VAEs and GANs

Mihaela Rosca
@elaClaudia

DeepMind

Thanks to Shakir Mohamed, Balaji Lakshminarayanan

What are the biggest problems with training GANs?

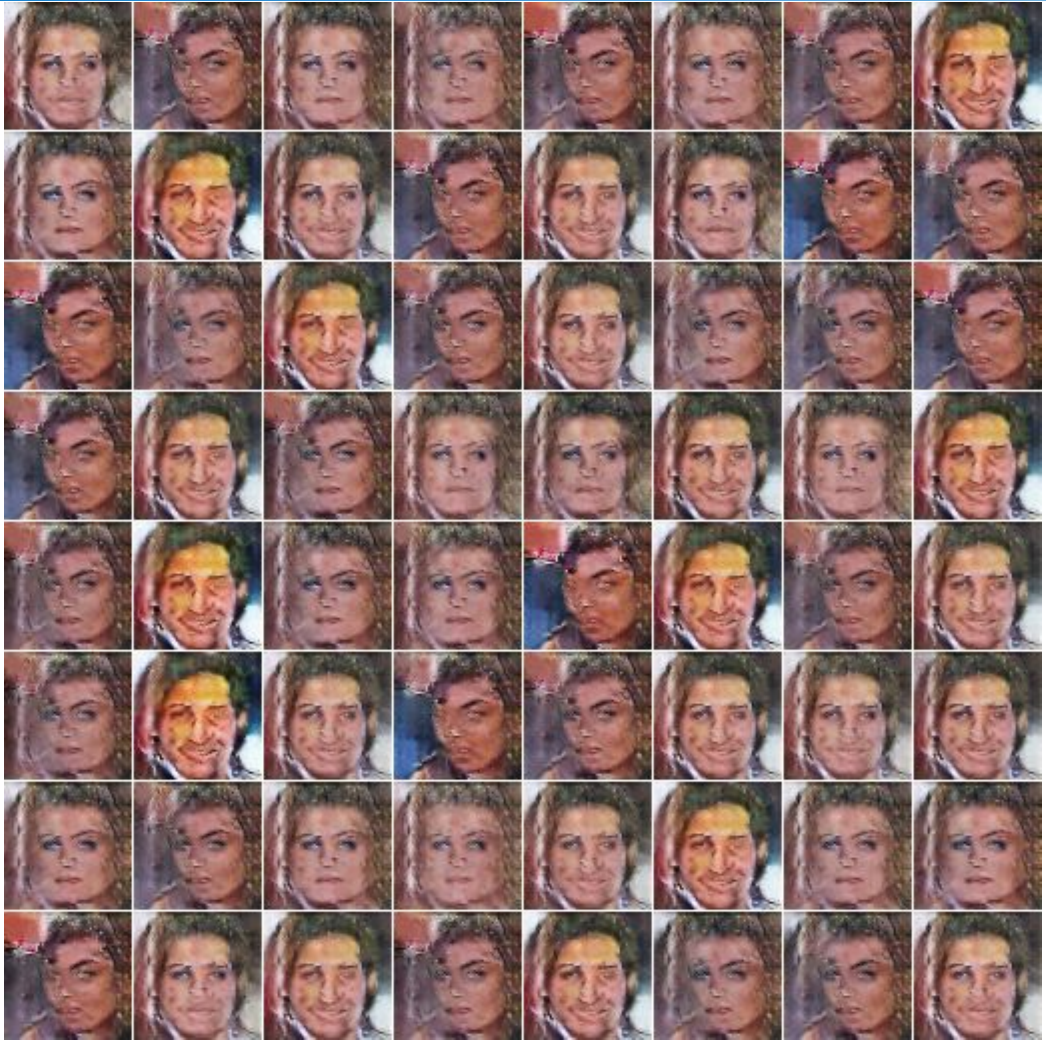

Maximum likelihood

Find the model which gives highest likelihood to the data.

$$\text{argmax}_\theta\ \mathcal{E}_{x \sim p^*} \log p_\theta(x)$$

Maximum likelihood

Find the model which gives highest likelihood to the data.

$$\text{argmax}_\theta \, \mathcal{E}_{x \sim p^*} \log p_\theta(x)$$

Maximum likelihood with latent variables

Leverage underlying data structure in generative process.

