

---

# Embedding-reparameterization procedure for manifold-valued latent variables in generative models

---

**Eugene Golikov**

Neural Networks and Deep Learning Lab  
Moscow Institute of Physics and Technology  
Russia  
golikov.ea@mipt.ru

**Maksim Kretov**

Neural Networks and Deep Learning Lab  
Moscow Institute of Physics and Technology  
Russia  
kretov.mk@mipt.ru

## Abstract

Conventional prior for Variational Auto-Encoder (VAE) is a Gaussian distribution. Recent works demonstrated that choice of prior distribution affects learning capacity of VAE models. We propose a general technique (embedding-reparameterization procedure, or ER) for introducing arbitrary manifold-valued variables in VAE model. We compare our technique with a conventional VAE on a toy benchmark problem. This is work in progress.

## 1 Introduction

Variational Auto-Encoder (VAE) [1] and Generative Adversarial Networks (GAN) [2] show good performance in modelling real-world data such as images well. The key idea of both frameworks is to map a simple distribution (typically Gaussian) of lower dimension to a high-dimensional observation space by a complex non-linear function (typically neural network). Most of research efforts are concentrated on the enhancement of training procedure and neural architectures giving rise to a variety of elegant extensions for VAE and GANs [3].

We consider prior distribution that is mapped to data distribution  $p(x)$  as one of design choices when building generative model. Its importance is highlighted in a number of works [4, 5, 6, 7]. Although [5] provides an extensive overview of usage of  $L_2$ -normalized latent variables (points lying on a hypersphere); this is clearly just one of the possible design choices for prior distribution in generative model.

Recent works [5, 6, 7] argued that manifold hypothesis for data [8] provides evidence in favor of using more complicated priors than Gaussian, for which the topology of latent space matches that of the data. The above mentioned works derived analytic formulas for reparameterization of probability density on manifold (hypersphere in [5] and Lie group  $SO(3)$  in [6]).

A somewhat less rigorous argument in favor of using manifold-valued latent variables is that we can represent generative process for data as having two sources of variation (see Figure 1): one is uniform sampling from a group of transformations that we consider as compact symmetry groups (for example group of rotations) and another one is all the rest. This favors the choice of such topology of the latent space that would match "real" generative process: choose uniform distribution on some compact symmetry group as a prior distribution for latent variables.

Once a universal procedure for fast prototyping of VAE with different manifold-valued variables is available, such VAE can be used for estimating the likelihood integral  $p(x|Model)$  (for example




Figure 1: Observed data  $X$  are generated by *uniform* sampling from compact symmetry group  $G$  and other independent factors of variation  $V$  (for example, label of the class).

using IWAE estimate [9]) and thus make conclusions about latent symmetries that are present in the data. This was one of the key motivations for the current work.

All of above brings to the focus the case of continuously differentiable symmetry groups (Lie groups), which is a special case of manifold-valued latent variables.

## 2 Manifold-valued latent variables

Let us make the following preliminary assumption:

*Data  $x \sim p(x)$ ,  $x \in S \subset \mathbf{R}^n$  are generated as on Figure 1 with Lie group  $G$  embedded in  $\mathbf{R}^m$  and there is a continuous mapping  $G \rightarrow S$ .*

When using images as a test bed it implies that images generated by "close" symmetry elements (say two similar rotation angles  $\phi_1$  and  $\phi_2$ ) are also close in the pixel space. It justifies using additional tricks such as continuity loss [6] for training VAE with manifold-valued latent variables.

### 2.1 Construction of VAE

Recall the optimization problem for VAE [1]:

$$\mathcal{L}(\phi, \psi) = \mathbb{E}_{x \sim \mathcal{D}} [\mathbb{E}_{z \sim q_\phi(z|x)} [\log p_\psi(x|z)] - \text{KL}(q_\phi(z|x) \| p(z))] \rightarrow \max_{\phi, \psi},$$

where  $\mathcal{D}$  denotes the data distribution,  $q_\phi(z|x)$  is a posterior distribution on latent space  $Z$ ,  $p(z)$  is the corresponding prior, and  $p_\psi(x|z)$  is the likelihood of a data point  $x$  given  $z$ . In order to construct a VAE with manifold-valued latent variables, we need the following:

1. An encoder that produces the posterior distribution  $q(z|x)$  from a parametric family of distributions on a manifold.
2. An ability to sample from this posterior distribution:  $z \sim q(z|x)$ .
3. An ability to compute KL-divergence between this posterior and a given prior.

