Structured Uncertainty Prediction Networks

Garoe Dorta[1,2] Sara Vicente[2] Lourdes Agapito[3] Neill D.F. Campbell[1] Ivor Simpson[2]
[1]University of Bath [2]Anthropics Technology Ltd. [3]University College London

CVPR 2018 SALT LAKE CITY • JUNE 18-22

Introduction

We are interested in modeling output uncertainty in generative DNN models.

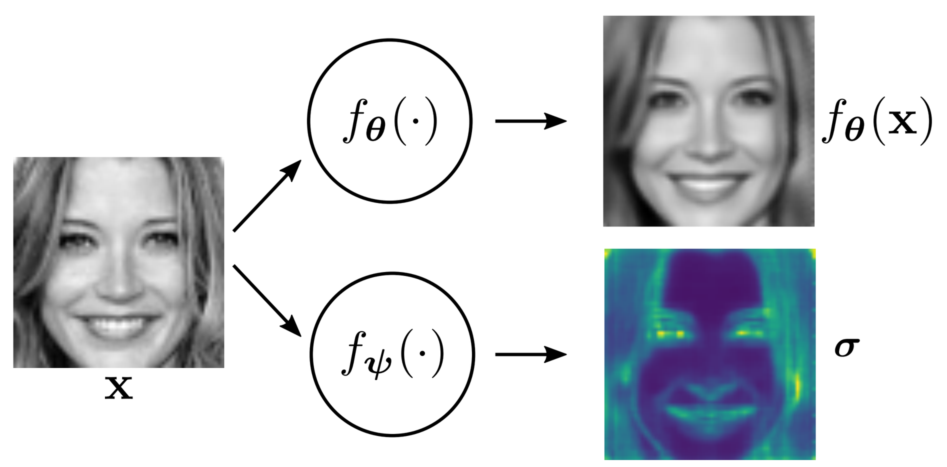

Predictive uncertainty can be useful in image editing or model-based decision making.

It is common [1, 2, 3] to use a factorized noise model, $\mathbf{y} = f_\theta(\mathbf{x}) + \epsilon$, where $\epsilon \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$.

For images, the factorized assumption does not hold, so we propose to use structured uncertainty, $\epsilon \sim \mathcal{N}(\mathbf{0}, \Sigma)$.