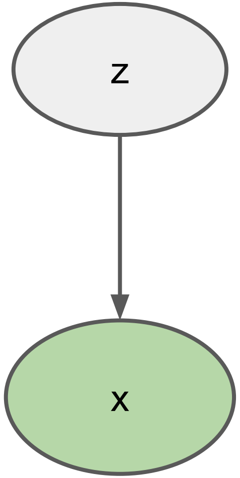

z = latents
x = observed

$$p_{\boldsymbol{\theta}}(\mathbf{x}) = \int p_{\boldsymbol{\theta}}(\mathbf{x}|\mathbf{z})p(\mathbf{z})d\mathbf{z}$$

Maximum likelihood with latent variables

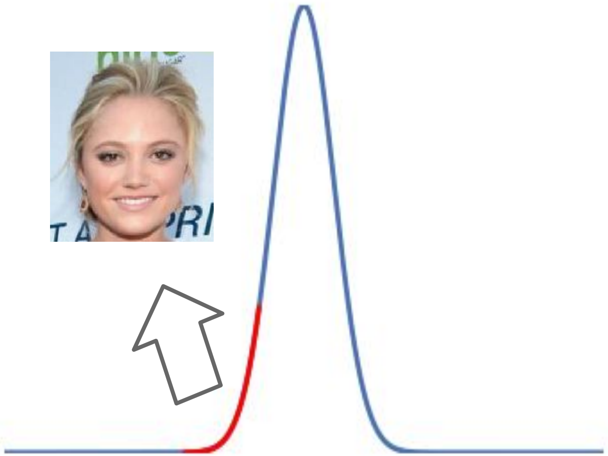

Maximum likelihood with latent variables

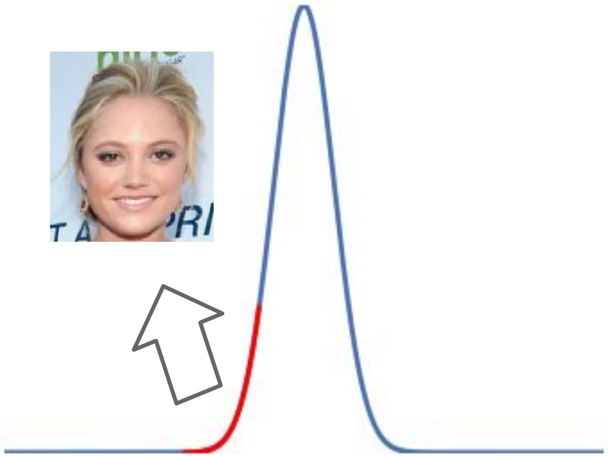

$$q(z|\boldsymbol{x})$$

Maximum likelihood with latent variables

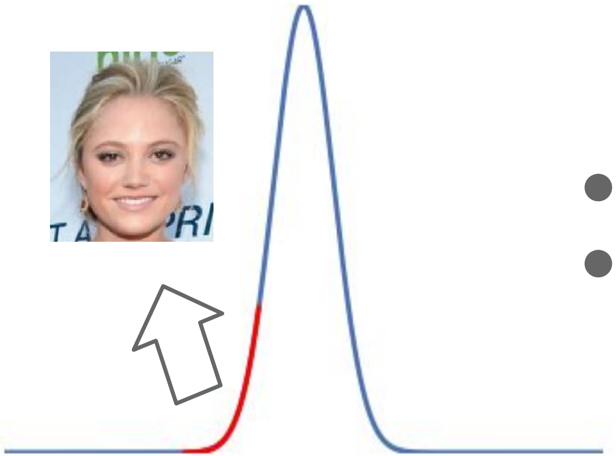

$q(z|\boldsymbol{x})$

- explain x
- stay close to p(z)