Recent works [5, 6] proposed approaches to working with manifold-valued latent variables that are similar in spirit to ours: they derive a reparameterization of probability density defined on smooth manifold and use it in VAE. Problem is that such derivation appears to be complicated and needs to be done for all manifolds of interest.

Our approach is the following. First of all, we introduce a hidden latent space  $Z_{hid}$ , such that  $\dim Z_{hid} = \dim \mathcal{M} = n$ , where  $\mathcal{M}$  is our manifold lying in a latent space  $Z$  of dimension  $m > n$ . Let  $p(z_{hid})$  be a prior distribution on  $Z_{hid}$ .

Suppose then, we have an embedding  $f : Z_{hid} \rightarrow Z$ , so that  $f(Z_{hid}) \subset \mathcal{M}$ . Being an embedding requires  $f$  to be a diffeomorphism with its image, in particular,  $f$  should be a differentiable injective map. We also pose an additional constraint on  $f$ : it should map the prior on  $Z_{hid}$  to a prior on the manifold  $\mathcal{M}$ ; in other words, if  $z_{hid} \sim p(z_{hid})$ , then  $f(z_{hid}) \sim p(z)$ .

Using this embedding  $f$ , we can construct a VAE with manifold-valued latent variables as depicted on the right part of Figure 2. In this case the posterior distribution  $q(z_{hid}|x)$  on  $Z_{hid}$  together with the embedding  $f$  induce a posterior distribution  $q(z|x)$  on  $\mathcal{M} \subset Z$ . We then have to compute KL-divergence between this induced posterior and the prior  $p(z)$  on the manifold. Despite the fact that in this case the probability mass is concentrated on the manifold  $\mathcal{M}$  and hence the probability density on  $Z$  is degenerate, we can define the manifold probability densities  $q_{\mathcal{M}}(z|x)$  and  $p_{\mathcal{M}}(z)$  (see Appendix 5.1 for details). Moreover, the corresponding KL-divergence is equivalent to the KL-divergence between distributions defined on  $Z_{hid}$  (Appendix 4.3):

$$\text{KL}(q_{\mathcal{M}}(z|x) \| p_{\mathcal{M}}(z)) = \text{KL}(q(z_{hid}|x) \| p(z_{hid}))$$

Hence the final optimization problem for model on the right part of Figure 2 becomes the following:

$$\mathcal{L}(\phi, \psi) = \mathbb{E}_{x \sim \mathcal{D}} [\mathbb{E}_{z_{hid} \sim q_\phi(z_{hid}|x)} [\log p_\psi(x|f(z_{hid}))]] - \text{KL}(q_\phi(z_{hid}|x) \| p(z_{hid})) \rightarrow \max_{\phi, \psi},$$

where  $\phi$  are parameters of VAE encoder, which encodes the object  $x$  into  $Z_{hid}$  space, and  $\psi$  are parameters of VAE decoder which maps the manifold  $\mathcal{M} \subset Z$  to data-manifold in feature space;  $\mathcal{D}$  is our data distribution.

Thereby working with probability distributions induced on manifold of interest is easy: both terms in VAE loss (reconstruction error and KL-divergence) are easily calculated in the original hidden space  $Z_{hid}$  that is further mapped on a manifold.

## 2.2 Learning manifold embedding

To apply the procedure described above, we have to construct an embedding  $f$ . In order to do this, we propose the following procedure:

1. Sample data from  $p(z)$  (distribution on  $\mathcal{M}$ ).
2. Train Wasserstein Auto-Encoder (WAE) [10] on the data from  $p(z)$  (feature space) and the latent space  $Z_{hid}$  with the prior  $p(z_{hid})$ : see the left part of Figure 2.
3. Use the decoder of this trained WAE as our embedding function  $f$ .

Our motivation is the following: since the dimension of latent space  $Z_{hid}$  and the dimension of manifold  $\mathcal{M}$  are the same, the reconstruction term in WAE objective constraints its decoder to be an injective map. Since it is represented with a neural network, it is also differentiable. The objective of WAE learning also forces its decoder to map a prior distribution on a latent space (in our case,  $p(z_{hid})$ ) to a distribution of data to the feature space (in our case,  $p(z)$ ). Hence WAE decoder is an ideal candidate for an embedding  $f$ .



Figure 2: **Left:** Scheme diagram of WAE with feature space  $Z$  and latent space  $Z_{hid}$  that learns the prior distribution  $p(z)$  on manifold  $\mathcal{M} \subset Z$ . **Right:** The generative model with manifold-valued latent variables  $z$ . The squared node is deterministic.

