Supplementary material for Optimizing the Latent Space of Generative Networks

Anonymous Authors[1]

Reconstruction. Table 1 gives the pSNR scores for all the methods. In particular below the lines, are DCGAN and GLO trained with Lap_1 but with images reconstruct with MSE. Since the MSE is directly related to the pSNR, this gives higher scores, even though it is the same model. The difference between DCGAN and GLO stays the same regardless of the reconstruction loss.

Figure 1 shows the distribution of pSNR scores for all the methods and across all the datasets. It is surprising to see that DCGAN pSNR scores have the same distribution as VAE or GLO. We could have expected a few images to be very well reconstruct from DCGAN if it was storing a few images and drop the others, as a crude mode-dropping. This result suggests that the " mode-dropping" of GANs is more subtle than simple matching a subset of images from the dataset. That being said, these plots were made on only 1K images, which may not be enough to conclude anything with certainty.

Probing the latent space. Figure 2 shows the cumulative spectrum of the covariance matrix of the latent variable. This measure shows how much space is used by the model to fit the dataset representations in the latent space. A plot close to the diagonal means that the model use as much space as available to store its latent variables, leading to a better reconstruction, but a sharper spectrum means that the model is using a smaller space to store the representation, which may either mean a worse reconstruction or a better compression of the dataset.

[1]Anonymous Institution, Anonymous City, Anonymous Region, Anonymous Country. Correspondence to: Anonymous Author <anon.email@domain.com>.

Preliminary work. Under review by the International Conference on Machine Learning (ICML). Do not distribute.

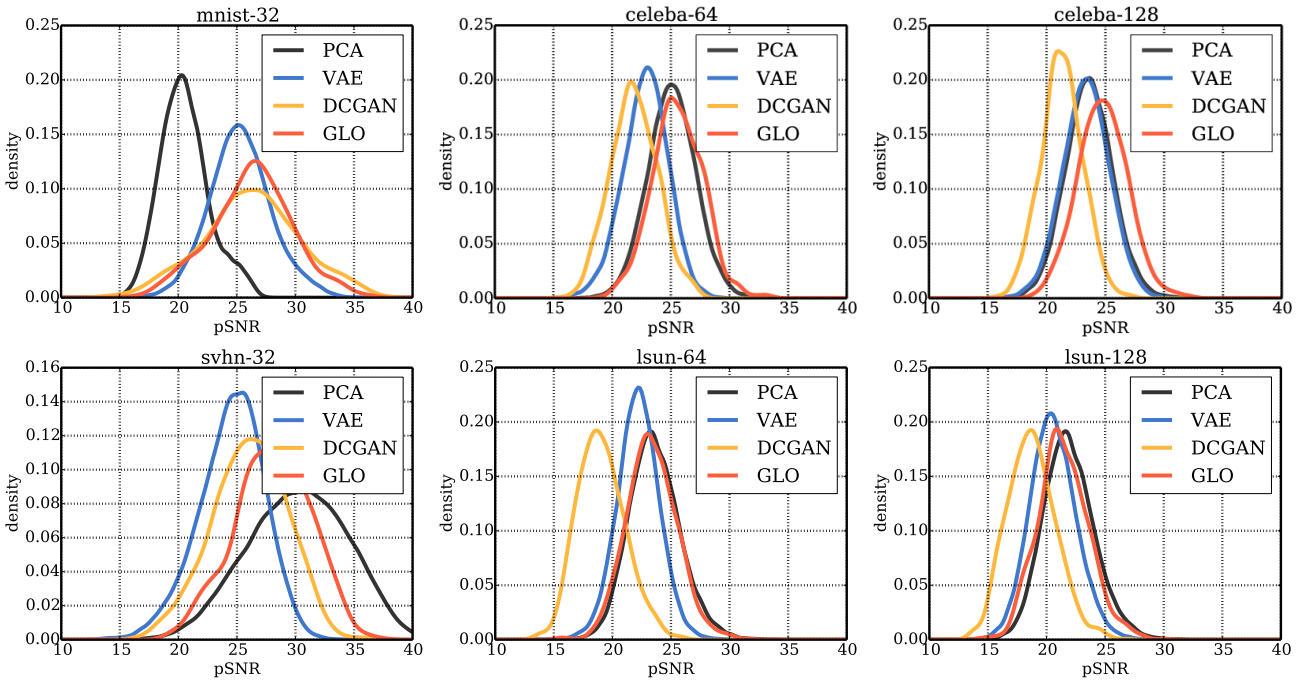

Figure 1. Distribution of the pSNR scores for all the methods on each dataset, out of 1K images. There are no obvious difference in the shape of the distributions of the different models.