Variational autoencoders

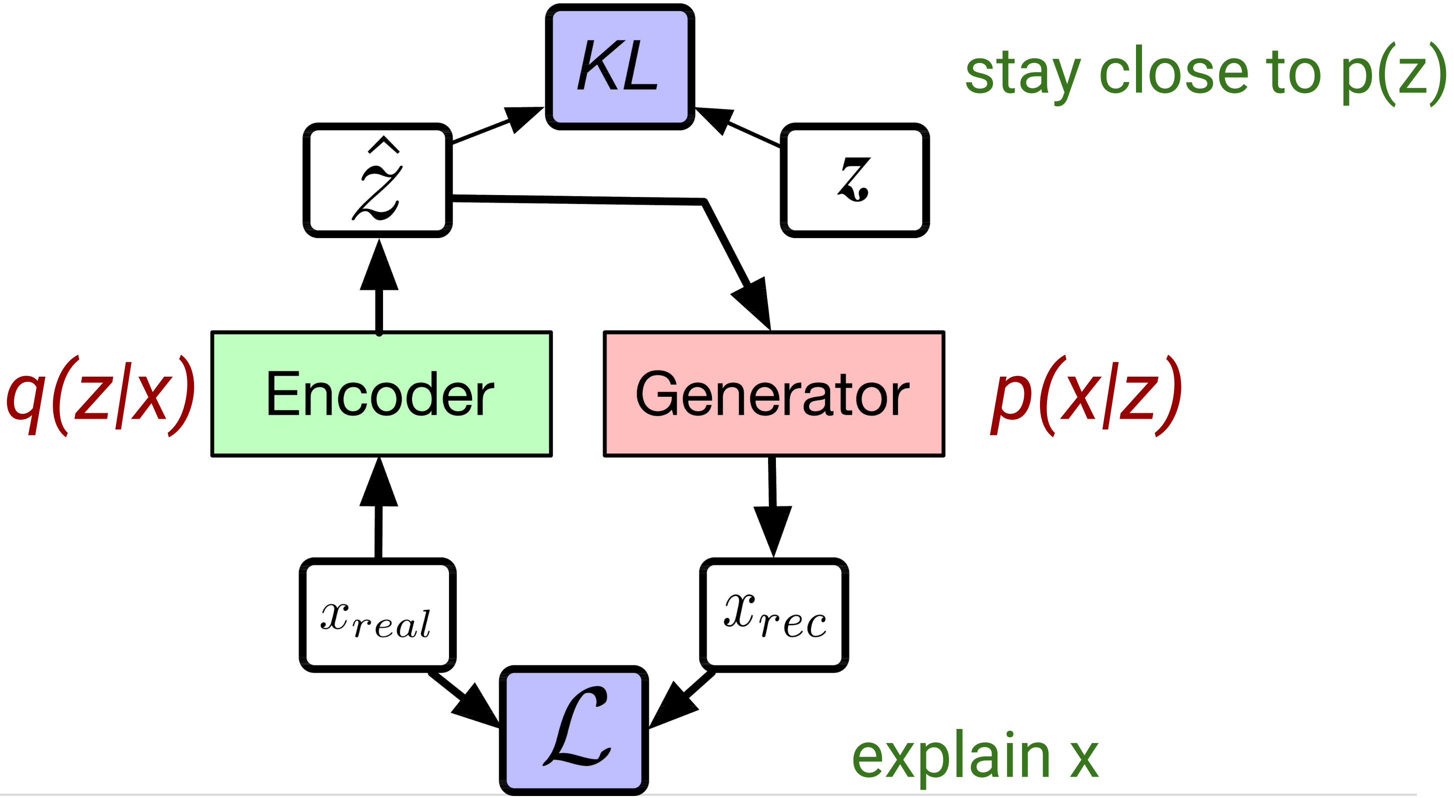

Variational autoencoders

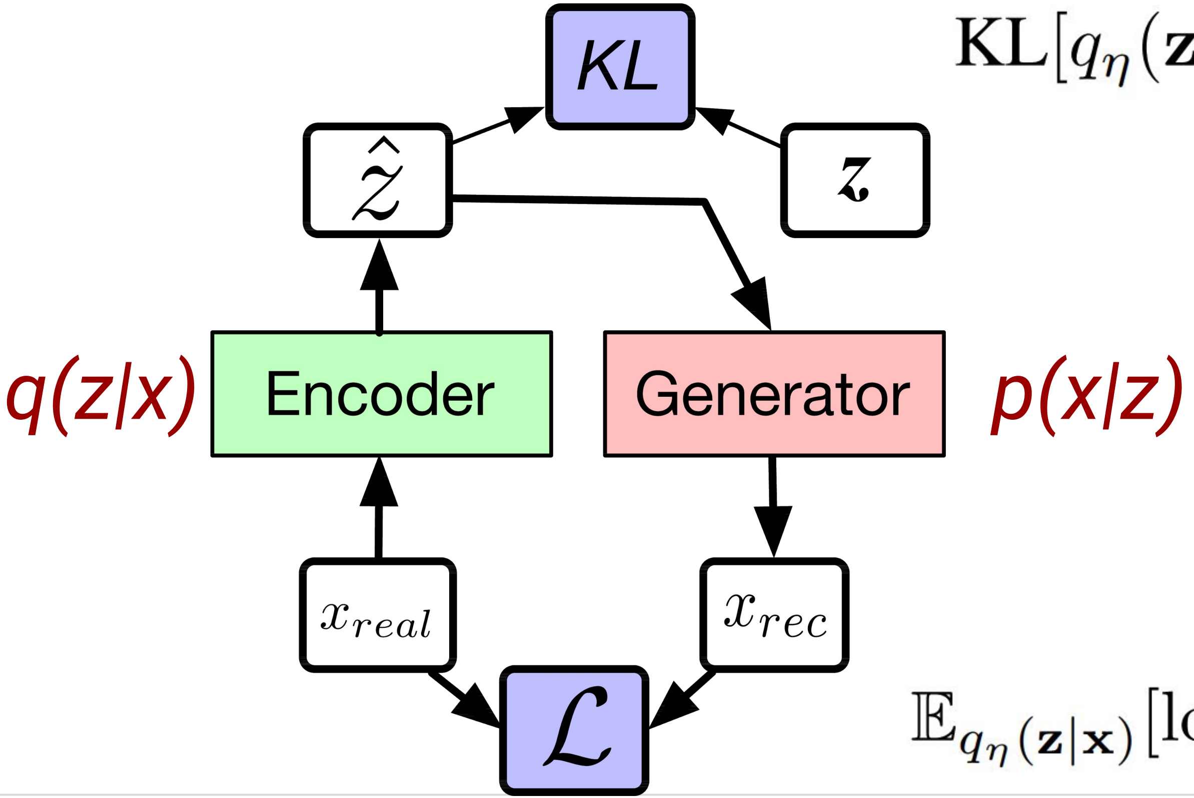

$$\mathrm{KL}[q_\eta(\mathbf{z}|\mathbf{x})||p(\mathbf{z})]$$

$$\mathbb{E}_{q_\eta(\mathbf{z}|\mathbf{x})}[\log p_{\boldsymbol{\theta}}(\mathbf{x}|\mathbf{z})]$$