At first glance the described model leaves quite similar questions as vanilla VAE: we "shifted" the complex task of learning non-homeomorphic manifolds of a different topology (latent space and data space) from the VAE decoder to sub-module of the same VAE but pretrained using WAE. Nevertheless, the procedure ensures better control over mapping to manifold and one can develop corresponding metrics to control the quality of mapping.

## 3 Introducing symmetries of latent manifold into encoder

Recall that in our scheme an encoder  $q(z_{hid}|x)$  together with embedding  $f : Z_{hid} \rightarrow Z$  induce a family of posterior distributions  $q(z|x)$  on  $\mathcal{M}$ ; let us call this family  $\mathcal{Q}$ .

A natural requirement to  $\mathcal{Q}$  is to have the same symmetries as  $\mathcal{M}$  has. Suppose we have a symmetry group  $G$  of  $\mathcal{M}$  acting on  $Z$ , i.e.

$$\forall z \in \mathcal{M}, \forall g \in G \quad gz \in \mathcal{M}.$$

For example, if  $\mathcal{M}$  is an  $n$ -dimensional sphere  $S^n$  in  $Z = \mathbf{R}^{n+1}$ ,  $G$  is a group of rotations  $SO(n+1)$ . We require  $G$  to also be a symmetry of  $\mathcal{Q}$  also:

$$\forall q \in \mathcal{Q}, \forall g \in G \quad \exists q' \in \mathcal{Q} : \forall z \in \mathcal{M} \quad q'(gz) = q(z).$$




Figure 3: The generative model with group action-encoder acting on manifold-valued latent variable  $z$ .

Table 1: Results on the toy task for different models.

Models	ELBO
VAE, $\dim Z = 1$	$183.98 \pm 11.66$
Manifold-latent VAE with learned $f$ , $\dim Z_{hid} = 1$	$197.19 \pm 20.46$
Manifold-latent VAE with $f = f_{proj}$ , $\dim Z_{hid} = 1$	$193.40 \pm 24.57$
Manifold-latent VAE with learned $f$ and group action encoder, $\dim Z_{hid} = 1$	<b><math>259.03 \pm 59.14</math></b>
VAE, $\dim Z = 2$	<b><math>356.53 \pm 22.96</math></b>

This means that if a symmetry  $g$  of  $\mathcal{M}$  acts on samples  $z$  from a distribution  $q \in \mathcal{Q}$ , we should get samples from another distribution  $q'$  from the same family  $\mathcal{Q}$ . Note that we did not pose this requirement while training  $f$ , hence it would not generally be satisfied. Therefore we have to symmetrize  $\mathcal{Q}$  explicitly.

In order to do this we introduce a group action encoder  $a(x)$ , see Figure 3. This group action encoder produces an element  $g = a(x)$  of the symmetry group  $G$  of  $\mathcal{M}$ , which further acts on a sample  $z = f(z_{hid})$ . This effectively enriches the posterior family  $\mathcal{Q}$  with  $q' : q'(gz) = q(z)$ .

This procedure has close connection with homeomorphic VAE [6]. Suppose our manifold  $\mathcal{M}$  is a compact Lie group. Then it is homeomorphic to its own symmetry group:  $\mathcal{M} \cong G$ . Then our group action-encoder is equivalent to  $R_\mu$  of [6].

## 4 Experiments and conclusions

We followed the same experimental setup as for a toy task in paper [5], but without noise. <sup>1</sup> Sampling of a batch from the dataset consisted of two steps:

1. We generated uniformly distributed points on a 1-dimensional unit sphere embedded in  $\mathbf{R}^2$ .
2. We applied a non-linear fixed transformation  $\mathbf{R}^2 \rightarrow \mathbf{R}^{100}$  implemented as a randomly initialized multilayer perceptron with one hidden layer of size 100 and ReLU nonlinearity. Xavier-uniform initialization scheme was applied to the hidden layer.

All models are VAEs with the posterior distribution  $q(z_{hid}|x)$  (Beta on  $[0, 1]^n$ ), the prior distribution  $p(z)$  (uniform of  $[0, 1]^n$ ) and the likelihood  $p(x|z)$  (Gaussian on  $\mathbf{R}^{100}$ ). As for the reparameterization function  $f(z_{hid})$ , it was either WAE-MMD or the exact mapping from segment  $[0, 1]$  into a 1-dimensional circle ("Projection") in the first layer of decoder:

$$f_{proj}(z_{hid}) = (\cos(2\pi z_{hid}) \quad \sin(2\pi z_{hid}))^T.$$

The dimensions of latent variables  $n$  were either 1 or 2.