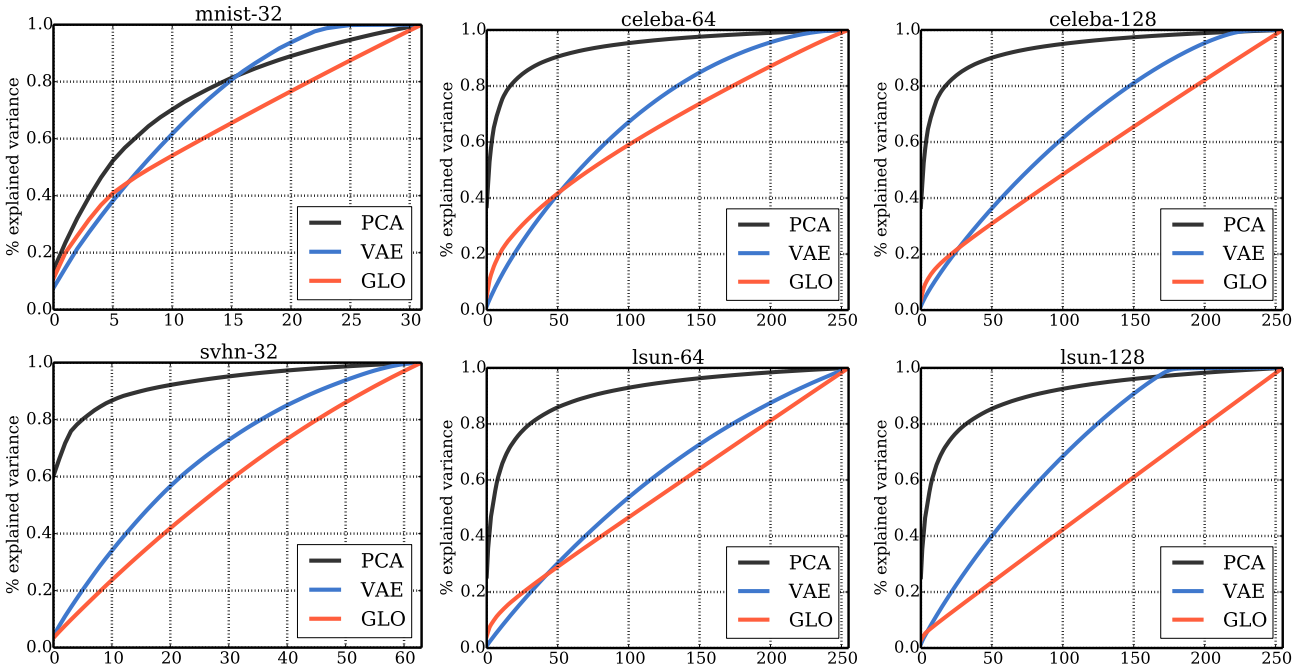

Figure 2. Cumulative spectrum of the covariance matrix of the latent variable. The closer it is to the diagonal, the more the information is well shared across the dimensions of the latent space. On the contrary, a cumulative that reach 1 rapidly means that the model only used a few directions for its latent space.

	MNIST		SVHN		CelebA				LSUN			
	32		32		64		128		64		128	
	train	test	train	test	train	test	train	test	train	test	train	test
PCA	20.6	20.3	30.2	30.3	25.1	25.1	23.6	23.6	23.6	23.7	21.9	22.0
VAE	25.3	25.0	24.5	24.5	22.8	22.8	23.4	23.2	22.1	22.1	20.6	20.6
DCGAN	25.8	26.2	26.0	26.0	21.9	21.9	21.3	21.3	19.0	19.1	18.7	18.7
GLO	26.2	26.2	27.9	28.0	25.5	25.6	24.7	24.8	23.3	23.4	21.4	21.4
DCGAN	26.9	27.2	30.2	30.1	25.0	25.0	23.5	23.5	21.8	21.9	20.8	20.9
GLO	27.0	27.2	30.7	30.7	27.7	27.7	26.4	26.4	24.8	24.9	22.0	22.1

Table 1. pSNR of reconstruction for different models. Above the line, the codes were found using Lap_1 loss (although the test error is still measured in pSNR). Below the line, the codes were found using mean square error. Note that the generators of the VAE and GLO models were trained to reconstruct in Lap_1 loss.