D.P. Kingma, M. Welling: Auto-Encoding Variational Bayes

Variational autoencoders

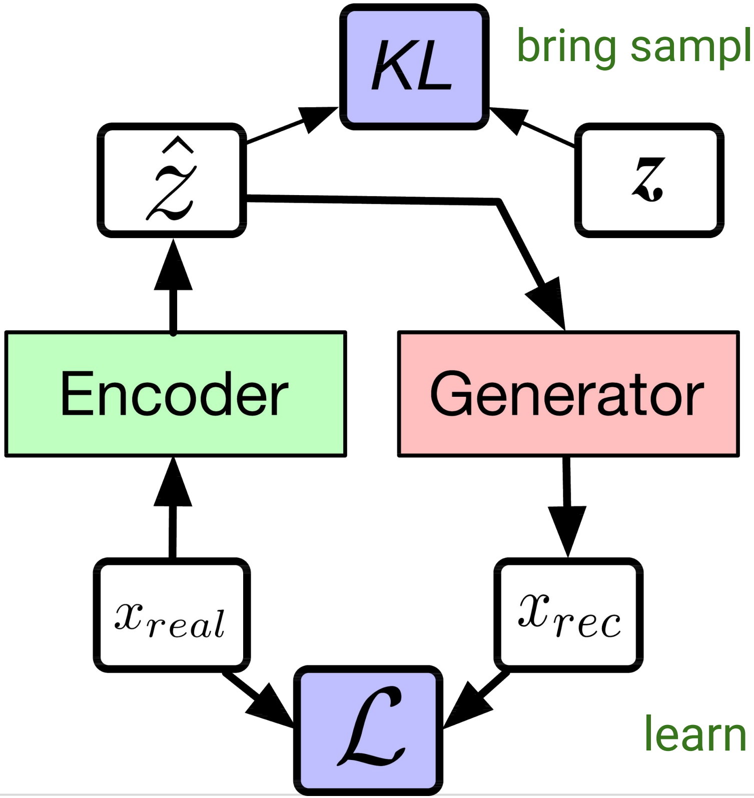

bring samples close to reconstructions

learn how to reconstruct

Inference

Learning distributions over representations.

Why:

- quantifying uncertainty
- imposing prior structure over learned representations