In a case when the group action encoder is used, it produces an angle (element of  $SO(2)$ ), which is further used to rotate the sample  $z = f(z_{hid}) \in \mathbf{R}^2$ .

The results are presented in Table 1. All decoder structures that include manifold mapping show better results than a vanilla VAE with 1-dimensional latent Gaussian space.

<sup>1</sup>Our code is available on GitHub: [https://github.com/varenick/manifold\\_latent\\_vae](https://github.com/varenick/manifold_latent_vae)

## Acknowledgments

This work was supported by National Technology Initiative and PAO Sberbank project ID 0000000007417F630002.

## References

- [1] Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. *CoRR*, abs/1312.6114, 2013.
- [2] 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.
- [3] Antonia Creswell, Tom White, Vincent Dumoulin, Kai Arulkumaran, Biswa Sengupta, and Anil A. Bharath. Generative adversarial networks: An overview. *CoRR*, abs/1710.07035, 2017.
- [4] Matthew D. Hoffman and Matthew J. Johnson. Elbo surgery: yet another way to carve up the variational evidence lower bound, 2016.
- [5] Tim R. Davidson, Luca Falorsi, Nicola De Cao, Thomas Kipf, and Jakub M. Tomczak. Hyper-spherical variational auto-encoders. *34th Conference on Uncertainty in Artificial Intelligence (UAI-18)*, 2018.
- [6] Luca Falorsi, Pim de Haan, Tim R. Davidson, Nicola De Cao, Maurice Weiler, Patrick Forré, and Taco S. Cohen. Explorations in homeomorphic variational auto-encoding. *CoRR*, abs/1807.04689, 2018.
- [7] Jiacheng Xu and Greg Durrett. Spherical latent spaces for stable variational autoencoders, 2018.
- [8] Mikhail Belkin, Partha Niyogi, and Vikas Sindhwani. Manifold regularization: A geometric framework for learning from labeled and unlabeled examples. *J. Mach. Learn. Res.*, 7:2399–2434, 2006.
- [9] Yuri Burda, Roger B. Grosse, and Ruslan Salakhutdinov. Importance weighted autoencoders. *CoRR*, abs/1509.00519, 2015.
- [10] Ilya Tolstikhin, Olivier Bousquet, Sylvain Gelly, and Bernhard Scholkopf. Wasserstein auto-encoders, 2017.

## 5 Appendix

### 5.1 Probability density functions with manifold support

Suppose we have a probability distribution on  $Z = \mathbf{R}^n$  with density  $p(z)$  and a diffeomorphism  $f : Z \rightarrow X$ , where  $X = \mathbf{R}^n$  as well. Then,  $f$  induces a probability distribution on  $X$  with the following density:

$$p(x) = p(f^{-1}(x)) |\det J_f(f^{-1}(x))|^{-1} = p(f^{-1}(x)) |\det J_{f^{-1}}(x)|.$$

Suppose now that  $X = \mathbf{R}^m$  with  $m > n$ , and  $f : Z \rightarrow X$  is a smooth embedding (which requires  $f$  to be a diffeomorphism between  $Z$  and  $f(Z)$ ). From this follows that  $f$  induces degenerate probability distribution on  $X$  since all the probability mass in  $X$  is concentrated on a manifold  $\mathcal{M} = f(Z)$ . The corresponding probability measure is trivial:

$$P(f(A)) = P(A)$$

for some event  $A$  on  $Z$ . Although we cannot define a valid probability density of  $X$ , we can define a manifold probability density on  $\mathcal{M} = f(Z)$  as follows:

$$\begin{aligned}
p_{\mathcal{M}}(f(z)) &:= \lim_{\text{Vol}(A) \rightarrow 0} \lim_{s.t. z \in A} \frac{P(f(A))}{\text{Vol}_{\mathcal{M}}(f(A))} \\
&= \lim_{\text{Vol}(A) \rightarrow 0} \lim_{s.t. z \in A} \frac{P(A)}{\text{Vol}_{\mathcal{M}}(f(A))} \\
&= \lim_{\text{Vol}(A) \rightarrow 0} \lim_{s.t. z \in A} \frac{\int_A p(z) dz_1 \dots dz_n}{\text{Vol}_{\mathcal{M}}(f(A))} \\
&= p(z) \lim_{\text{Vol}(A) \rightarrow 0} \lim_{s.t. z \in A} \frac{\int_A dz_1 \dots dz_n}{\text{Vol}_{\mathcal{M}}(f(A))},
\end{aligned}$$

