


β - Annealed Variational Autoencoders for glitches

Sivaramakrishnan Sankarapandian, Brian Kulis
Electrical and Computer Engineering



Abstract:


Gravitational wave detectors such as LIGO and Virgo are susceptible to various types of instrumental and environmental disturbances known as glitches which can mask and mimic gravitational waves. While there are 22 classes of non-Gaussian noise gradients currently identified, the number of classes is likely to increase as these detectors go through commissioning between observation runs. Since identification and labelling new noise gradients can be arduous and time-consuming, we propose Beta-Annealed VAEs to learn representations from spectrograms in an unsupervised way. Using the same formulation as [1], we view Bottleneck-VAEs [2] through the lens of information theory and connect them to Beta-VAEs [3]. Motivated by this connection, we propose an annealing schedule for the hyperparameter β in Beta-VAEs which has advantages of: 1) One fewer hyperparameter to tune, 2) Better reconstruction quality, while producing similar levels of disentanglement.

VAE, Beta-VAE, Bottleneck-VAE

Generative model of data with observed data \mathbf{x} and latent variables \mathbf{z} ,

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

Posterior $p(\mathbf{z}|\mathbf{x})$ is approximated using q_ϕ



- VAEs optimize Evidence LowerBound (ELBO),

$$\mathcal{L}(\theta, \phi; \mathbf{x}, \mathbf{z}) = \mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x})}[\log p_\theta(\mathbf{x}|\mathbf{z})] - D_{KL}(q_\phi(\mathbf{z}|\mathbf{x})||p(\mathbf{z}))$$

