



HAL
open science

Variational Auto-Encoder: not all failures are equal

Victor Berger, Michèle Sebag

► **To cite this version:**

Victor Berger, Michèle Sebag. Variational Auto-Encoder: not all failures are equal. 2020. hal-02497248

HAL Id: hal-02497248

<https://inria.hal.science/hal-02497248v1>

Preprint submitted on 3 Mar 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Variational Auto-Encoders: Not all failures are equal

Victor Berger¹ and Michele Sebag¹

¹TAU, CNRS – INRIA – Univ. Paris-Saclay, France

Abstract

We claim that a source of severe failures for Variational Auto-Encoders is the choice of the distribution class used for the observation model. A first theoretical and experimental contribution of the paper is to establish that even in the large sample limit with arbitrarily powerful neural architectures and latent space, the VAE fails if the sharpness of the distribution class does not match the scale of the data. Our second claim is that the distribution sharpness must preferably be learned by the VAE (as opposed to, fixed and optimized offline): Autonomously adjusting this sharpness allows the VAE to dynamically control the trade-off between the optimization of the reconstruction loss and the latent compression. A second empirical contribution is to show how the control of this trade-off is instrumental in escaping poor local optima, akin a simulated annealing schedule. Both claims are backed upon experiments on artificial data, MNIST and CelebA, showing how sharpness learning addresses the notorious VAE blurriness issue.

1 Introduction

A huge leap forward to the identification of generative models from data, Variational Auto-Encoders (VAEs) [12, 23] can be analysed as tackling two complementary problems: i) identifying the data manifold (Pb (1)); ii) identifying a distribution on this manifold that actually matches the data distribution (Pb (2)). The VAE training criterion tackles both problems through the joint optimization of a data fitting term and a regularization term enforcing the smoothness of the generative model w.r.t.

the chosen distribution class (section 2).

The first claim of the paper is that Pb (1) is the most severe bottleneck for the VAE: if the data manifold is not properly identified, the VAE fails to learn any relevant probability distribution. This failure can be due to several reasons, e.g. insufficiently powerful neural architectures or reconstruction losses [14, 25, 10]. We posit that a necessary condition for a successful VAE training is an appropriate choice of the distribution class of the observation model, and that an inappropriate choice *cannot be compensated for*. This claim is backed upon theoretical and experimental results on a simple artificial problem,¹ i) showing that even in the large sample limit with an infinite representation power of the neural architecture, a VAE does fail if the sharpness of the distribution class does not match the data scale; and ii) analytically characterizing the mismatch between the distribution sharpness and the data scale.

The second claim of the paper is that a most advantageous option is to let the VAE learn the appropriate sharpness of the observation model, defining the L-VAE setting. Empirical evidence shows that a L-VAE – considering a flexible distribution class p_θ and learning the best hyper-parameter θ^* thereof – behaves significantly better than a VAE considering the fixed distribution class p_{θ^*} . The inspection of the training dynamics shows that the autonomous adjustment of the sharpness allows the L-VAE to control the trade-off between the data fitting and the regularization terms. In contrast, considering a distribution with a fixed though proper sharpness makes the data fitting term dominate the optimization in the early learning phase, engaging the VAE

¹ Quoting Ali Rahimi (2017): *Simple experiments, simple theorems are the building blocks that help us understand more complicated systems.*

in unpromising regions w.r.t. the optimization of the regularization objective. These observations, backed upon experiments on artificial data, MNIST and CelebA, are interpreted in terms of a simulated annealing process, as analyzed in [9] w.r.t. the dimensionality of a neural net. The lesson learned is that the learning process should better operate in a larger space than the one containing the eventual solution, for the sake of a less constrained and more efficient optimization process.

The paper is organized as follows. Section 2 introduces the state of the art. Section 3 presents a theoretical analysis of the impact of the observation model, and describes the L-VAE setting. Section 4 reports and discusses empirical evidence of the merits of L-VAE. The concluding remarks present some perspectives for further research.

2 Formal Background and Related Work

Now a classical approach for generative modelling, Variational Auto-Encoders [23, 11] search for a generative distribution $p_\theta(x)$, expressed w.r.t. some latent variable z :

$$p_\theta(x) = \int p_\theta(x|z)p_\theta(z)dz \quad (1)$$

The encoder module yields the parameters of distribution $q(z|x)$, defined on the latent space and used to sample z . The decoder module yields the parameters of distribution $p_\theta(x|z)$, defined on the input space and used to generate (ideally) a new sample of the data distribution. The VAE loss is composed of a data fitting term (maximizing the data likelihood) and a latent compression term (enforcing that $q(z|x)$ stay close to the latent prior $p_\theta(z)$):

$$\mathcal{L}_{VAE}(x) = \mathbb{E}_{z \sim q} -\log p_\theta(x|z) + D_{KL}(q(z|x)||p_\theta(z)) \quad (2)$$

The data fitting term is optimized by taking advantage of the Evidence Lower Bound (ELBO), stating that for any distribution $q(z|x)$:

$$\log p_\theta(x) \geq \mathbb{E}_{z \sim q(z|x)} \log \frac{p_\theta(x|z)p_\theta(z)}{q(z|x)} \quad (3)$$

The minimization of Eq. (2) enables the end-to-end training of the VAE (Fig. 1), specifically the encoder

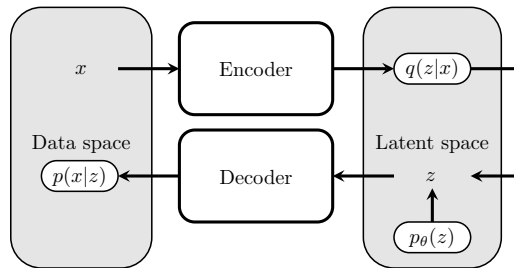


Figure 1: VAE architecture: the encoder (resp. decoder) module is a neural network taking x (resp z) as input and computing the parameters of distribution $q(z|x)$ (resp., $p_\theta(x|z)$).

$q(z|x)$ (a.k.a. inference model) and decoder $p_\theta(x|z)$ (a.k.a. observation model), both implemented as neural networks. Both encoder and decoder take a value (x or z) as input and produce the parameters of a distribution ($q(z|x)$ or $p_\theta(x|z)$). The mainstream class of distribution chosen for the inference model is a Normal distribution, the mean and variance of which are computed by the encoder along the known re-parametrization trick, $q(z|x) = \mathcal{N}(\mu(x), \sigma(x))$ with $\mu(x)$ and $\sigma(x)$ provided by the encoder module.

The maximization of the data likelihood (Eq. (1)) does not prevent the observation model from generating out-of-distribution samples [26]. In practice, VAEs are observed to generate unrealistic samples more often than Generative Adversarial Networks [21], prompting the introduction of adversarial losses in VAEs (see for instance [14, 3]).

Inference model. The role of the inference model is to enforce the quality of the lower bound (Eq. (3)). In the (ideal) equality case, $q(z|x) = p_\theta(z|x)$, that is, the inference model is a probabilistic inverse of the observation model relative to the prior $p_\theta(z)$. In such a case, the dataset would be perfectly represented (compressed) in the latent space, with $q(z|x)$ and $p_\theta(x|z)$ potentially being very complex maps. When the inference model is poor, the ELBO is quite loose, the input space is mapped onto a poorly compressed latent space, increasing the risk of generating out-of-distribution samples.

As said, the inference model is classically built upon a Normal distribution through the reparametrization trick, while prior model $p_\theta(z)$ is set to $\mathcal{N}(0; \mathbf{I})$. The limitations of this setting are related to its smoothness, that might not reflect

the data, opening mostly two directions for VAE extensions.