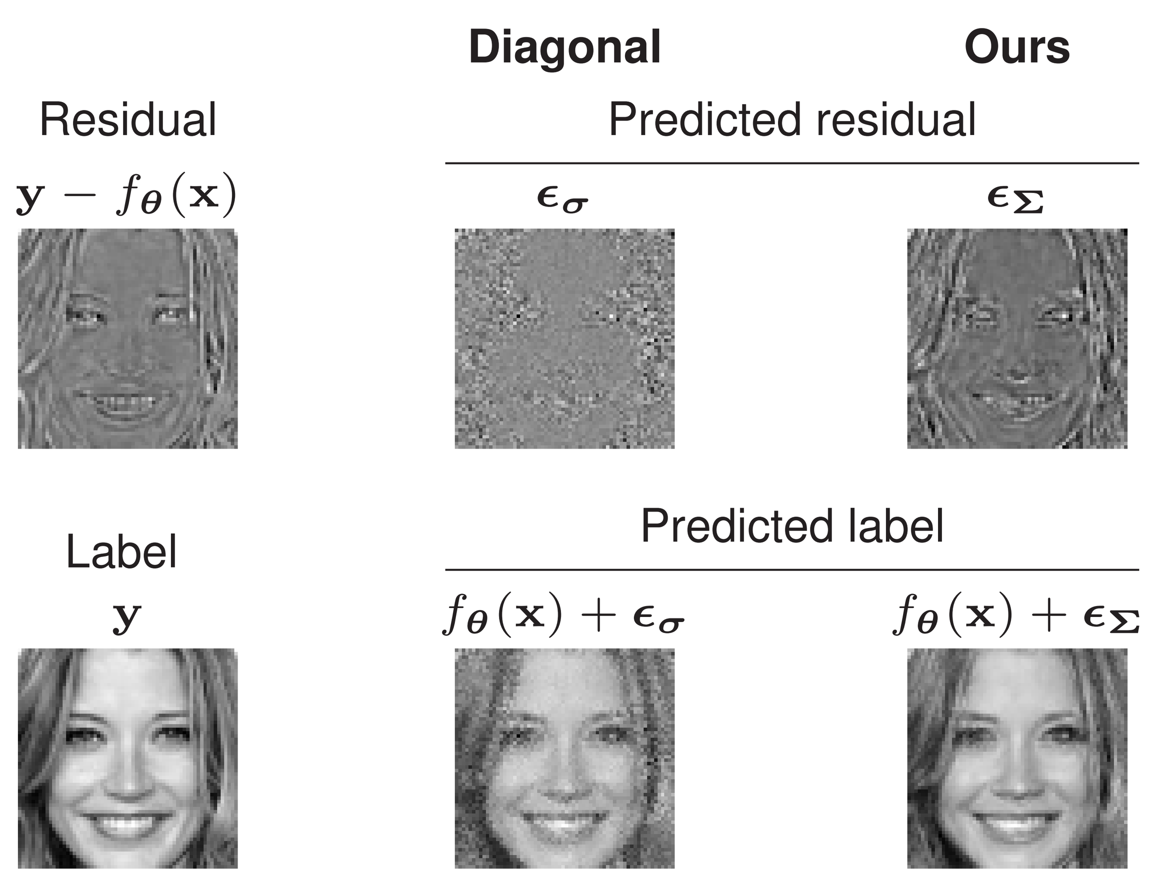

The estimation of the residual covariance matrix, Σ, is complicated by only having a single residual image, ϵ, per input, $\mathbf{x}$.

Methodology

We design a structured uncertainty prediction network for a pretrained VAE [1] model:

$$p_\theta(\mathbf{x} \mid \mathbf{z}) = \mathcal{N}\left(\mu(\mathbf{z}), \Sigma_\psi(\mathbf{z})\right).$$

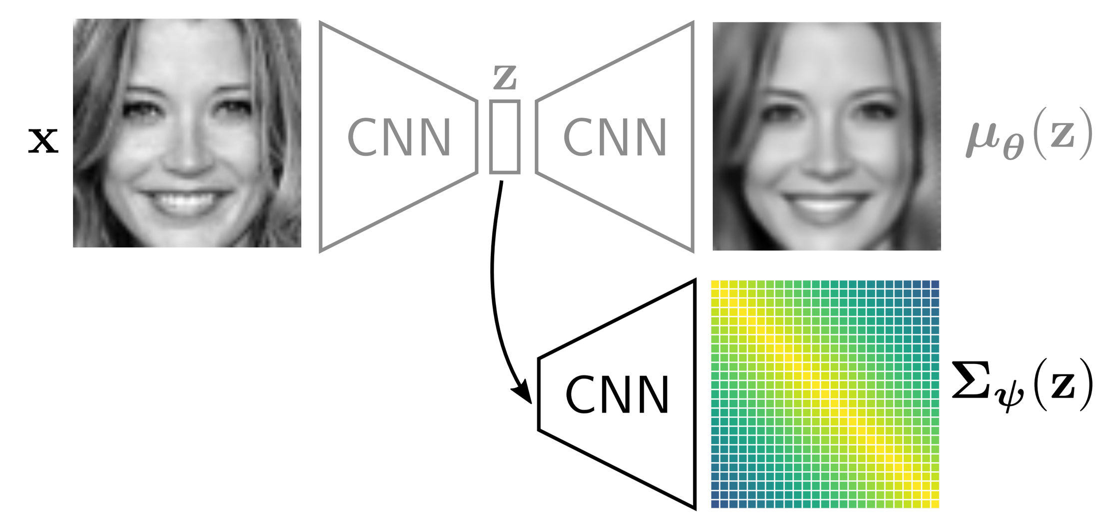

The model is trained via maximum likelihood:

$$\min_\psi \ \log\left(|\Sigma_\psi|\right) + (\mathbf{x} - \mu)^{\intercal} (\Sigma_\psi)^{-1} (\mathbf{x} - \mu).$$

Inference and learning is made tractable by imposing sparsity on the matrix $\mathbf{L}$, where $\Sigma^{-1} = \mathbf{L}\mathbf{L}^{\intercal}$ is a Cholesky decomposition.