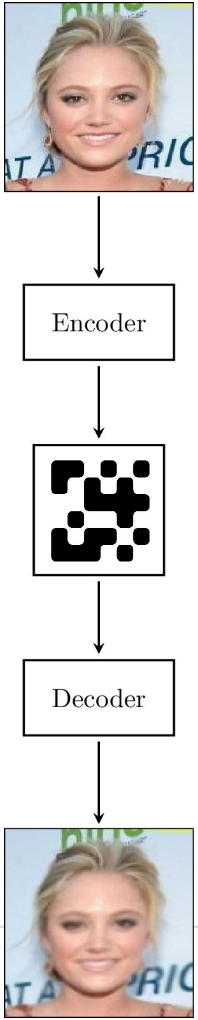

Imposing prior structure over representations

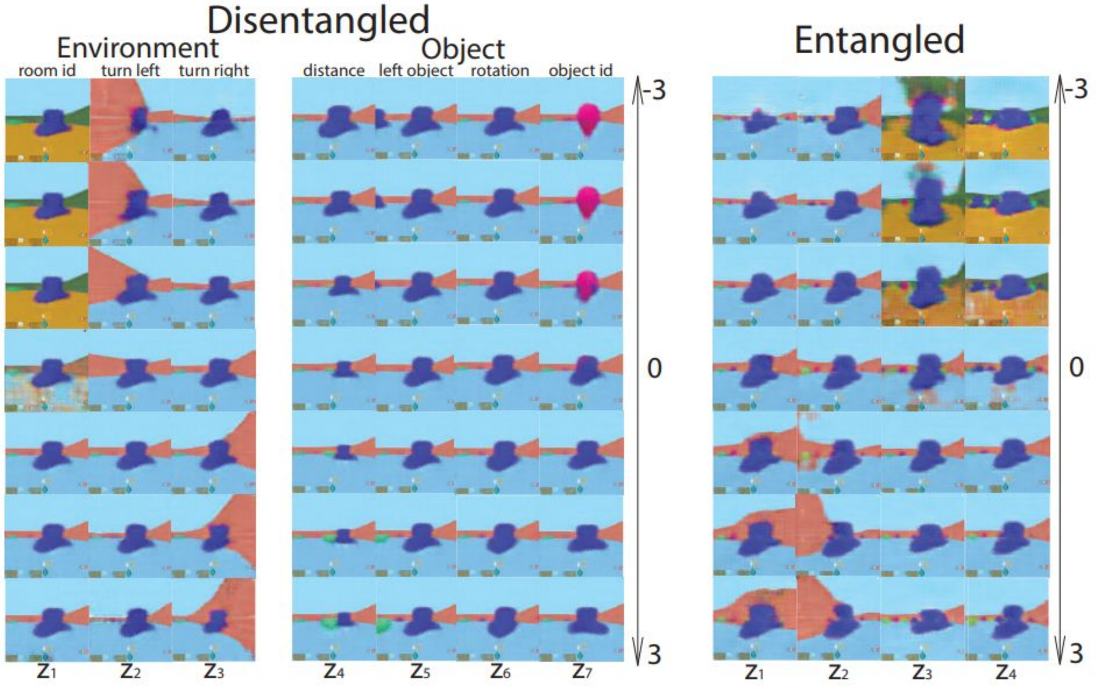

VAE distribution matching in visible space

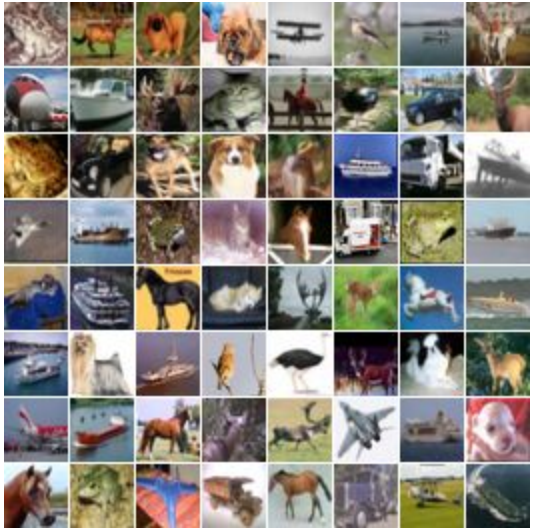

Data

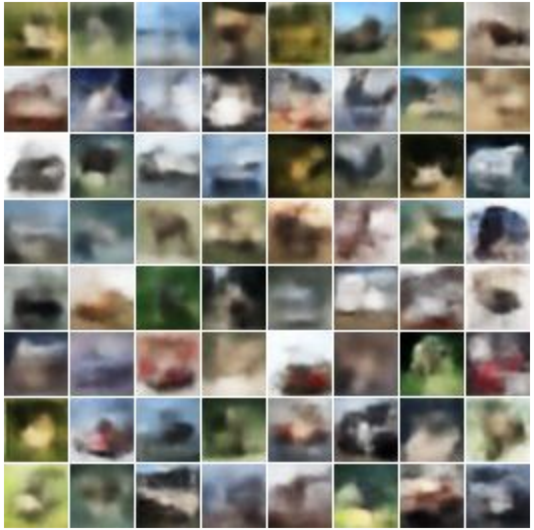

VAE samples

VAE distribution matching in latent space

Data

Low posterior VAE samples

VAEs match **marginal** distributions by matching **conditional** distributions.