- Beta VAEs optimize,

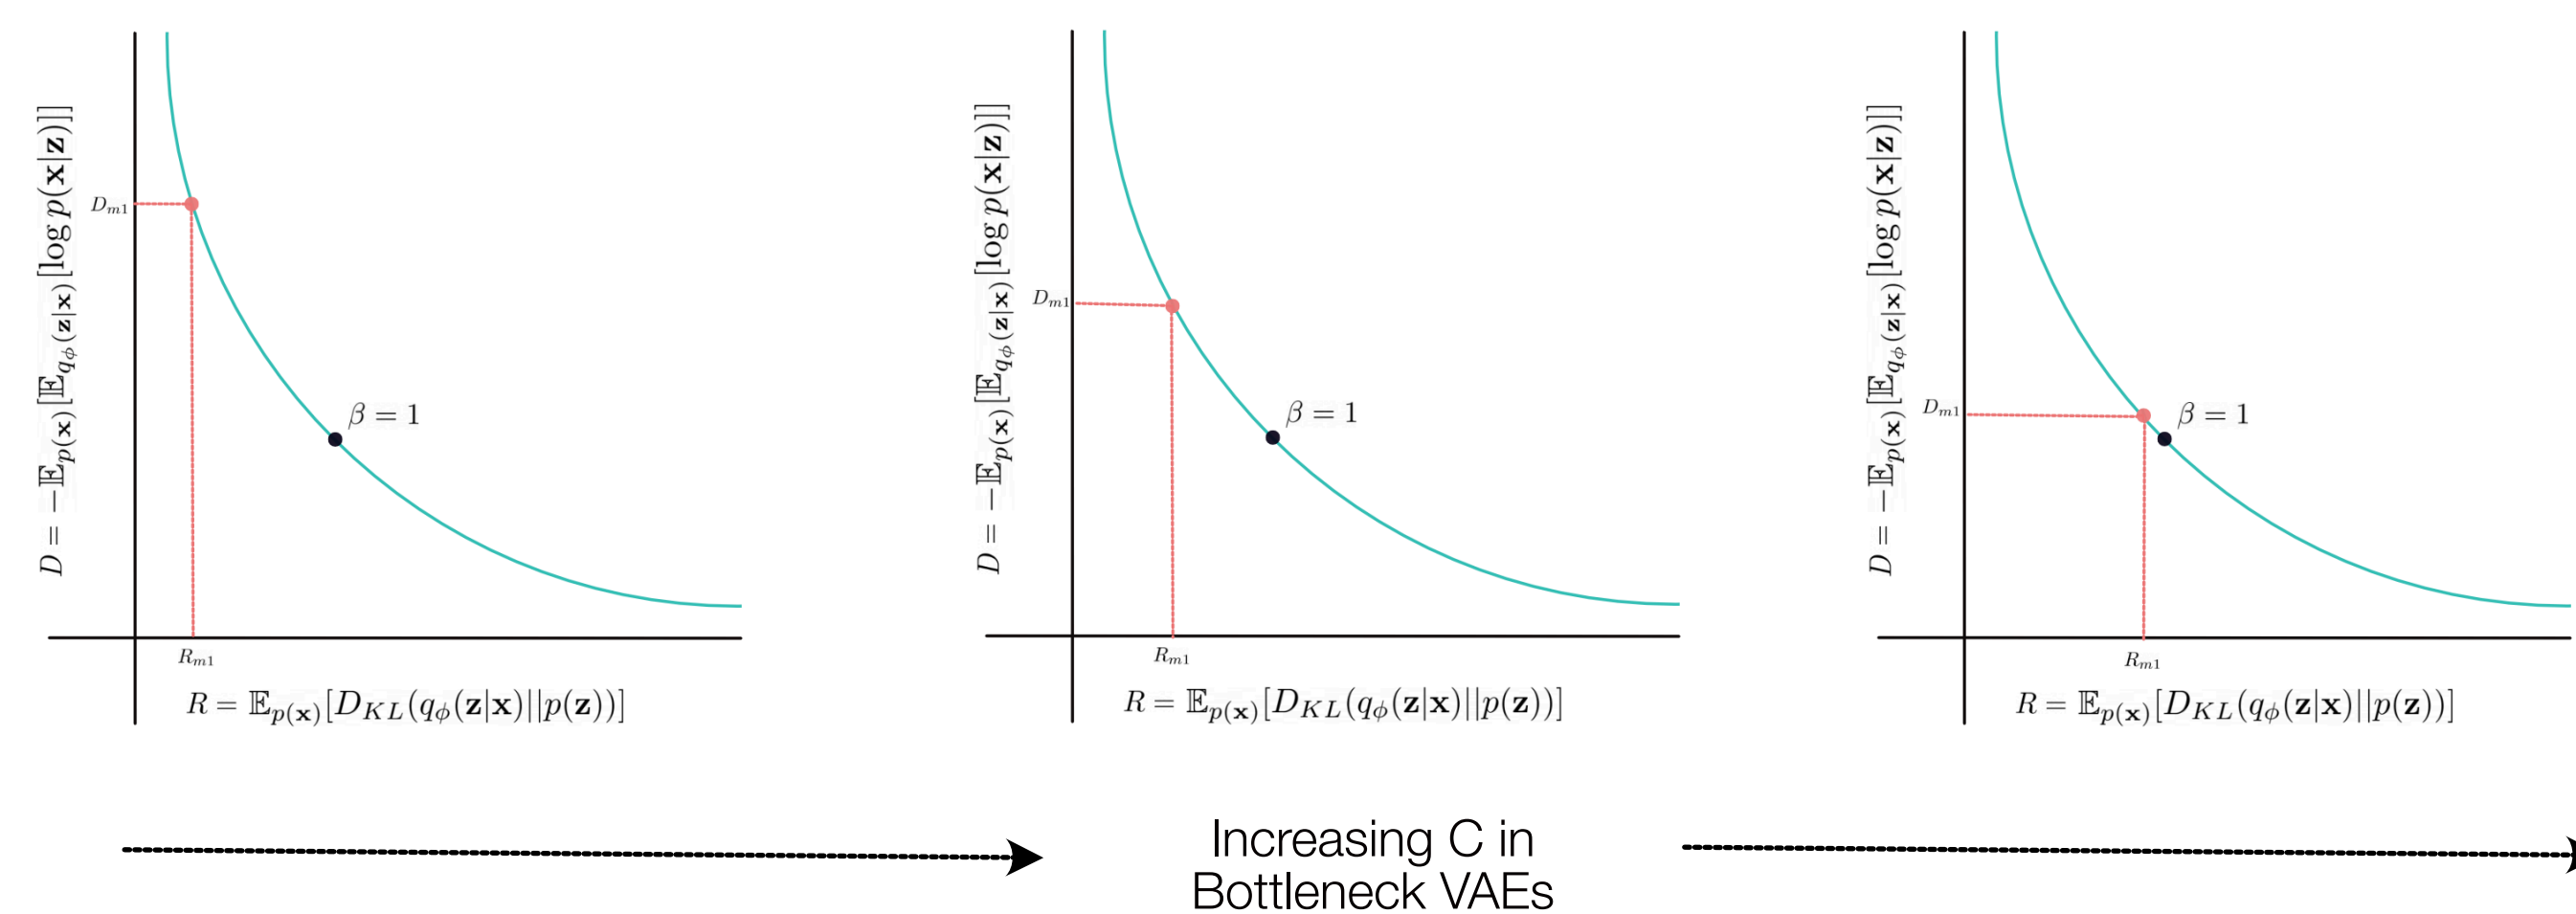
$$\mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x})}[\log p_\theta(\mathbf{x}|\mathbf{z})] - \beta D_{KL}(q_\phi(\mathbf{z}|\mathbf{x})||p(\mathbf{z}))$$