In the case where the dataset is structured, e.g. involving disjoint clusters, the problem is that Gaussian encoders cannot build sharp boundaries in the latent space. A natural option thus is to also create clusters in the latent space to ensure a good reconstruction, via a more complex prior. More complex distribution classes, e.g. involving categorical variables, are considered to handle structures in the dataset [12, 13], as well as auto-regressive priors [20, 27, 22].

Another direction for VAE extensions stems from the following remark. As the decoder defines a continuous map from the latent z to the data x , it is necessary that the latent space topology jointly defined by $q(z|x)$ and $p(z)$ matches the dataset topology. This remark calls for using hierarchical latent structures [25] or exploring non-Euclidean latent spaces [19, 18].

Observation model. Besides a good inference model $q(x|z)$, the VAE success also requires a good observation model $p_\theta(x|z)$, efficiently untangling a compressed $q(x|z)$ and mapping it back onto the data space.

The choice of the distribution class for the observation model governs the reconstruction loss ($-\log p_\theta(x|z)$), with a significant impact on the quality of the generated samples. Typically, the use of a Normal distribution class with fixed variance boils down to using a mean square error in the input space. Such a model is clearly inappropriate in the domain of computer vision, as setting an independent noise on the image pixels hardly yields a realistic image.

Two main research directions have been investigated to improve the observation model while keeping computational complexity under control. The former one leverages domain knowledge to augment the VAE criterion, e.g. using perceptual metrics [14, 3]. The latter one considers significantly more complex distribution classes [5, 24], e.g. involving autoregressive models over the pixels. A difficulty with such models is that their representation power is sufficient to learn the data distribution while ignoring the latent space. This phenomenon, referred to as posterior collapse [6, 1], is blamed (in the linear VAE case) on the presence of spurious optima

in the log marginal likelihood [17].

3 Impact of the observation model

The claim that the most severe cause of VAE failures is a bad approximation of the data manifold (Pb (1)) is argued as follows. On the one hand, the samples generated by the VAE can hardly be of better quality than the samples reconstructed from the data; and the reconstructed samples can only be as good as the manifold approximation. On the other hand, the approximation of the data distribution on the manifold (Pb (2)) is directly addressed by the ELBO criterion (Eq (3)), accounting for the diverse modes of the data; actually, VAEs do not suffer from the notorious mode-dropping phenomenon observed in GANs [2].

As said, this paper focuses on one among the possible causes for the bad approximation of the data manifold, namely the choice of the distribution family p_θ . This section first presents a case of non-identifiability of the data manifold, only due to the choice of the distribution family p_θ . Some approaches used in the literature to avoid this failure case are discussed and an alternative is proposed.

3.1 A non-identifiability case

Let us consider the usual choice of a Normal distribution, where the observation model $p(x|z)$ is defined as $\mathcal{N}(\mu(z), \sigma^2 \mathbf{I})$ with $\mu(z)$ the output of the decoder network and σ a fixed scalar variance. While in this setting the reconstruction term conveniently boils down to the squared error loss, the choice of σ has a dramatic impact on the identification of the data manifold. Let us consider a dataset located on the hyper-sphere of radius R in dimension D . The following result establishes that a VAE cannot capture the data manifold for large values of hyper-parameter σ :

Theorem 1. *Let us assume arbitrarily powerful neural architectures for the encoder and decoder networks, as well as an arbitrarily powerful latent space. Let us further assume that the data is an infinite sample size of the D -dimensional hyper-sphere of radius R noted $\mathbf{S}_{D-1}(R)$.*

Then, a VAE with observation model $\mathcal{N}(\cdot, \sigma^2 \mathbf{I})$ fails to characterize the data manifold $\mathbf{S}_{D-1}(R)$ if

$$\sigma \geq \frac{R}{\sqrt{D-1}}$$

Proof. Let $\mu(z)$ denote the output of the decoder network of the VAE. By definition, $p_\theta(x|z) = \mathcal{N}(\mu(z), \sigma^2 \mathbf{I})$. Under the large sample and arbitrarily powerful neural architecture assumptions, the VAE can exactly characterize $q(z|x)$ and $p(z)$. Distribution $p(z)$ can be replaced by the distribution over the (deterministic) decoder output $p(\mu)$. The ELBO equation therefore becomes an equality and $p_\theta(x)$ can be computed exactly:

$$p_\theta(x) = \int p(x|\mu)p_\theta(\mu)d\mu \quad (4)$$

with

$$p(x|\mu) = \frac{1}{(\sqrt{2\pi}\sigma)^D} \exp\left(-\frac{1}{2\sigma^2}\|x - \mu\|^2\right) \quad (5)$$

Under the assumptions made, the optimal distribution $p_\theta(\mu)$ is derived as follows. The data being radially symmetrical, the resulting distribution will necessarily be so as well. Distribution $p_\theta(\mu)$ thus only depends on the distance r from μ to the center of the sphere. Furthermore, $p_\theta(x)$ being the expectation of function $p(x|\mu)$ under $p_\theta(\mu)$, it is maximized when $p_\theta(\mu)$ has all its mass where $p(x|\mu)$ reaches its maximum. Accordingly, $p_\theta(\mu)$ ranges among the uniform distributions over some $S_{D-1}(r)$, hyper-sphere of some radius r in dimension D with same center as the dataset. Taking normalization into account, it comes:

$$p_\theta(\mu) = \frac{1}{r^{D-1}A_{D-1}} \mathbf{1}_{(\|\mu\|=r)} \quad (6)$$

with A_{D-1} the area of a unit D -dimensional sphere. Using spherical coordinates,

$$\begin{aligned} \|x - \mu\|^2 &= R^2 + r^2 - 2rR \cos(\phi) \\ d\mu &= A_{D-2}r^{D-1} \sin^{D-2}(\phi)drd\phi \end{aligned} \quad (7)$$

with ϕ the angle between x and μ wrt the center of the sphere and A_{D-2} the area of the unit hyper-sphere in dimension $D-2$ (over which the integration is trivial as it has no impact on other terms).

Removing constant terms wrt to r and ϕ it comes:

$$p_\theta(x) \propto e^{-\frac{r^2}{2\sigma^2}} \int_{\phi=0}^{\pi} \sin^{D-2}(\phi) \exp\left(\frac{rR}{\sigma^2} \cos(\phi)\right) d\phi \quad (8)$$

It is easily shown that the derivative of $p_\theta(x)$ has same sign as:

$$\int_0^\pi \frac{r}{\sigma^2} \left(\frac{R^2 \sin^2(\phi)}{(D-1)\sigma^2} - 1 \right) \sin^{D-2}(\phi) e^{rR \cos(\phi)} d\phi \quad (9)$$

Accordingly, an extremum is reached for $r = 0$, and $r = 0$ is the global maximum for $\frac{R^2}{(D-1)\sigma^2} \leq 1$. Therefore, if the radius R of the data hyper-sphere is small compared to standard deviation σ ($R \leq \sigma\sqrt{D-1}$), in the most favorable case of large sample limit and arbitrary power of the neural architectures, the VAE can but approximate the data manifold by a normal distribution with same center as the dataset hypersphere. In other words, it grossly characterizes the volume of the hyper-sphere instead of its area. \square

Inspecting the second order derivative of Eq (8) shows that while $r = 0$ is a local optimum, it is a local maximum for $R < \sigma\sqrt{D}$, and a local minimum for larger values of R . In the latter case, the optimal value of r asymptotically converges toward $R = r$. In other words, the generative model $p_\theta(x)$ is expressed as a mixture of Normal distributions, the mean of which is located on the data manifold: the generative model is accurately paving the data manifold.

For $D = 2$, Eq. (8) can be computed analytically; its heat map depending on $\frac{R}{\sigma}$ and $\frac{r}{\sigma}$ is depicted in Fig. 2. For small values of $\frac{R}{\sigma}$, $r = 0$ is the only global maximum. As $\frac{R}{\sigma}$ increases, the maximum moves toward the $r = R$ diagonal; graphically, this change occurs for $\frac{R}{\sigma} \approx \sqrt{2}$.