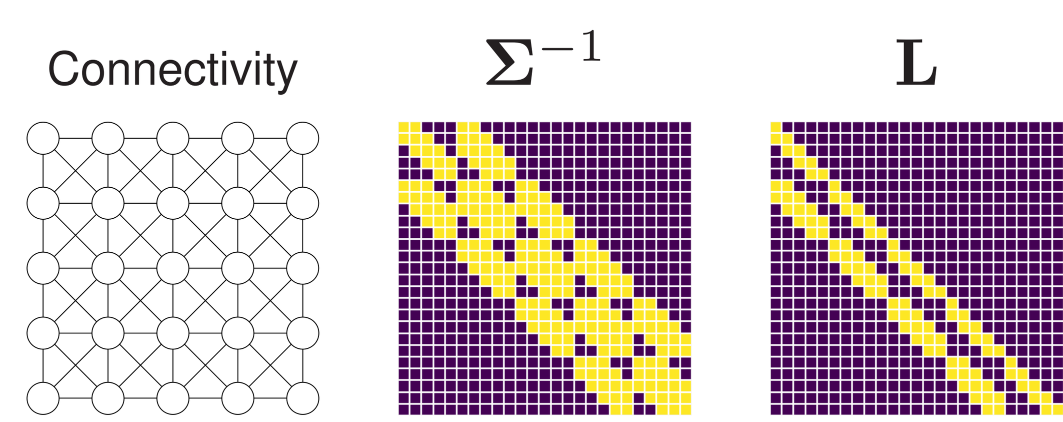

Efficient convolutional implementation

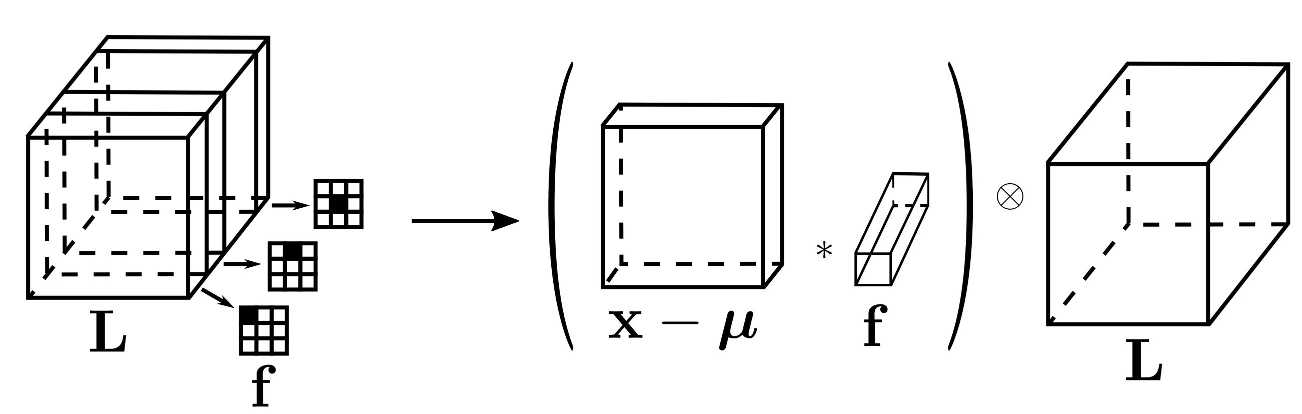

Reconstructions

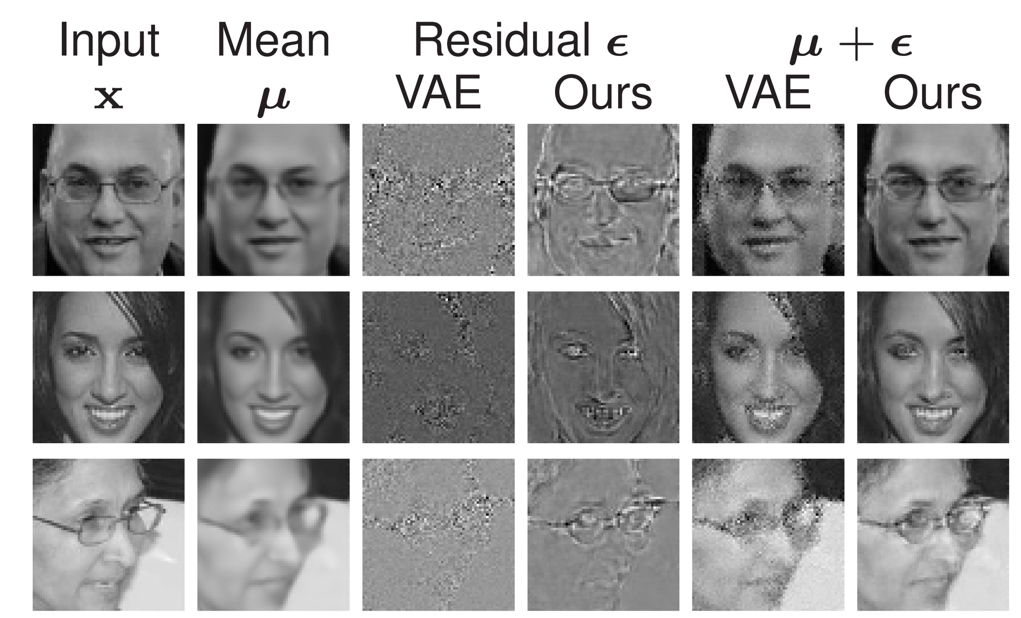

Model	NLL	$-\log p(\mathbf{x} \mid \mathbf{z})$
VAE [1]	-5378 ± 931	-6079 ± 936
Ours	-7753 ± 1323	-8386 ± 1339