- Bottleneck VAEs optimize,

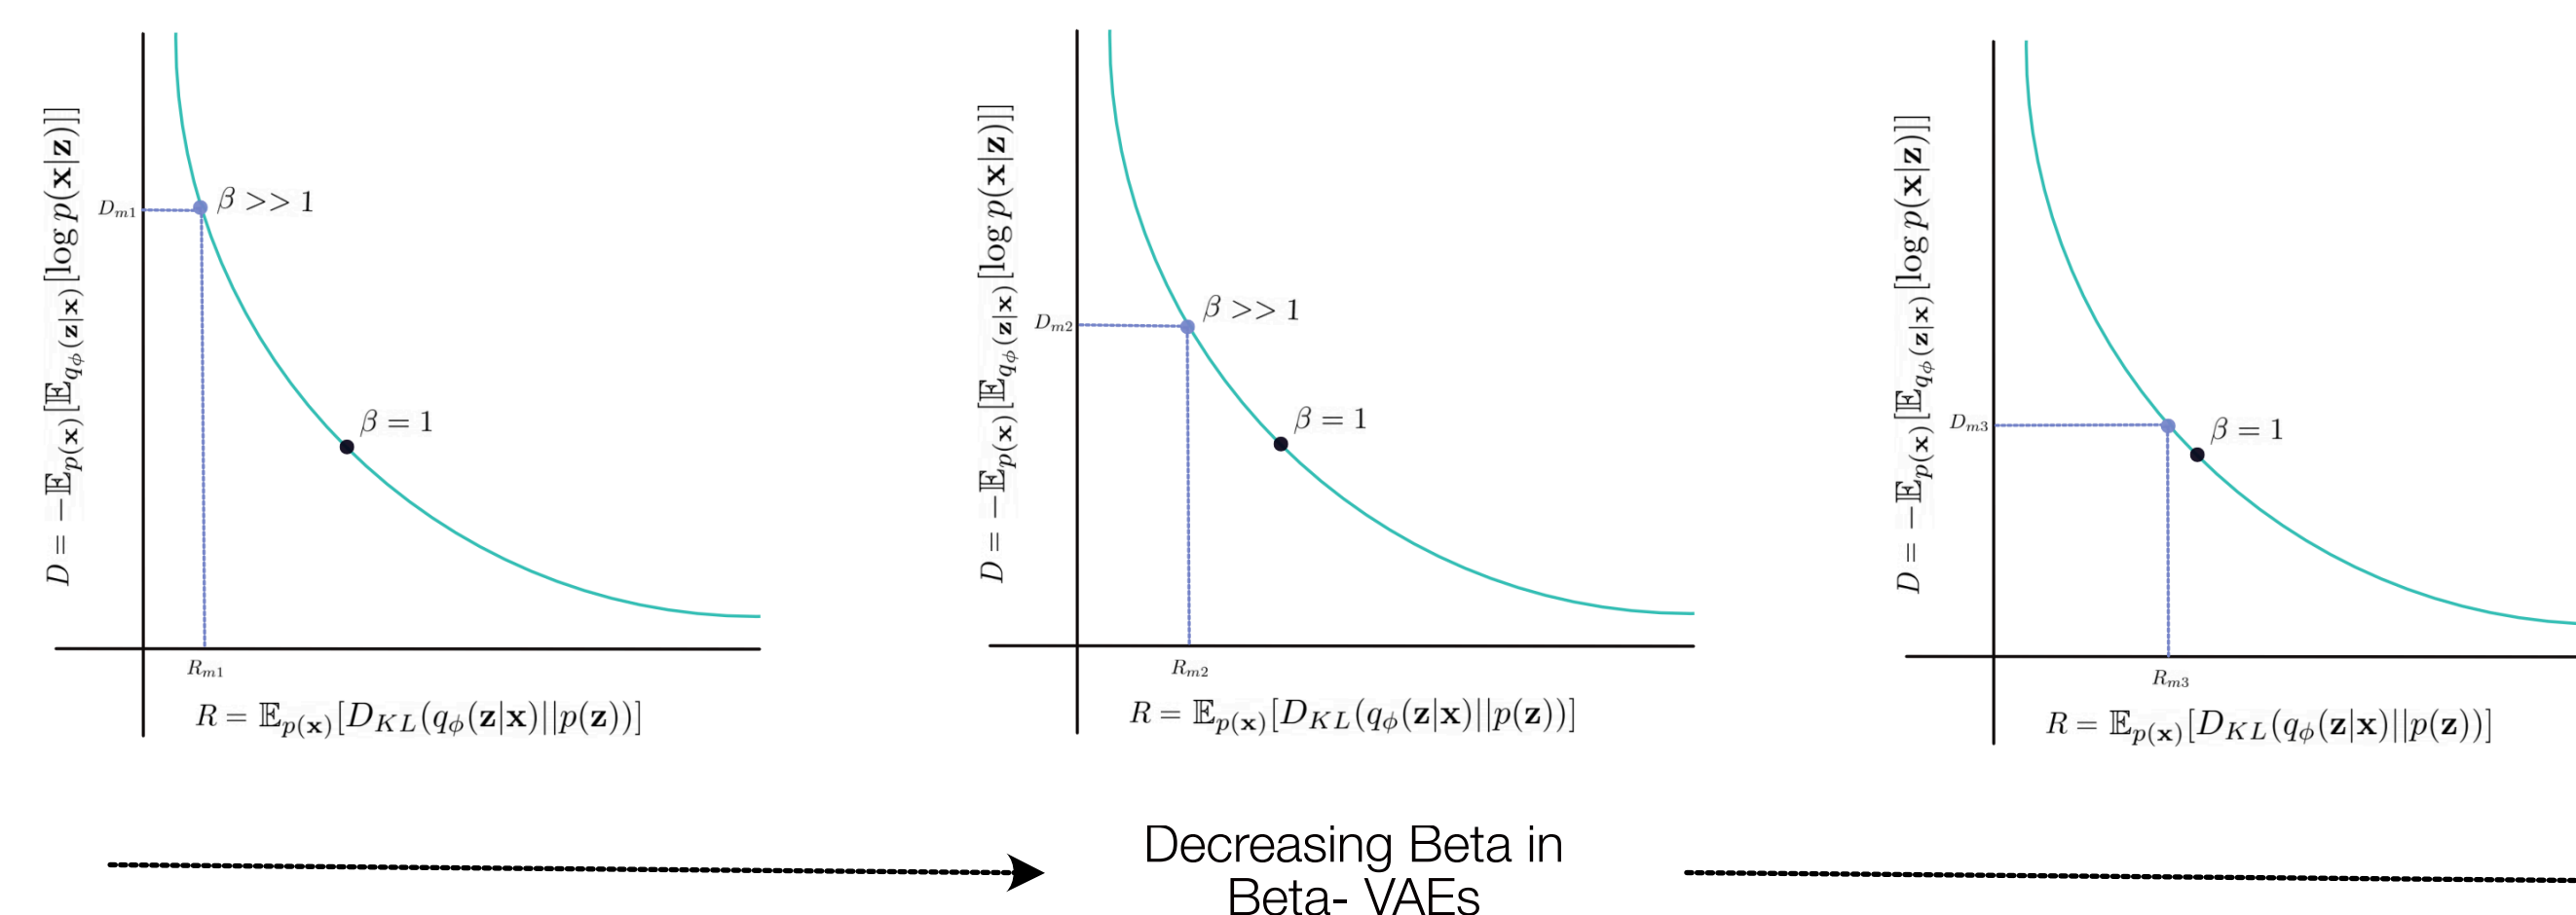
$$\mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x})}[\log p_\theta(\mathbf{x}|\mathbf{z})] - \gamma |D_{KL}(q_\phi(\mathbf{z}|\mathbf{x})||p(\mathbf{z})) - C|$$