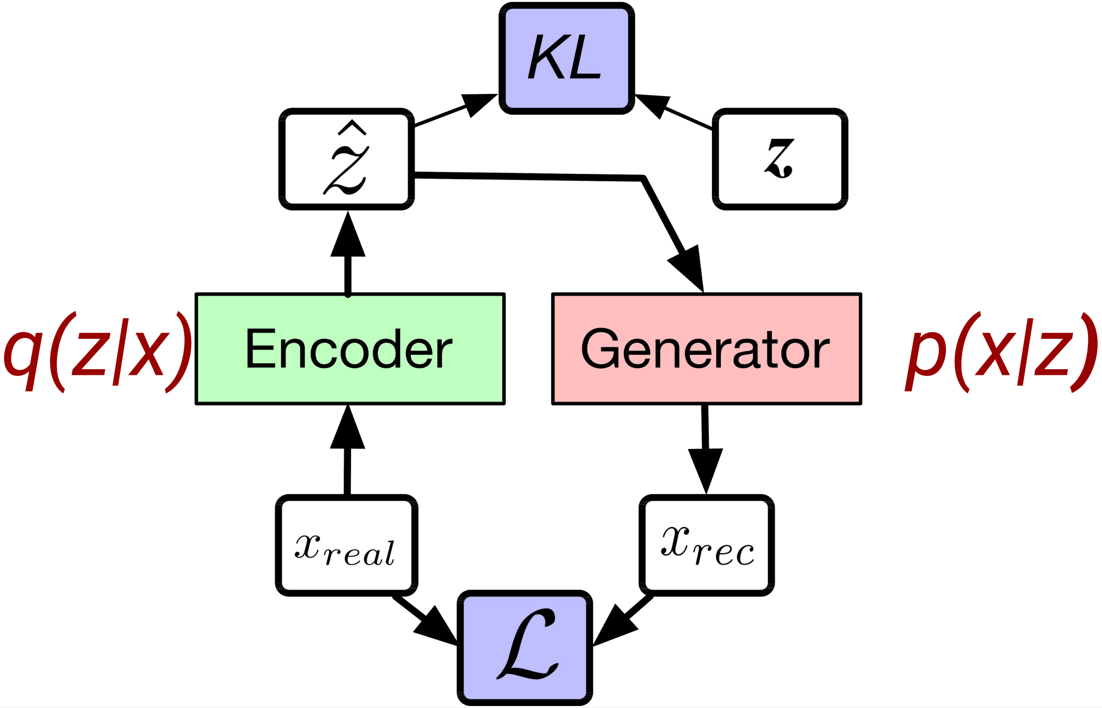

VAEs match **marginal** distributions by using **explicit** distributions.

