf{y},\mathbf{u})) + \mathbb{E}_{q(\mathbf{u})} \log p_{\theta}(\mathbf{y}|\mathbf{u}) - \mathcal{K}\mathcal{L}(q(\mathbf{u}|\mathbf{y})||p(\mathbf{u}))$$
$$= \mathbb{E}_{q(\mathbf{z}|\mathbf{x},\mathbf{y}) q(\mathbf{y}|\mathbf{x})} \log p_{\theta}(\mathbf{x}|\mathbf{y},\mathbf{z}) - \mathcal{K}\mathcal{L}(q(\mathbf{z}|\mathbf{x},\mathbf{y})||p(\mathbf{z}|\mathbf{y},\mathbf{u})) + \mathbb{E}_{q(\mathbf{u}|\mathbf{y})q(\mathbf{y}|\mathbf{x})} \Big[ \log \frac{p(\mathbf{y}|\mathbf{u})p(\mathbf{u})}{q(\mathbf{u}|\mathbf{y})q(\mathbf{y}|\mathbf{x})} \Big].$$

#### A.2. Neural Network Architecture

In Figure 4 is depicted the main architecture of the VAE with the bijective prior as well as the optimization choices. The design choices for the encoder and decoder form the building blocks to every model that was trained and evaluated.

## A.3. Supplementary results

The datasets CIFAR-10 and ImageNet32 were split as described in (Hoogeboom et al., 2019). We notice that some papers in the literature use different and, in our opinion, unfair data division.

Quantitative Results Additional quantitative results are presented in Tables 2 and 3.

**Qualitative Results** Additional qualitative results of the VAE and the srVAE for CIFAR-10 are illustrated in Figures 5, 6, 7 and 10. For ImageNet32, the images are presented in Figures 8, 9 and 11.



Figure 4: Architecture of our autoencoder. On the right, there are some basic buildings block of the network. The notation as 'G' on the Conv2D channels indicate the growth rate of the densely connected network. The  $\epsilon$  indicates a random variable drawn from a standard Gaussian, which helps us to make use of the *reparametrization* trick. Until z, we refer to this architecture as *Encoder NN* and thereafter as *Decoder NN*. The former and the later form the *building blocks* to every model that we train and evaluate. All models that are evaluated had ~ 35M trainable parameters, were trained for 2 thousands epochs, using AdaMax optimizer (Kingma & Ba, 2014) and the dimensionality of all the latent variables kept at 8 × 8 × 16. We applied weight normalization on the parameters with data-depended initialisation (Salimans & Kingma, 2016).

Table 2: Generative modelling performance in bits per dimension. The symbol  $\dagger$  on the ImageNet32 data denotes that the obtained results were produced using a different downsampling method from the one that was introduced by (van den Oord et al., 2016a), indicating not a fair comparison with the other methods. In the case of Flow++, we provide the results of the variational dequantization in the brackets.

Model Family	Model	CIFAR-10	ImageNet 32x32
	PixelCNN (van den Oord et al., 2016a)	3.14	_
Autoregressive	PixelRNN (van den Oord et al., 2016a)	3.00	3.86
	Gated PixelCNN (van den Oord et al., 2016b)	3.03	3.83
	PixelCNN++ (Salimans et al., 2017)	2.92	_
	Image Transformer (Parmar et al., 2018)	2.90	3.77
	PixelSNAIL (Chen et al., 2017)	2.85	3.80
	RealNVP (Dinh et al., 2016)	3.49	4.28
	DVAE++ (Vahdat et al., 2018)	3.38	_
	Glow (Kingma & Dhariwal, 2018)	3.35	4.09
Non autorograssivo	IAF-VAE (Kingma et al., 2016)	3.11	_
	BIVA (Maaløe et al., 2019)	3.08	3.96
	Flow++ (Ho et al., 2019)	3.29 (3.08)	- (3.86)
	VAE with bijective prior (ours)	3.51	$3.80^{\dagger}$
	srVAE (ours)	3.65	$4.00^{\dagger}$

Table 3: FID scores obtained from different models trained on CIFAR-10. Lower FID implies better sample quality. All results except ours are taken from (Chen et al., 2019). In the case of our VAEs, we provide the values obtained on the test set and the training set (in brackets).

Model	FID
PixelCNN (van den Oord et al., 2016b)	65.93
PixelIQN (Ostrovski et al., 2018)	49.46
iResNet Flow (Liang et al., 2017)	65.01
GLOW (Kingma & Dhariwal, 2018)	46.90
Residual Flow (Chen et al., 2019)	46.37
DCGAN (Radford et al., 2015)	37.11
WGAN-GP (Gulrajani et al., 2017)	36.40
VAE with bijective prior (ours)	41.36 (37.25)
srVAE (ours)	34.71 (29.95)



Figure 5: Qualitative results on CIFAR-10 of the VAE with the bijective prior. (A) Interpolation (B) Reconstructions (OG: real images, RS: reconstructions) and (C) Unconditional Generations.



Figure 6: Qualitative results on CIFAR-10 of the srVAE. (A) Interpolation (B) Reconstructions (OG: real images, RS: reconstructions) and (C) Super-Resolution (CM: downscaled images, OG: real images, RS: conditional generations (super-resolution).



Figure 7: Unconditional generations of the srVAE trained on CIFAR10. The model initially generates the  $\times 2$  downscaled (compressed) representation of the image which intends to captures the "global" structure and then adds the local structure (the "details") while increasing its receptive field. These results indicate that the model successfully captures the global structured information from data on the first stage with the latent variable **u** and then adds a local structure with the help of the latent variable **z**, whose responsibility is to capture the missing information between the compressed and the original data.



Figure 8: Reconstruction results on ImageNet<sup>†</sup> from the VAE with the bijective prior. The results show that the model can successfully reconstruct natural images from a different source though its  $8 \times 8 \times 16$  dimensional latent space. Top row indicates the ground truth samples and the second rows present the results after reconstruction.



Figure 9: Super-Resolution results of the srVAE on ImageNet<sup>†</sup>. Even though the model was trained on CIFAR10, its performance showcases its robustness capabilities. The top, second and third row illustrate the  $2 \times$  downscaled image, the original and the after the super resolution result of the model, respectively.



Figure 10: Comparison on the image reconstruction on CIFAR10 between the VAE with the bijective prior (second row) and the srVAE (third row). Even though both models output images very similar to the original one (top row), the srVAE seems to preserve more details from the ground truth (take a look at images 5, 10 and 15).



Figure 11: Comparison on the image reconstruction on ImageNet<sup> $\dagger$ </sup> between the VAE with the bijective prior (second row) and the srVAE (third row). Again, we can notice that the last row seems to be more similar to the ground truth samples (top row).