3.2 External control of the model sharpness

Most generally, the sharpness of the distribution class governs the information extracted from the data. For a given fixed sharpness, the observation model is bound to ignore all details that are small relatively to this sharpness: the gain in terms of reconstruction loss is smaller than the cost incurred by encoding these details in the latent variables,

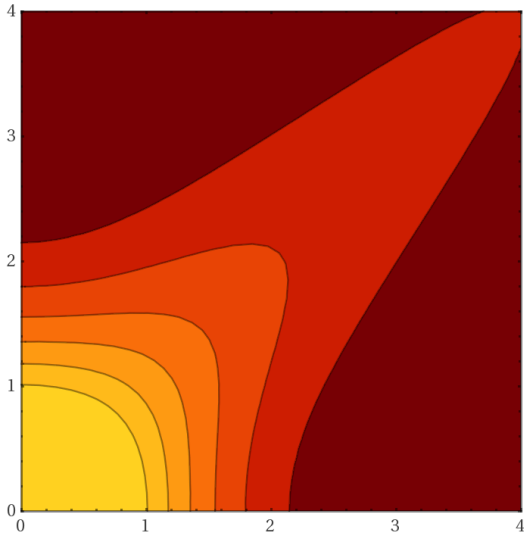


Figure 2: Heat map of $p_\theta(x)$ vs $\frac{R}{\sigma}$ (horizontal axis) and $\frac{z}{\sigma}$ (vertical axis) for $D=2$ (the lighter the higher; better seen in color). See comments in text.

regardless of the data amount and of the representation power of the neural encoder and decoder.

The sharpness impact is often hidden by the fact that generated samples are derived as the expectation of distribution $p_\theta(x|z)$, particularly so in the case of computer vision [14, 3, 8]: the generated sample simply is the raw output of the decoder, that is the mean² of $p_\theta(x|z)$. This generation process indeed mitigates the impact of large σ values; still, it does not address the potential mis-specifications of the observation model. In the particular case of images, the squared error loss corresponds to a Normal distribution with variance $\sigma^2 = \frac{1}{2}$. This very large variance (considering that pixels are usually normalized in the $[0, 1]$ interval) causes most data details to be discarded, explaining the blurriness of the generated images.

Most generally, the VAE learning criterion aims at a generative process sampling (as opposed to, averaging) $p_\theta(x|z)$. The averaging strategy thus creates some discrepancy between the well-founded VAE design and its actual usage.

Another strategy is to reconsider the VAE training loss itself and design a more effective one, be it derived from a principled probabilistic setting or

²The interpretation is more complex for other losses such as the Bernoulli loss [16].

not. The simplest option is to control the trade-off between the reconstruction loss and the latent compression by means of some penalization weight λ on the latent term. In the general case, increasing the weight of the latent term cannot be rigorously interpreted in terms of the sharpness of the observation model due to normalization issues,³ but informally speaking, introducing a λ term in front of the latent term of the loss is similar to raising the observation model distribution at the power $\frac{1}{\lambda}$: depending on the value of λ the model might capture finer-grained details of the data and increase the latent loss ($\lambda < 1$); or discard the details and achieve a higher compression of the latent variables ($\lambda > 1$). In the case of a Normal distribution with fixed variance σ^2 , using a penalization weight λ on the latent loss exactly corresponds to multiplying σ by $\sqrt{\lambda}$.

3.3 Learning the model sharpness

As the VAE success depends on a good approximation of the data manifold, and the quality of the approximation depends in part on the observation model, one possibility is to allow this model to be arbitrarily sharp, e.g. by making the variance of the Normal distribution a learned parameter, thus defining the L-VAE setting.⁴ The reconstruction term of the L-VAE loss then reads (up to an additive constant), with D the dimension of the input space:

$$-\log p_\theta(x|z) = \frac{1}{2\sigma^2} \|x - \mu(z)\|^2 + D \log \sigma \quad (10)$$

Maximizing the likelihood of the generated samples based on Eq. (10) allows the L-VAE to autonomously adjust the trade-off between the reconstruction and the latent losses (as opposed to, manually setting a fixed penalization weight on the latent term). As will be seen (section 4), the desirable balance between both terms evolves along the learning trajectory, making all the more important to let the VAE control the trade-off.

Normal distributions with learned σ permit in principle an arbitrarily accurate approximation of

³The importance of taking into account such normalization issues, and the impact of neglecting those are discussed in [16].

⁴This setting must be distinguished from the well-known re-parameterization trick [4], operating on the latent distribution $q(z|x)$.

the manifold, by selecting arbitrarily small values of σ , while the $D \log \sigma$ term (Eq (10)) is an incentive to choose σ as small as possible, to the extent permitted by the rest of the model.

A straightforward refinement of the above isotropic Normal distribution is to consider a diagonal covariance matrix $\mathcal{N}(\mu(z), \sigma(z))$, with σ a D -dimensional vector computed by the decoder akin μ . The associated training loss reads:

$$-\log p_{\theta}(x|z) = \sum_{i=1}^D \frac{1}{2\sigma_i(z)^2} (x_i - \mu_i(z))^2 + \log \sigma_i(z) \quad (11)$$

This refinement might still be insufficient when the data manifold is not aligned with the axes of the input space data, requiring a Normal distribution with full covariance matrix to accurately describe the manifold. Learning the full covariance matrix however raises strong scalability issues, and only the scalar and vectorial L-VAE settings (Eqs. (10) and (11)) are considered in the following.

4 Inspecting the VAE dynamics

This section investigates in more depth the impact of the observation model on the generative process. An artificial dataset, used to provide empirical evidence and backup our first claim, enables to inspect the VAE learning dynamics. The main lesson learned, and the second claim of the paper, is that learning the sharpness of the observation model allows the VAE to escape bad local optima of the optimization landscape. The claims and the proposed interpretation of the results on the artificial dataset are confirmed by experiments on MNIST and CelebA [15].

4.1 Experimental setting

The main goal of experiments is to assess the impact of the observation model and its sharpness, be it fixed or learned, on the reconstruction and generation abilities of a VAE.

An artificial problem is designed in \mathbf{R}^2 (Fig. 3): data are located on a 1D manifold, structured in three regions with moderate, small and large noise. This design aims to independently assess whether

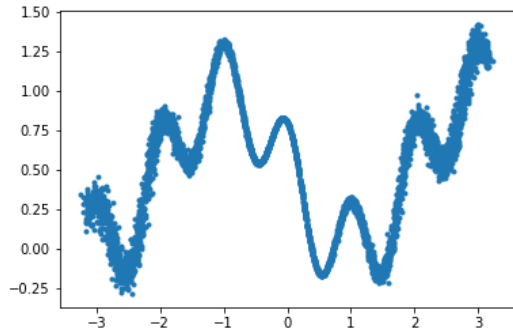


Figure 3: Artificial dataset in \mathbf{R}^2 , located on a 1D manifold with heteroscedastic noise.

the considered VAE settings enable to identify the manifold and its fine-grained details.

The neural architecture used for both the encoder and the decoder is a ResNet [7] with 4 residual blocks of 200 neurons. Considering the data complexity, this architecture is expected to offer a "sufficient" representation capacity (satisfying the assumptions of Thm 1), and make the identification of the data manifold the only bottleneck for VAE. Likewise, the considered dataset satisfies the large sample limit assumption, with 100 points sampled anew from the data distribution in each batch.

Several VAE settings are experimentally compared, only differing in the observation model $\mathcal{N}(\cdot, \sigma)$, with no penalization weight on the latent term:

- i) F-VAE considers a fixed scalar σ , ranging in $\{1.0, 0.3, 0.1, 0.03, 0.01\}$; ii) LS-VAE considers a learned scalar σ (initialized as $\sigma = 1$); iii) LV-VAE considers a learned vector $\sigma(z)$, as an output of the decoder network; iv) finally, F*-VAE considers a fixed scalar σ set to the final value learned by LS-VAE.