Samples

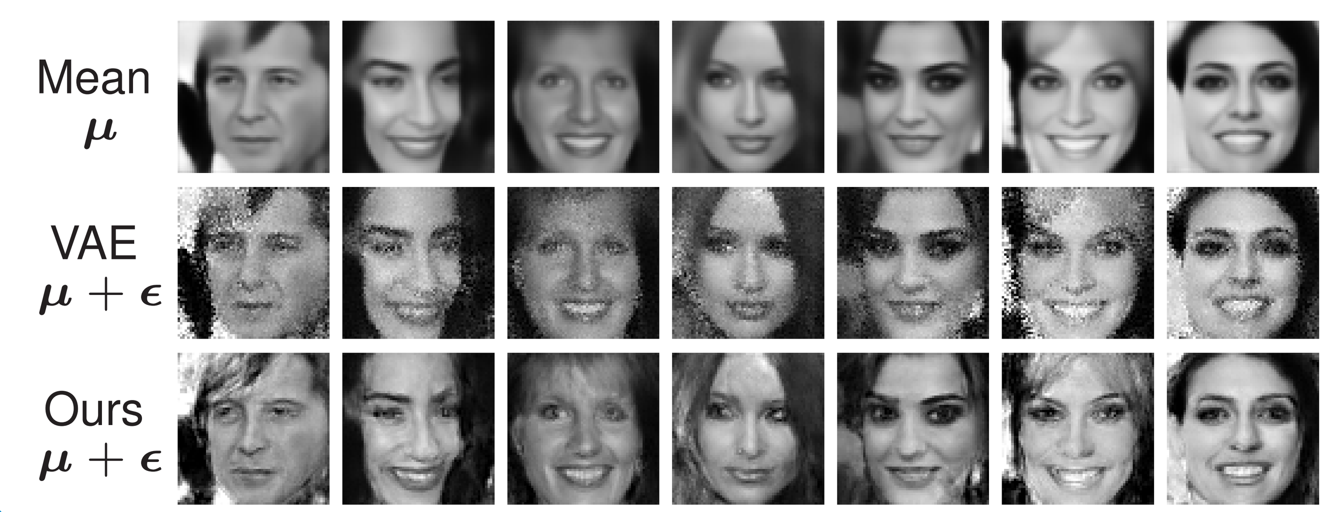

Interpolations

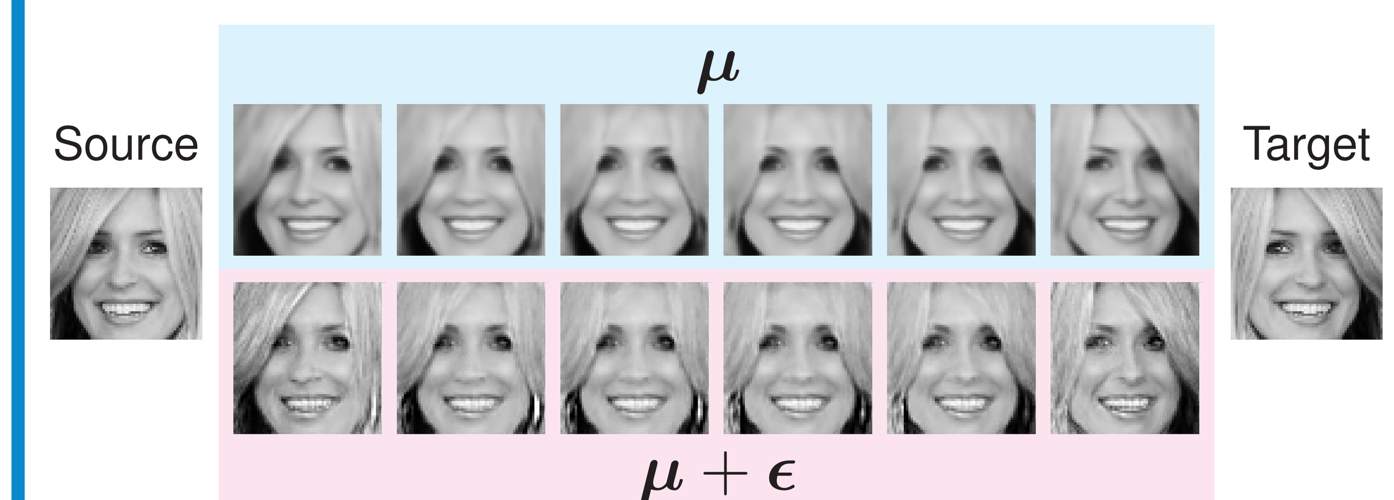

Samples while interpolating on the latent space $\mathbf{z}$ from the source to the target image.