Beta-Annealed VAEs

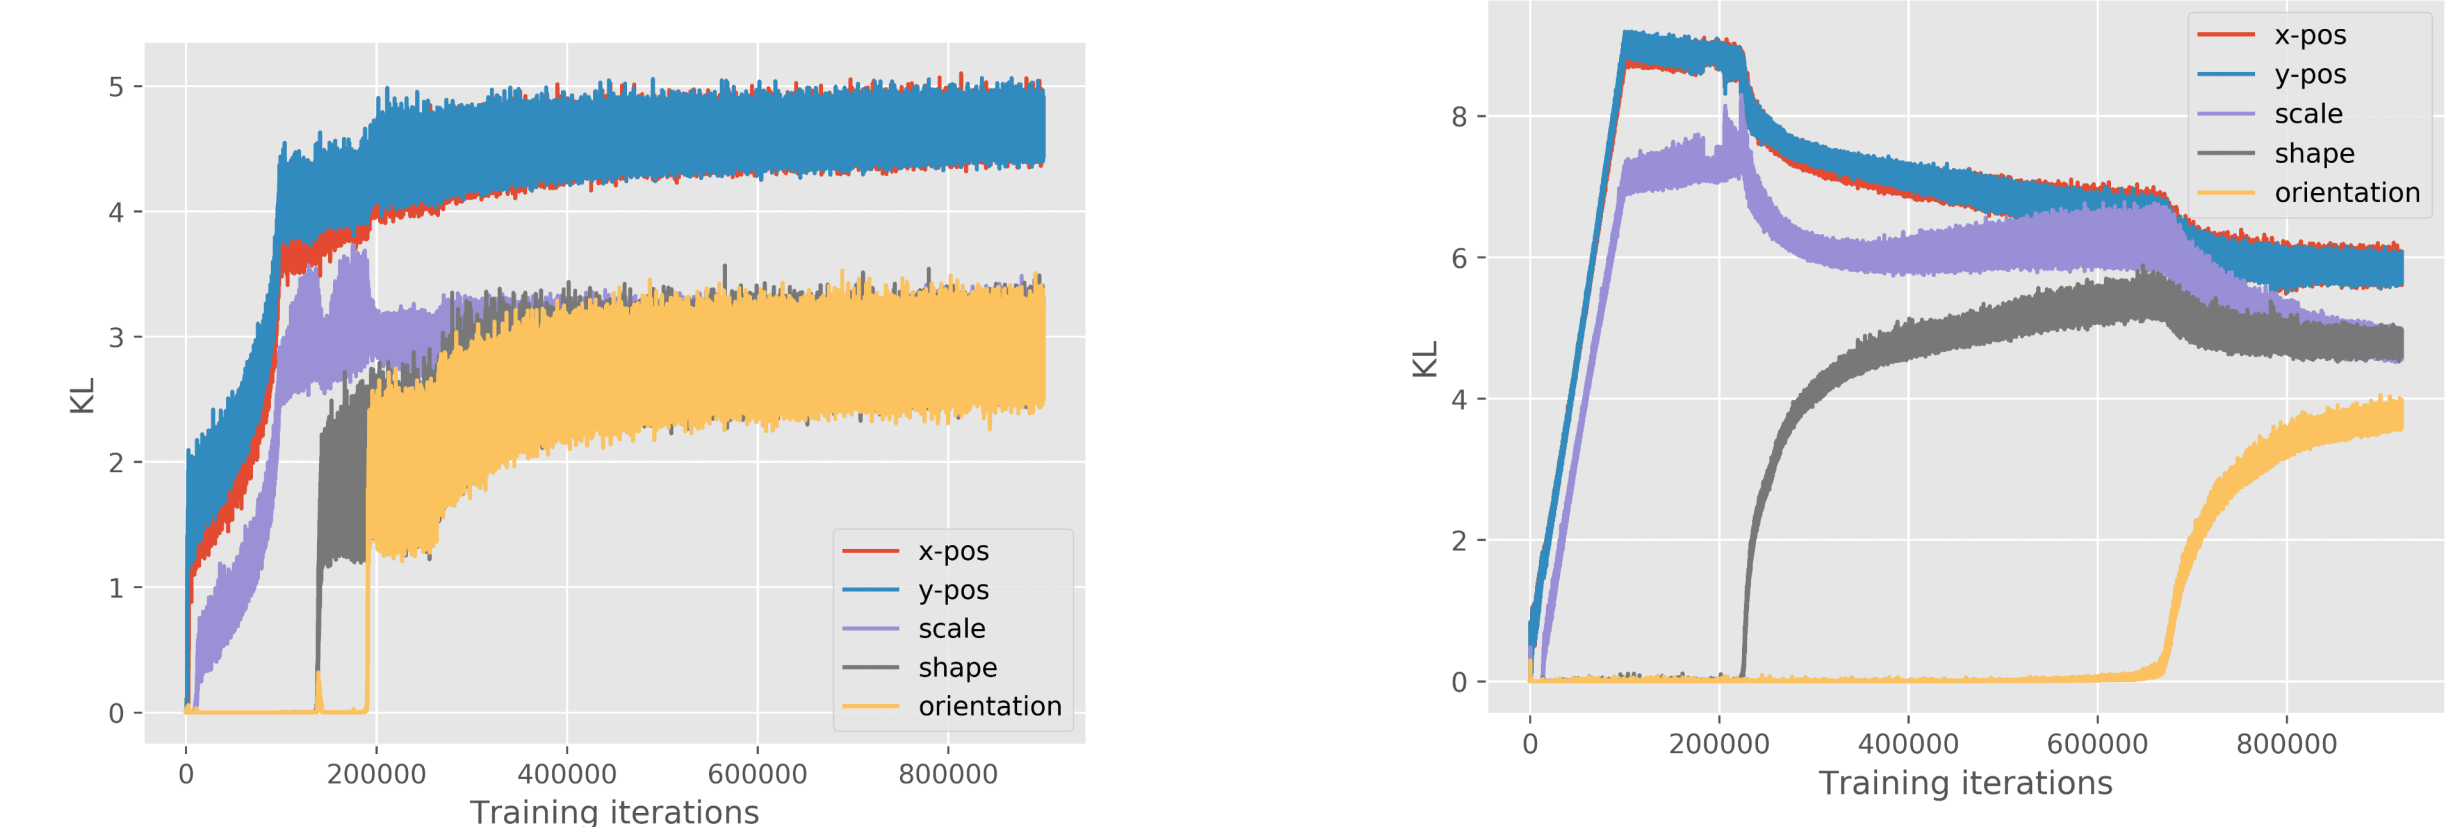
- Beta-VAEs try to find the optimum distortion(D) and rate(R) for a fixed $\beta = \frac{\partial D}{\partial R}$ by optimizing $\min_{q_\phi(\mathbf{z}|\mathbf{x}), p_\theta(\mathbf{x}|\mathbf{z}), p_\theta(\mathbf{x}|\mathbf{z})} D + \beta R$
- Bottleneck-VAEs (with $\gamma = 1$) with $C = R_m$ optimize distortion for a constant rate R_m
- Different increasing values of C, corresponds to relaxing the constraint variational posterior needs to be closer in terms of KL-divergence to the prior over the latent variables
- If we want to replicate the effects of linearly increasing C in case of Bottleneck VAEs, a Beta-VAE can be trained with monotonically decreasing β from $\beta \gg 1$ to $\beta \ll 1$.