The first performance indicator is the ELBO value⁵, quantitatively assessing the generative model. Two other qualitative performance indicators are proposed: The reconstruction performance is qualitatively assessed by displaying the reconstructed samples (the mean of $p_{\theta}(x|z)$) on the top of the initial samples; a perfect reconstruction is characterized by the fact that the reconstructed

⁵ The ELBO values reported for F-VAE are augmented with a $2 \log \sigma$ term, for a fair comparison with LS-VAE and LV-VAE, after Eqs (10) and (11).

Setting	σ	ELBO
F-VAE, Fixed σ	$\sigma = 1.0$	-1.14 ± 0.02
	$\sigma = 0.3$	-0.063 ± 0.03
	$\sigma = 0.1$	0.773 ± 0.09
	$\sigma = 0.043$	0.335 ± 0.08
	$\sigma = 0.03$	0.141 ± 0.10
	$\sigma = 0.01$	-0.961 ± 0.12
LS-VAE, scalar σ	$\sigma^* = 0.043$	1.26 ± 0.09
LV-VAE, vector σ		1.62 ± 0.08

Table 1: ELBO values (averaged over 10 runs) of F-VAE, LS-VAE and LV-VAE; the higher the better.

samples exactly hide the initial ones. Likewise, the generative performance is qualitatively assessed by comparing the generated samples with the initial dataset.

4.2 Results and analysis

Table 1 displays the ELBO indicator for all considered VAE settings (averaged on 10 runs). These results show the decisive impact of learning the variance of the observation model: the option significantly dominating all others is LV-VAE, learning a vector σ ; the second best is LS-VAE, learning a scalar σ .

Most interestingly, when fixing σ to the optimal scalar value found by LS-VAE ($\sigma = .043$) the ELBO obtained by F*-VAE is significantly degraded compared to LS-VAE; the ELBO of F*-VAE also is significantly lower than for F-VAE with $\sigma = 0.1$. These results suggest that the dynamics of the learning process matters as much as the actual sharpness σ of the observation model.

A more detailed analysis is permitted by the qualitative performance indicators, visually displayed in Fig. 4. Regarding F-VAE, the reconstruction is poor for high values of σ (Fig. 4.(a), Left; all details are lost, following the analysis in section 3); as σ decreases, the reconstruction gradually improves and the reconstructed samples exactly hide the initial ones for $\sigma \leq .03$ (at the expense of the compression and generation quality). Note that for medium values of σ , the reconstructed samples are closer to the manifold than the initial ones, particularly so in the regions with high or moderate noise (Fig. 4.(b), Left). This phenomenon is likewise explained by the fact that the VAE discards details that are

small comparatively to the model sharpness (section 3): when the data noise is smaller than σ (e.g. for $\sigma = .1$), the noise is discarded, i.e. the reconstruction operates like a denoiser.

Regarding F-VAE, the generation is poor in all cases but for $\sigma = .1$ (Fig. 4.(a-c), Right). For high values of σ , the details are missed, with a roughly cosine-like shape generated for $\sigma = .3$. For small values of σ , the generated samples also form a rough shape with many samples far from the manifold, all the more so as σ decreases. Despite the blurriness of the generated samples, their average does accurately capture the data manifold. This situation is similar to that observed on, e.g., the CelebA dataset, where the generated samples are seen as random noise, while their average rightly resembles a face.

Regarding L-VAE, in both scalar and vector cases the reconstruction is good although the denoising effect is observed (reconstructed samples are closer to the manifold), which is attributed to the comparatively large value of the learned σ (Fig. 4.(d),(f), Left). In both scalar and vector cases, the generation is quite good compared to that of F-VAE at its best. In the scalar case, the generated samples are located on a single-width band centered on the manifold: the observation model does not have the representation power needed to fit the varying data noise. In the vector case, the observation model allows to encode both the manifold and its thickness, and the generated samples exactly match the initial samples (perfect generation).

In order to disentangle the impacts of the model sharpness and of the optimization process, additional experiments are conducted with a fixed- σ VAE, where σ is set to the value learned by FS-VAE ($\sigma = .043$). The reconstruction and generation results are displayed on Fig. 4.(e). The reconstruction is decent, with no denoising effect as σ is sufficiently small to capture the details. The generation is much poorer than for L-VAE, and even poorer than for F-VAE with $\sigma = .1$, with the generated samples located on a thick band centered on the manifold.

This comparison suggests that besides the sharpness of the observation model, the second key issue for the VAE performance is the trade-off between the reconstruction and the latent losses, and how this trade-off is dynamically controlled along the learning trajectory. It is observed that FS-VAE decreases the learned σ value by jumping from time to time to lower σ plateaus, and these jumps coincide

with leaps of the ELBO (more in appendix).

The dynamics of the learned σ along the VAE training can be interpreted in terms of an annealing process, taking inspiration from the analysis of [9]. In the early learning stage, the large σ value compensates for the large reconstruction loss and lowers the energy landscape. At this stage, the latent compression term of the loss dominates, and the model only aims to match the global shape of the data manifold. As the reconstruction at the large scale improves, the model can then afford to lower σ and improve the overall loss. This has the effect of gradually making the energy landscape steeper, raising more details of the data to the attention of the model, and allowing σ to be again lowered as the reconstruction loss decreases. The process continues until any improvement of the reconstruction term would be cancelled out by a degradation of the latent term. At this point σ stabilizes and the self-controlled annealing ends. This gradual evolution of the energy landscape allows the VAE to iteratively build and refine its latent representation, eventually yielding a better compression than the one that would be reached by running the optimization in a fixed energy landscape.

Complementary experiments (in appendix) show that an externally set σ schedule yields the same results (quantitative ELBO, visual reconstruction and generation) as the LS-VAE, confirming the importance of a dynamic schedule to adjust the reconstruction *vs* latent optimization trade-off.

4.3 Image datasets

The significance of the above discussion is examined by comparing the F-VAE and LV-VAE settings on the MNIST and CelebA datasets.

Note that a usual practice in computer vision is to set the reconstruction loss to the mean square error, which is equivalent to considering a fixed Normal distribution with variance $\sigma^2 = 1/2$. As said, this large variance is expected to prevent the reconstruction from capturing the fine-grained details and to cause generated samples to look like random ones (all things considered, the situation is similar to that of $\sigma = .3$ in the 1D artificial problem). The generative process thus can only fall back on the average of the generated samples, that is realistic although blurry (Fig. 5.(a-b), Left).