Application: Denoising

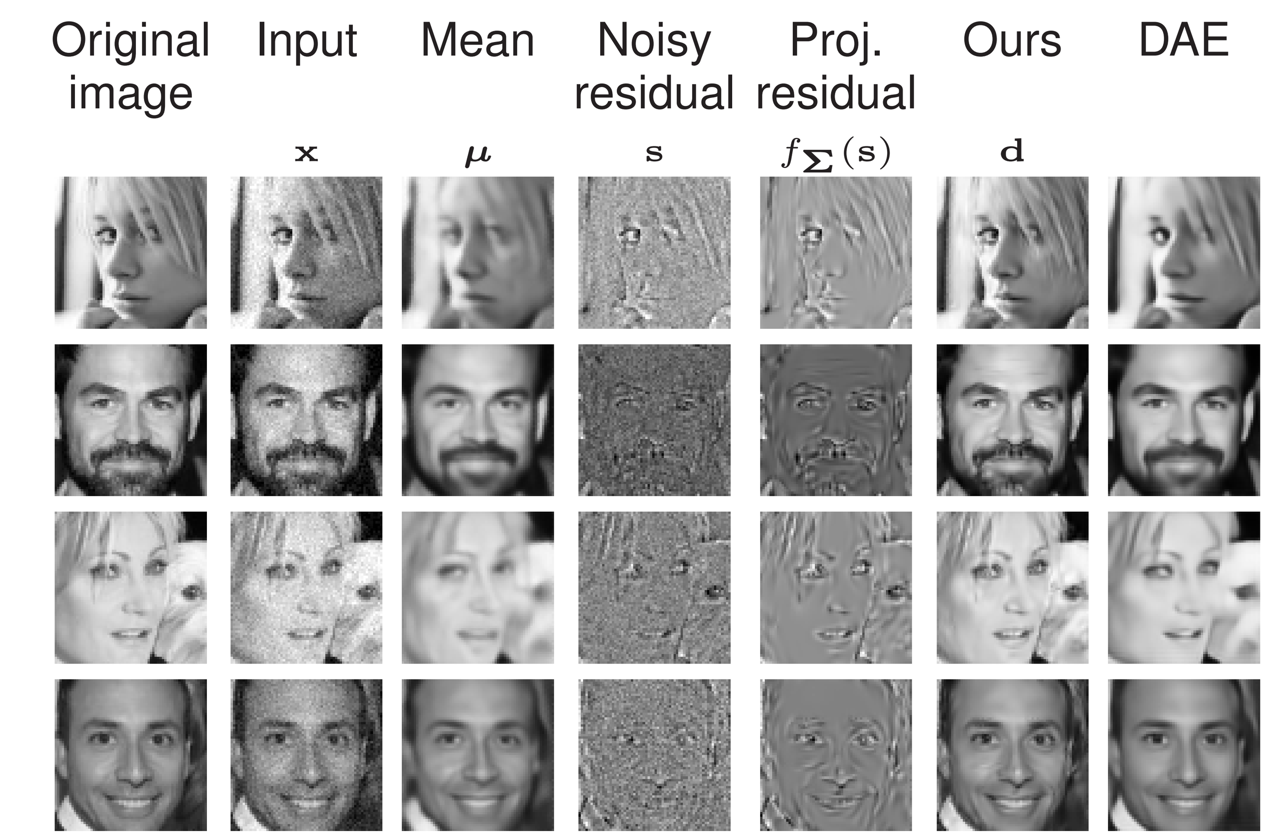

Denoising by recovering a more likely residual under $\mathcal{N}(\mathbf{0}, \Sigma)$, by projecting the noisy residual onto the principal eigenvectors of Σ.

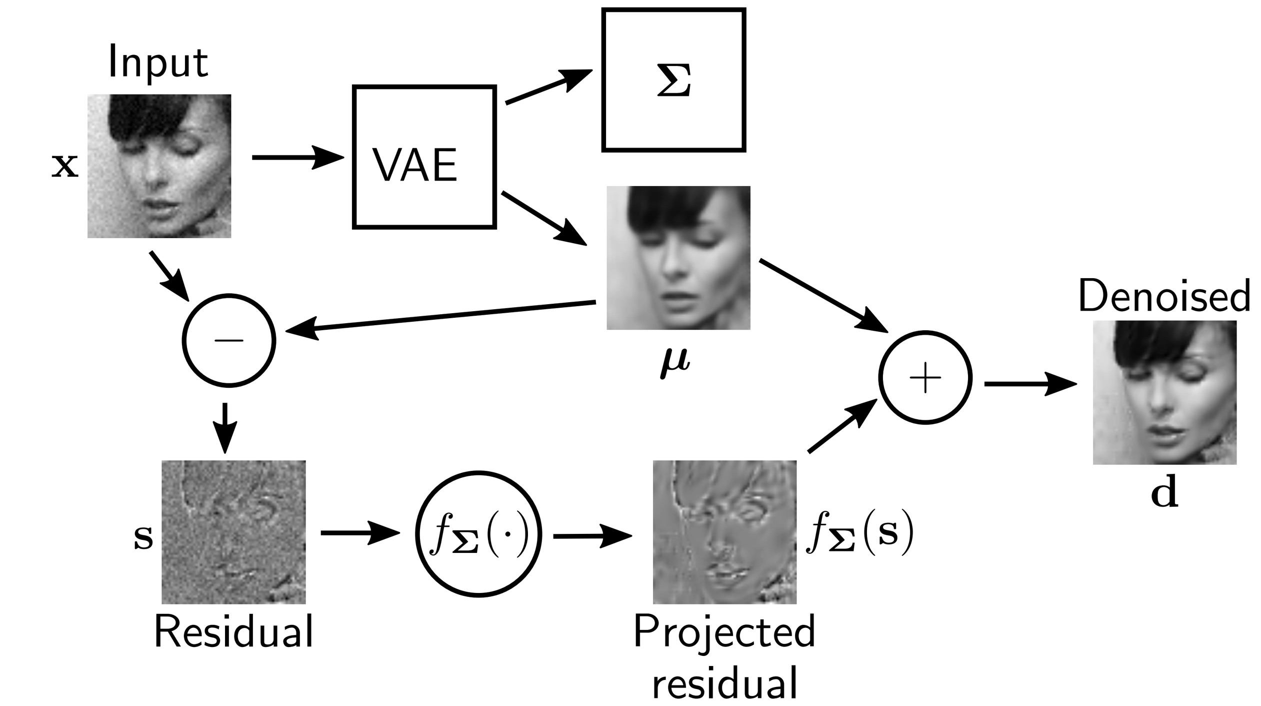

Model	MSE
DAE	5.13e-3 ± 2.52e-3
Ours	2.99e-3 ± 7.98e-4

References

[1] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *ICLR*, 2014.
[2] Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and approximate inference in deep generative models. *ICML*, 2014.
[3] 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. In *ICLR*, 2017.