where by  $\text{Vol}_{\mathcal{M}}(f(A))$  we denote an  $n$ -dimensional volume of  $f(A) \subset \mathcal{M}$ ; let us define this volume. Let  $\Omega$  be an open subset of  $Z$ . Then its image under embedding  $f$  is an open subset of a manifold  $f(\Omega)$  (open in terms of the topology of  $\mathcal{M}$ ). If  $Z$  is a Euclidean space, then the "volume" of  $\Omega$  is given simply as:

$$\text{Vol}(\Omega) = \int_{\Omega} dz_1 \dots dz_n.$$

Since  $\mathcal{M}$  is embedded into  $X$ , and  $X$  is a Euclidean space, we can measure an  $n$ -dimensional "volume" of  $f(\Omega) \subset \mathcal{M}$ . It is given as:

$$\text{Vol}_{\mathcal{M}}(f(\Omega)) = \int_{f(\Omega)} \sqrt{|\det G(z)|} dz_1 \dots dz_n,$$

where  $G(z)$  is a metric tensor on  $Z$ , induced by the scalar product  $\langle \cdot \rangle$  on  $X$  and the embedding  $f$ :

$$G_{ij}(z) = \left\langle \frac{df(z)}{dz^i}, \frac{df(z)}{dz^j} \right\rangle.$$

Returning to our formula for probability density on  $\mathcal{M}$ , we now have:

$$\begin{aligned}
p_{\mathcal{M}}(f(z)) &= p(z) \lim_{\text{Vol}(A) \rightarrow 0} \lim_{s.t. z \in A} \frac{\int_A dz_1 \dots dz_n}{\int_{f(A)} \sqrt{|\det G(z)|} dz_1 \dots dz_n} \\
&= p(z) |\det G(z)|^{-1/2}.
\end{aligned}$$

Or,

$$p_{\mathcal{M}}(x) = p(f^{-1}(x)) |\det G(f^{-1}(x))|^{-1/2}.$$

## 5.2 Calculation of KL divergence in the case of normalizing flow

$$\begin{aligned}
\text{KL}(q_{\mathcal{M}}(z|x) \| p_{\mathcal{M}}(z)) &= \mathbb{E}_{z \sim q_{\mathcal{M}}(z|x)} (\log q_{\mathcal{M}}(z|x) - \log p_{\mathcal{M}}(z)) \\
&= \mathbb{E}_{z_{hid} \sim q(z_{hid}|x)} (\log q(z_{hid}|x) + \log |\det J_f(z_{hid})|^{-1} \\
&\quad - \log p(z_{hid}) - \log |\det J_f(z_{hid})|^{-1}) \\
&= \mathbb{E}_{z_{hid} \sim q(z_{hid}|x)} (\log q(z_{hid}|x) - \log p(z_{hid})) \\
&= \text{KL}(q(z_{hid}|x) \| p(z_{hid})).
\end{aligned}$$

where  $q(z_{hid}|x)$  is the posterior distribution (i.e. fully-factorized Gauss or Beta) on latent variables of WAE, which we use for manifold embedding,  $p(z_{hid})$  is the corresponding prior (i.e. standard Gauss or Uniform),  $f$  is the decoder of the WAE, which we use to transform the latent space of WAE into manifold  $\mathcal{M}$ , and  $J_f(z_{hid})$  is the Jacobian of this transformation. As we see, log-determinants of Jacobians cancel out, and we are left with the KL-divergence on latent space of WAE.

### 5.3 Calculation of KL divergence in case of embedding map

$$\begin{aligned}\text{KL}(q_{\mathcal{M}}(z|x)||p_{\mathcal{M}}(z)) &= \mathbb{E}_{z \sim q_{\mathcal{M}}(z|x)}(\log q_{\mathcal{M}}(z|x) - \log p_{\mathcal{M}}(z)) \\ &= \mathbb{E}_{z_{hid} \sim q(z_{hid}|x)}(\log q(z_{hid}|x) + \log |\det G(z_{hid})|^{-1/2} \\ &\quad - \log p(z_{hid}) - \log |\det G(z_{hid})|^{-1/2}) \\ &= \mathbb{E}_{z_{hid} \sim q(z_{hid}|x)}(\log q(z_{hid}|x) - \log p(z_{hid})) \\ &= \text{KL}(q(z_{hid}|x)||p(z_{hid})),\end{aligned}$$

where  $G$  denotes the metric tensor of the embedding  $f$ . As in Appendix 4.2, the corresponding terms cancel out.