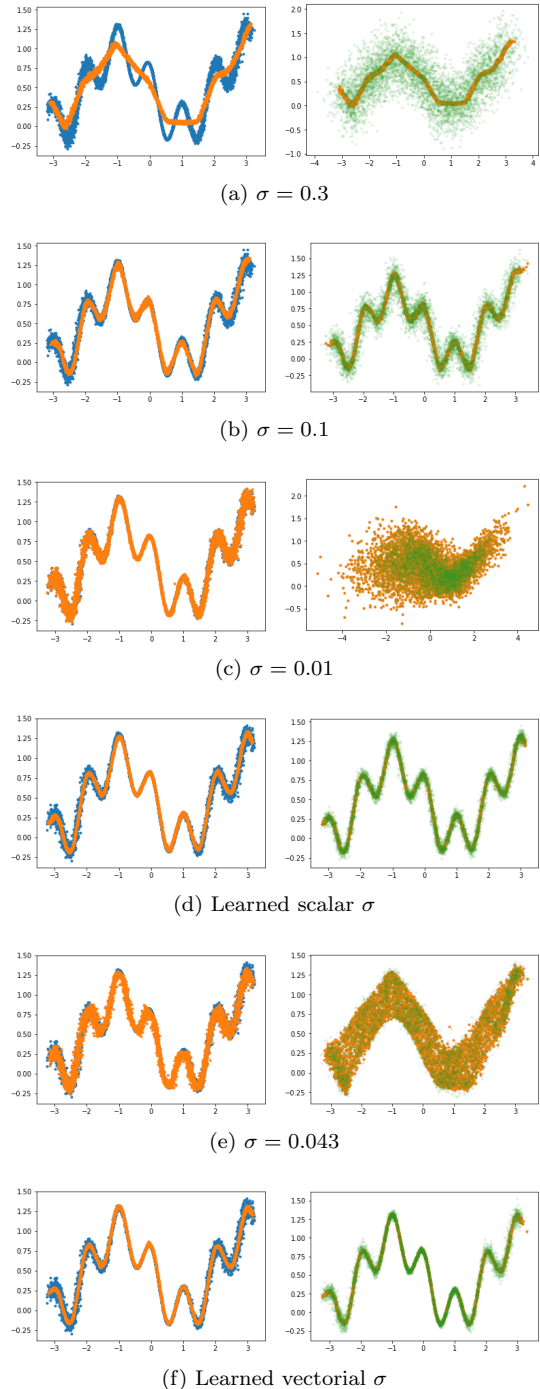
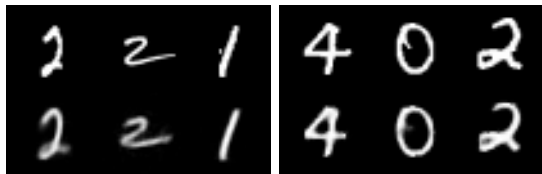


Figure 4: (a)-(c): F-VAEs performance with fixed σ ranging from 1 to .01 from top to bottom ; (d): LS-VAE; (e): F-VAE with same σ value as LS-VAE; (f): LV-VAE (better seen in color).

Left: Reconstructed samples (in orange) and initial samples (in blue). Right: Generated samples (in green) and means of $p_\theta(x|z)$ (in orange).

Dataset	F-VAE	LV-VAE
MNIST	0.06	4.89
CelebA	5.52	57.51

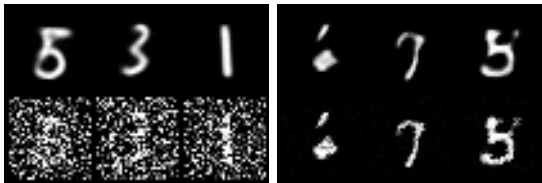
Table 2: Comparison of F-VAE (mean square error) and LV-VAE on MNIST and CelebA: ELBO (in bits per pixel, accounting for the $\log \sigma$ term for a fair comparison).



(a) MNIST reconstruction



(b) CelebA reconstruction



(c) MNIST generation



(d) CelebA generation

Figure 5: Reconstruction and generation on MNIST and CelebA.

(a)-(b): Reconstruction; Top: initial images. Bottom Left: F-VAE (MSE loss); Bottom Right: LV-VAE (Eq. 11).

(c)-(d): Generation; Left: F-VAE (MSE loss), Right: LV-VAE (learned vectorial σ). Means of $p_\theta(x|z)$ are shown in the top rows; samples are shown in the bottom rows.

The F-VAE and LV-VAE approaches are compared in terms of reconstruction on Fig. 5.(a-b). As LV-VAE controls the observation sharpness, it achieves a quasi pixel-perfect reconstruction (Fig. 5.(a-b), Right), much better than F-VAE. The quantitative ELBO indicator (Table 2) confirms that F-VAE is dominated by LV-VAE by about one order of magnitude.⁶

Regarding the generation performance, the samples $p_\theta(x|z)$ generated by F-VAE are viewed as random noise (Fig. 5.(c-d), Left, bottom), but their averages (the mean of $p_\theta(x|z)$) are quite similar to the inverse image of z , suggesting that while the VAE identifies a much simplified manifold, its latent space manages to efficiently compress this manifold (Fig. 5.(c-d), Left, top). Quite the contrary, the samples generated by LV-VAE (Fig. 5.(c-d), Right, bottom) are much better than random images; the precision is similar to that of reconstructed images. Henceforth, the samples and their averages are very similar; unfortunately they are less convincing than the averaged samples generated by F-VAE (some of the MNIST digits are distorted, and the CelebA faces lack structure). A tentative interpretation for this weakness is as follows. As LV-VAE approximates the data manifold with a better accuracy, more details need be represented, making the latent space more difficult to compress. Eventually the LV-VAE faces another bottleneck: the insufficient power of representation of the latent space.⁷

5 Conclusion and Perspectives

The main contributions of the paper is to shed more light on the VAE bottlenecks, to propose some principles in order to address these bottlenecks, and to illustrate the practical efficiency thereof.

Our first claim is that the primary VAE bottleneck lies in the identification of the data manifold. In particular, an observation model unable to pave the data manifold is bound to grossly approximate the data and discard important information, irre-

⁶As in Table 1, the ELBO values reported for F-VAE are augmented with a $\log \sigma$ term for the sake of a fair comparison.

⁷This representation power can be limited by both the capacity of the encoder and decoder neural networks, as well as the probabilistic structure chosen for the latent variable z , here set to the traditional factorized Gaussian.

spective of the expressive power of the other VAE elements.

Our second claim is that a good identification of the manifold can be achieved in a theoretically sound and algorithmically efficient way: the proposed L-VAE learns the sharpness of the observation model, thereby dynamically adjusting the optimization trade-off between the reconstruction and latent terms akin a simulated annealing schedule, ending up with a richer and better compressed latent space than allowed with a fixed sharpness. The L-VAE achieves a quasi pixel-perfect reconstruction, and removes the generation noise associated with a too large fixed variance. In counterpart, as the L-VAE considers more information, it exerts more pressure on the latent space and its compression capabilities, decreasing the visual realism of the raw output of the decoder (mean of $p_{\theta}(x|z)$).

A first perspective for further research is to revisit the diverse latent space structures and/or neural encoder and decoder architectures proposed in the literature, and to assess them *when coupled with a good identification of the data manifold*. The rationale is that a poor approximation of the data manifold might cause compound failures of the overall VAE, obfuscating the true impact of new latent spaces and neural architectures.

A second research perspective is to design observation models better tailored to the specifics of the dataset, able to identify the local tangent to the data manifold though with a better scalability than based on a full-rank covariance matrix.

References

- [1] Alexander A. Alemi, Ben Poole, Ian Fischer, Joshua V. Dillon, Rif A. Saurous, and Kevin Murphy. Fixing a Broken ELBO. In *ICML*, 2017.
- [2] Martin Arjovsky and Léon Bottou. Towards Principled Methods for Training Generative Adversarial Networks. *International Conference on Learning Representations*, (2017), January 2017. arXiv: 1701.04862.
- [3] Alexey Dosovitskiy and Thomas Brox. Generating Images with Perceptual Similarity Metrics based on Deep Networks. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems 29*, pages 658–666. Curran Associates, Inc., 2016.
- [4] 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. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 2672–2680. Curran Associates, Inc., 2014.
- [5] Ishaan Gulrajani, Kundan Kumar, Faruk Ahmed, Adrien Ali Taïga, Francesco Visin, David Vázquez, and Aaron C. Courville. PixelVAE: A Latent Variable Model for Natural Images. *ArXiv*, abs/1611.05013, 2016.
- [6] Junxian He, Daniel Spokoyny, Graham Neubig, and Taylor Berg-Kirkpatrick. Lagging Inference Networks and Posterior Collapse in Variational Autoencoders. *ArXiv*, abs/1901.05534, 2019.
- [7] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, June 2016. ISSN: 1063-6919.
- [8] Irina Higgins, Loic Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick, Shakir Mohamed, and Alexander Lerchner. beta-VAE: Learning Basic Visual Concepts with a Constrained Variational Framework. *ICLR*, 2(5):6, 2017.
- [9] Chin-Wei Huang, Shawn Tan, Alexandre Lacoste, and Aaron C Courville. Improving Explorability in Variational Inference with Annealed Variational Objectives. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31*, pages 9701–9711. Curran Associates, Inc., 2018.
- [10] Huaibo Huang, zhihang li, Ran He, Zhenan Sun, and Tieniu Tan. IntroVAE: Introspective Variational Autoencoders for Photographic