- When compared to Bottleneck-VAEs, a linearly decreasing schedule of β in Beta-VAEs (which we call Beta-Annealed VAEs) offers advantages such as: 1) without having to set C, our proposed schedule have one less hyperparameter to tune; 2) in all of our experiments, we linearly decreasing from $\beta \gg 1$ to $\beta = 1$ during training, which can be interpreted as Beta-VAEs are trained as vanilla VAEs during later stages in training leading to better reconstruction error.




Results



KL-divergence of each latent dimension with respect to a unit Gaussian during training on dSprites. **(Left)** In Beta-Annealed VAE, β is decreased as the training progresses **(Right)** In Bottleneck-VAE, C is increased as the training progresses to increase the information capacity



ELBO **(Left)** and reconstruction error **(Right)** for different hyperparameters in Bottleneck and Proposed VAEs on dSprites.



(First two rows) Original data and their reconstructions **(Other rows)** all 10 latent dimension traversals with captured attribute indicated in the sides. Greyed out rows indicate dead dimensions.

Results on unsupervised representation learning of non Gaussian noise transients that occur in gravitational wave detectors

Model	Accuracy
β -VAE	61.26%
Bottleneck-VAE	80.01%
Proposed-VAE	81.60%

References

- Alemi, A., Poole, B., Fischer, I., Dillon, J., Saurous, R. A., & Murphy, K. (2018, July). Fixing a broken ELBO. In International Conference on Machine Learning (pp. 159-168). PMLR.
- Burgess, C. P., Higgins, I., Pal, A., Matthey, L., Watters, N., Desjardins, G., & Lerchner, A. (2018). Understanding disentangling in Beta-VAE. arXiv preprint arXiv:1804.03599.
- 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