- Image Synthesis. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31*, pages 52–63. Curran Associates, Inc., 2018.
- [11] Diederik P. Kingma and Max Welling. Auto-Encoding Variational Bayes. *International Conference on Learning Representations*, (2014), May 2014. arXiv: 1312.6114.
- [12] Durk P Kingma, Shakir Mohamed, Danilo Jimenez Rezende, and Max Welling. Semi-supervised Learning with Deep Generative Models. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 3581–3589. Curran Associates, Inc., 2014.
- [13] Durk P Kingma, Tim Salimans, Rafal Jozefowicz, Xi Chen, Ilya Sutskever, and Max Welling. Improved Variational Inference with Inverse Autoregressive Flow. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems 29*, pages 4743–4751. Curran Associates, Inc., 2016.
- [14] Anders Boesen Lindbo Larsen, Søren Kaae Sønderby, Hugo Larochelle, and Ole Winther. Autoencoding Beyond Pixels Using a Learned Similarity Metric. In *Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48*, ICML’16, pages 1558–1566. JMLR.org, 2016. event-place: New York, NY, USA.
- [15] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep Learning Face Attributes in the Wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015.
- [16] Gabriel Loaiza-Ganem and John P. Cunningham. The continuous Bernoulli: fixing a pervasive error in variational autoencoders. *Conference on Neural Information Processing Systems*, (2019), October 2019. arXiv: 1907.06845.
- [17] James Lucas, George Tucker, Roger B Grosse, and Mohammad Norouzi. Don’t Blame the ELBO! A Linear VAE Perspective on Posterior Collapse. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d\textquotesingle Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 9403–9413. Curran Associates, Inc., 2019.
- [18] Emile Mathieu, Charline Le Lan, Chris J. Maddison, Ryota Tomioka, and Yee Whye Teh. Continuous Hierarchical Representations with Poincaré Variational Auto-Encoders. *arXiv:1901.06033 [cs, stat]*, November 2019. arXiv: 1901.06033.
- [19] Yoshihiro Nagano, Shoichiro Yamaguchi, Yasuhiro Fujita, and Masanori Koyama. A Wrapped Normal Distribution on Hyperbolic Space for Gradient-Based Learning. In *ICML*, 2019.
- [20] Eric Nalisnick and Padhraic Smyth. Stick-breaking variational autoencoders. In *International Conference on Learning Representations (ICLR)*, 2017.
- [21] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks. *CoRR*, abs/1511.06434, 2015.
- [22] Ali Razavi, Aaron van den Oord, and Oriol Vinyals. Generating Diverse High-Fidelity Images with VQ-VAE-2. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d\textquotesingle Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 14837–14847. Curran Associates, Inc., 2019.
- [23] Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic Backpropagation and Approximate Inference in Deep Generative Models. *International Conference on Learning Representations*, (2014), May 2014. arXiv: 1401.4082.
- [24] Hossein Sadeghi, Evgeny Andriyash, Walter Vinci, Lorenzo Buffoni, and Mohammad H. Amin. PixelVAE++: Improved PixelVAE with Discrete Prior. *ArXiv*, abs/1908.09948, 2019.

- [25] Casper Kaae Sønderby, Tapani Raiko, Lars Maaløe, Søren Kaae Sønderby, and Ole Winther. Ladder Variational Autoencoders. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems 29*, pages 3738–3746. Curran Associates, Inc., 2016.
- [26] Lucas Theis, Aäron van den Oord, and Matthias Bethge. A note on the evaluation of generative models. *International Conference on Learning Representations*, (2016), April 2016. arXiv: 1511.01844.
- [27] Aaron van den Oord, Oriol Vinyals, and Koray Kavukcuoglu. Neural Discrete Representation Learning. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 6306–6315. Curran Associates, Inc., 2017.

A Detailed analysis on the 2D problem

A.1 Impact of learning sharpness σ

The reconstruction (Fig. 6, Left) and generation (Fig. 6, Right) achieved by VAE when learning sharpness σ (LS-VAE) are compared with those of F-VAE (Fig. 7), where the fixed σ varies from 1. to .01.

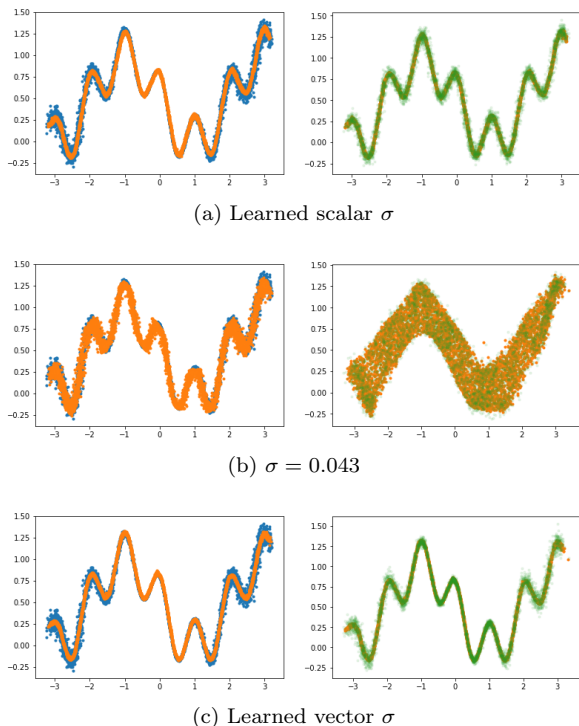


Figure 6: VAE with learned and fixed σ (better seen in color). Top: LS-VAE, learned scalar σ ; Medium: F-VAE, fixed σ with same value as the eventual σ learned by LS-VAE. Bottom: LV-VAE, learned vector σ . Left: Reconstructed samples (in orange) and initial samples (in blue). Right: Generated samples (in green) and their average (in orange).

Fig. 7 illustrates the transition of F-VAE as σ decreases: A rough approximation of the manifold is obtained for large σ values ($\sigma \geq .3$, Fig. 7.(a-b), Left); A better approximation is obtained for $\sigma = .1$, though F-VAE misses the details and achieves a denoising reconstruction (Fig. 7.(c), Left); as σ decreases, F-VAE learns the data noise as being

part of the manifold (Fig. 7.(d-e), Left). In the meanwhile, the generation is acceptable only for $\sigma = .1$ (Fig. 7.(c), Right); for higher σ values, the generation fails due to a bad approximation of the manifold (for $\sigma > .1$); for smaller σ values, the generation failure is due to the fact that the latent space encodes a rich information which is insufficiently compressed.

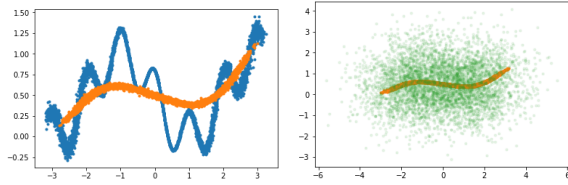
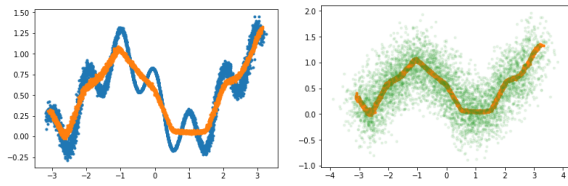
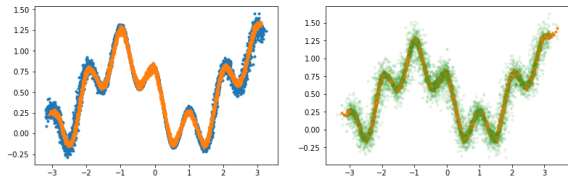
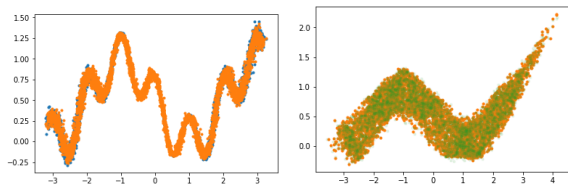
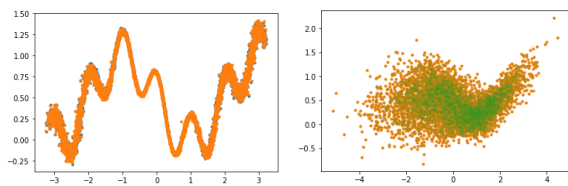
(a) $\sigma = 1.0$ (b) $\sigma = 0.3$ (c) $\sigma = 0.1$ (d) $\sigma = 0.03$ (e) $\sigma = 0.01$

Figure 7: F-VAE: VAE with fixed σ ranging from 1 to .01 from top to bottom (better seen in color). Left: Reconstructed samples (in orange) and initial samples (in blue). Right: Generated samples (in green) and their average (in orange).

A.2 Training dynamics

Likewise, Figs. 8 and 9 display the evolution of the training losses for respectively the fixed σ case (F-VAE, Fig. 8) and the learned one (LS-VAE, Fig. 9).

In the early learning stage, the reconstruction is bad. For a fixed, small value of σ (Fig. 8), the reconstruction loss is amplified by the $\frac{1}{2\sigma^2}$ factor:

$$-\log p_{\theta}(x|z) = \frac{1}{2\sigma^2} \|x - \mu(z)\|^2 + \log \sigma \quad (12)$$

The reconstruction then quickly improves as the model learns to auto-encode the data, and all losses stabilize as the auto-encoding is fine-tuned during the rest of the training period.

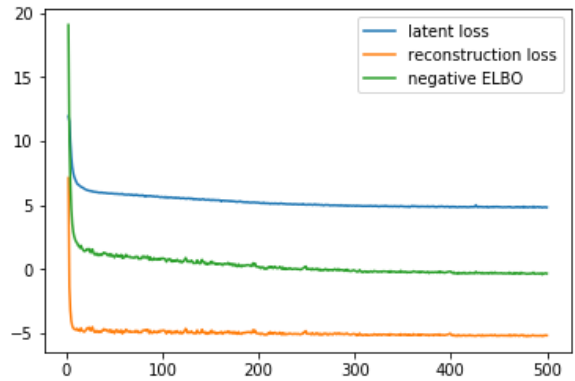


Figure 8: F-VAE: Reconstruction and latent losses along the training process for a fixed $\sigma = 0.043$.

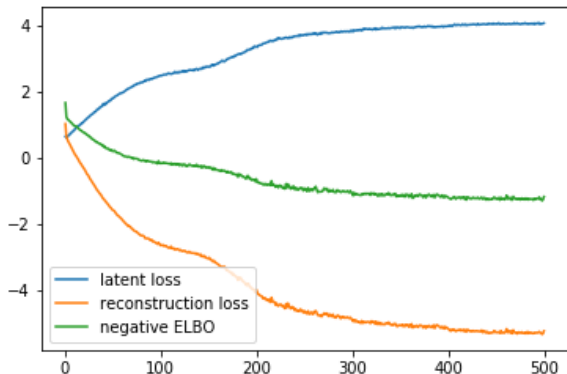


Figure 9: LS-VAE: Reconstruction and latent losses along the training process for a learned scalar σ .

On the contrary, when σ is learned (Fig. 9), the reconstruction loss is moderate in the early learning stage, as σ is initialized to 1. All losses are moderate at this early stage. The reconstruction loss gently decreases along the training period. Its decreases coincide with the decreases of the learned σ , illustrated on Fig. 10. Simultaneously, the latent loss also gently increases. Eventually, the reconstruction loss is the same for F-VAE and LS-VAE (circa -5); but the latent loss is significantly lower for LS-VAE than for F-VAE (≈ 4 compared to ≈ 5), yielding a better final ELBO, and therefore a higher generation quality. Our tentative interpretation is that the learning of σ enables a gradual increase of the information stored in the latent space, allowing the VAE to converge toward a more compressed latent representation.

This tentative interpretation is confirmed by complementary experiments, as follows. F-VAE is launched with a fixed σ , the value of which is set according to the dynamic schedule given on Fig. 10.

For all runs (with different random initialization of the encoder and decoder modules), F-VAE with this dynamic schedule of σ consistently yields the same loss curves, final ELBO and observed generation quality as LS-VAE. These experiments tend to confirm that the σ learning scheme allows LS-VAE to set and follow an efficient annealing optimization scheme.

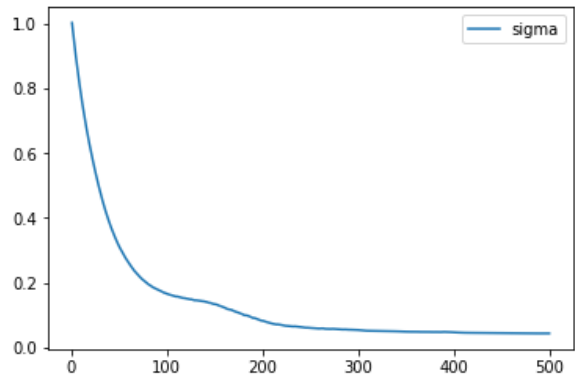


Figure 10: LS-VAE: Evolution of the learned value for σ .

B Complementary experiments on MNIST and CelebA

This section presents in more detail the results obtained by F-VAE and LV-VAE on MNIST and CelebA.

Same neural network architectures are used for F-VAE and LV-VAE: the only difference is that LV-VAE considers twice as much output channels as F-VAE (associating to each image pixel a pair $(\mu, \log \sigma)$ instead of μ).

It is emphasized that the training cost is identical for both F-VAE and LV-VAE.

B.1 MNIST

Neural architectures. The encoder and decoder architectures are symmetrical, with 3 convolution layers followed by 4 residual blocks for the encoder, and 4 residual blocks followed by 3 transposed convolutions for the decoder. The latent space dimension is 256. The training time is 5 hours (1,000 epochs).

Figs. 11 and 12 respectively illustrate the reconstruction with F-VAE (square error loss, equivalent to $\sigma = \frac{1}{\sqrt{2}}$) and LV-VAE (where σ is learned as an output of the decoder). The difference is moderate, as MNIST involves high-contrast images with few details. Still, F-VAE reconstructs the images as a smoothed version of themselves, with smoother edges and imperfections removed. In contrast, LV-VAE more closely matches all the details of the

original data.



Figure 11: F-VAE (fixed $\sigma = \frac{1}{\sqrt{2}}$): Reconstruction on MNIST.

Figs. 13 and 14 likewise illustrate the generation with F-VAE and LV-VAE, comparing for each model the mean of the output $p_\theta(x|z)$ with actual samples from it. A striking finding is the difference between the samples and their expectation for F-VAE (Fig. 13): the former ones are close to random noise while the latter ones are smoothed realistic images.

As said (section 3.2), this random noise effect explains why the means are preferred to samples, inducing some discrepancy between the VAE design (F-VAE being optimized to sample $p_\theta(x|z)$) and its usage.

On the contrary, LV-VAE yields precise distri-



Figure 12: LV-VAE (σ learned as decoder output): Reconstruction on MNIST.

butions (Fig. 14), as illustrated by the similarity between the mean of $p_\theta(x|z)$ and its samples. The slight differences (noise on the edge of the shapes) are explained as LV-VAE retains larger σ values at the frontier between the black and white regions in the image.

As a counterpart for its ability to reconstruct details, LV-VAE stores more information in the latent space, making it more difficult to compress, and leading to less realistic generated images. The bottleneck here comes from the insufficient expressivity of the encoder and decoder architectures; a more powerful model is needed to appropriately compress the data.

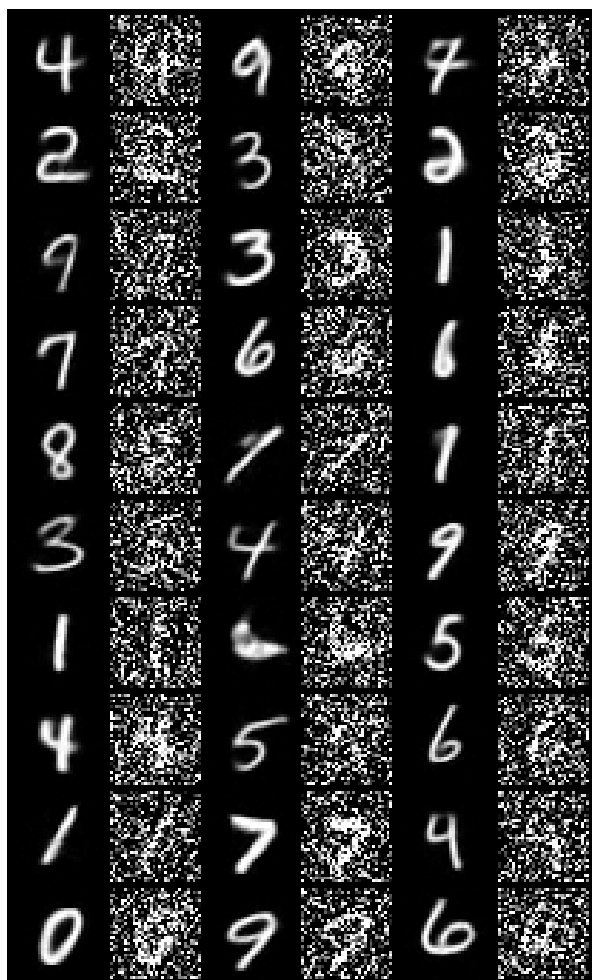


Figure 13: F-VAE (fixed $\sigma = \frac{1}{\sqrt{2}}$): Generations (mean and sample of $p_\theta(x|z)$) on MNIST.



Figure 14: LV-VAE (σ learned as decoder output): Generations (mean and sample of $p_\theta(x|z)$) on MNIST.

B.2 CelebA

Neural architectures. The encoder and decoder architectures are similar to those used for MNIST, but with 4 convolutions layers, 5 residual blocks, and a latent space dimension of 2048. The training time is ca 24 hours (200 epochs).

The reconstruction and generation performances of F-VAE and LV-VAE on CelebA are interpreted in much the same way as on MNIST, with the fact that the human eye is more able to identify unrealistic data when it comes to human faces. Figs. 15 and 16 show the reconstructions from F-VAE and LV-VAE. Like for MNIST, the F-VAE reconstructions (Fig. 15) are smoothed out, and lacking detail (see e.g. the hair or the tip of the nose). Again, the LV-VAE reconstructions (Fig. 16) are more detailed and almost pixel perfect.

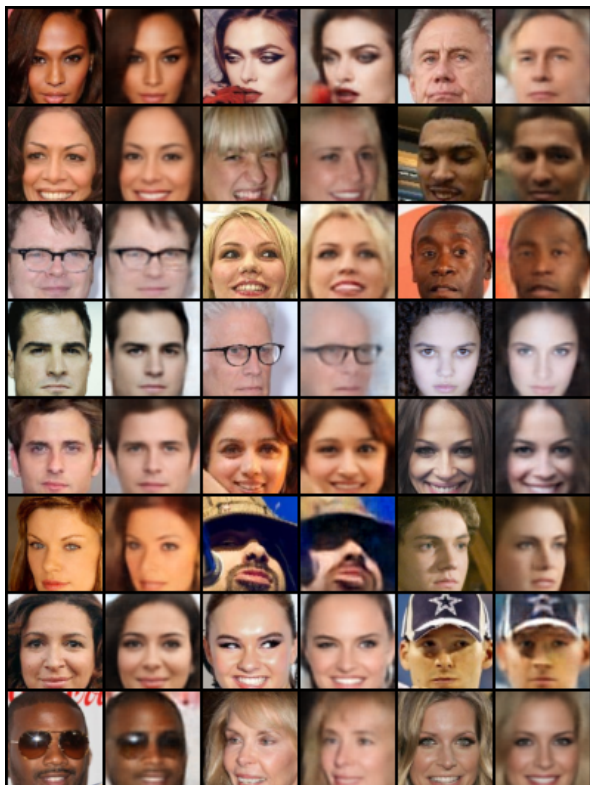


Figure 15: F-VAE (fixed $\sigma = \frac{1}{\sqrt{2}}$): Reconstructions on CelebA.



Figure 16: LV-VAE (σ learned as decoder output): Reconstructions on CelebA.

Figs. 17 and 18 likewise illustrate the generation with F-VAE and LV-VAE, showing actual samples of $p_\theta(x|z)$ and their expectation for each model. The remarks made for MNIST still hold: for F-VAE, the samples show a very high level of noise, while the means of $p_\theta(x|z)$ have a similar quality as the reconstructions. For LV-VAE, the samples are very close to the means, indicating that the learned variance is low. While the generated images are detailed, these details however lack a global consistency, harming the overall realism of the generated images. This limitation is likewise explained from the insufficient expressivity of the encoder and decoder neural networks, as well as the latent space structure.

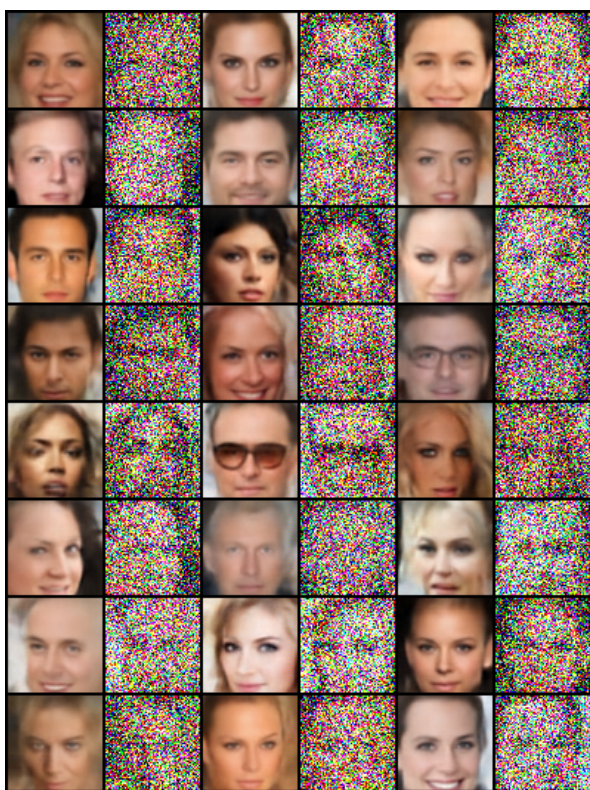


Figure 17: F-VAE (fixed $\sigma = \frac{1}{\sqrt{2}}$): Generations (mean and sample of $p_{\theta}(x|z)$) on CelebA.



Figure 18: LV-VAE (σ learned as decoder output): Generations (mean and sample of $p_{\theta}(x|z)$) on CelebA.