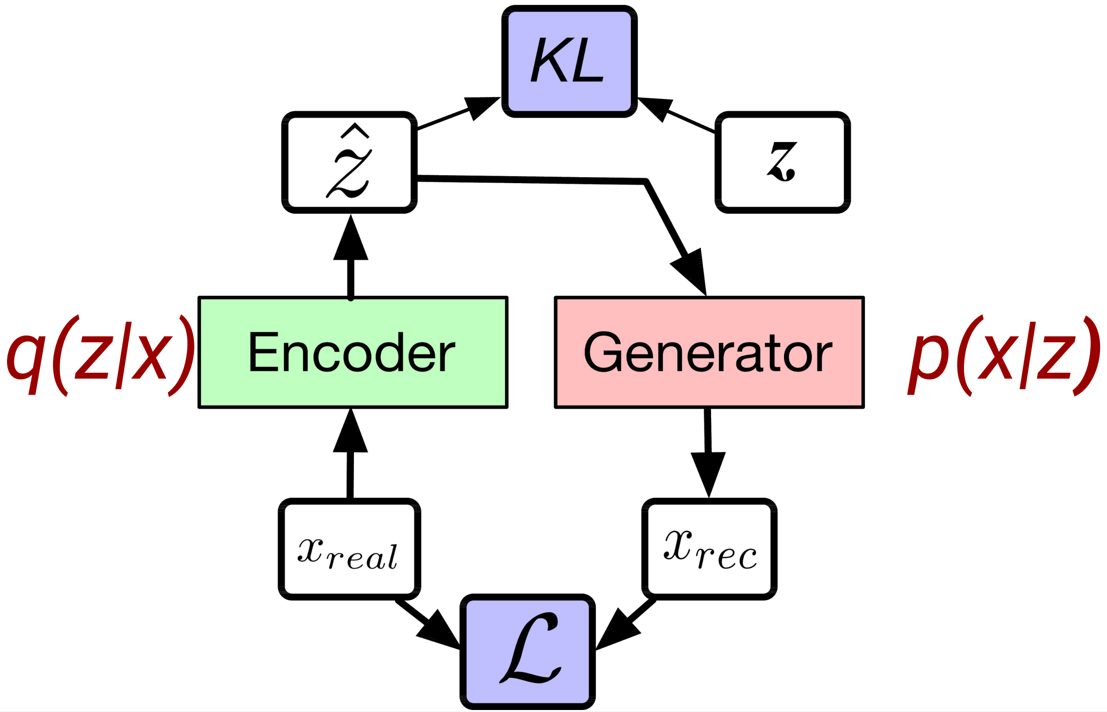

GANs

- Marginal distribution matching
- Implicit distributions

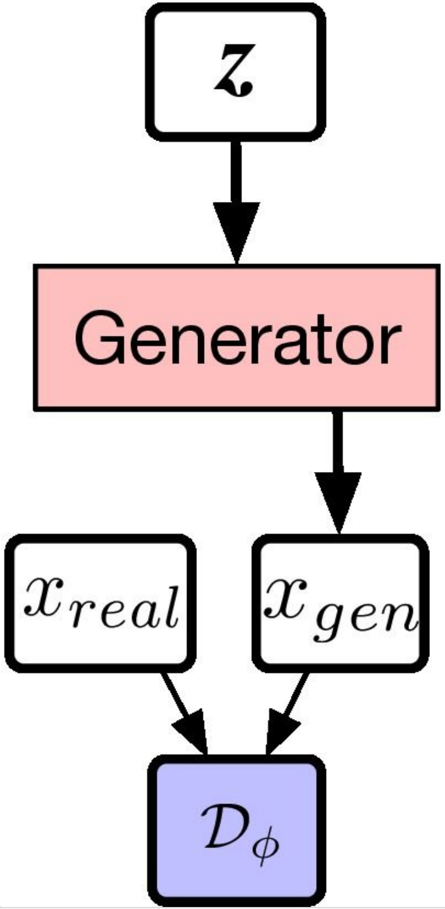

Combining GANs and VAEs

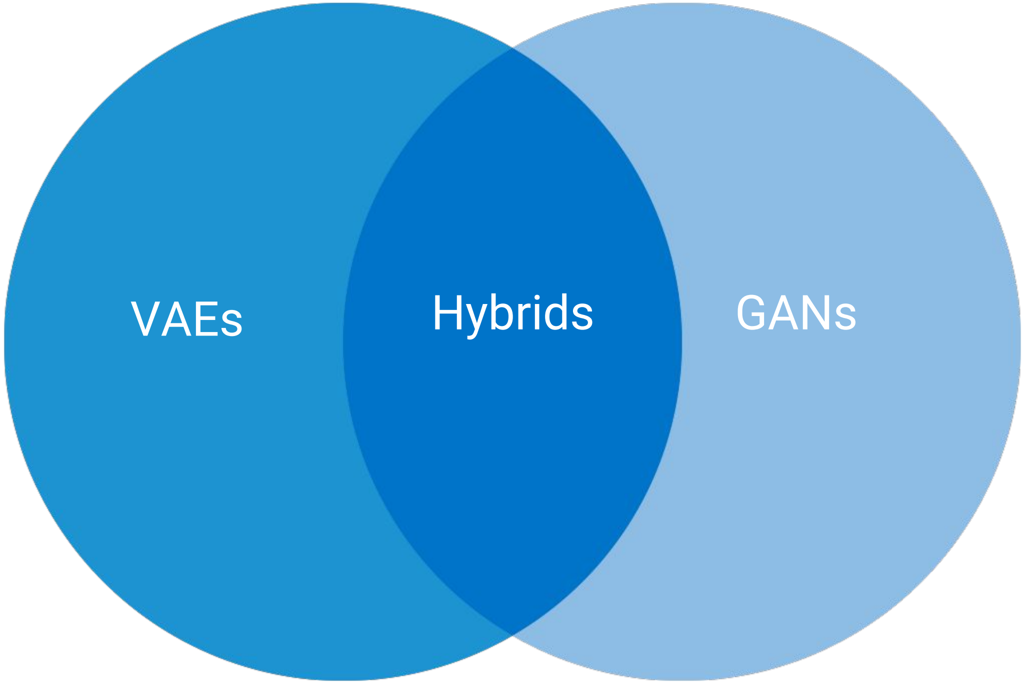

The promise of VAE-GAN hybrids

Improve
sample quality

Improve
representation learning

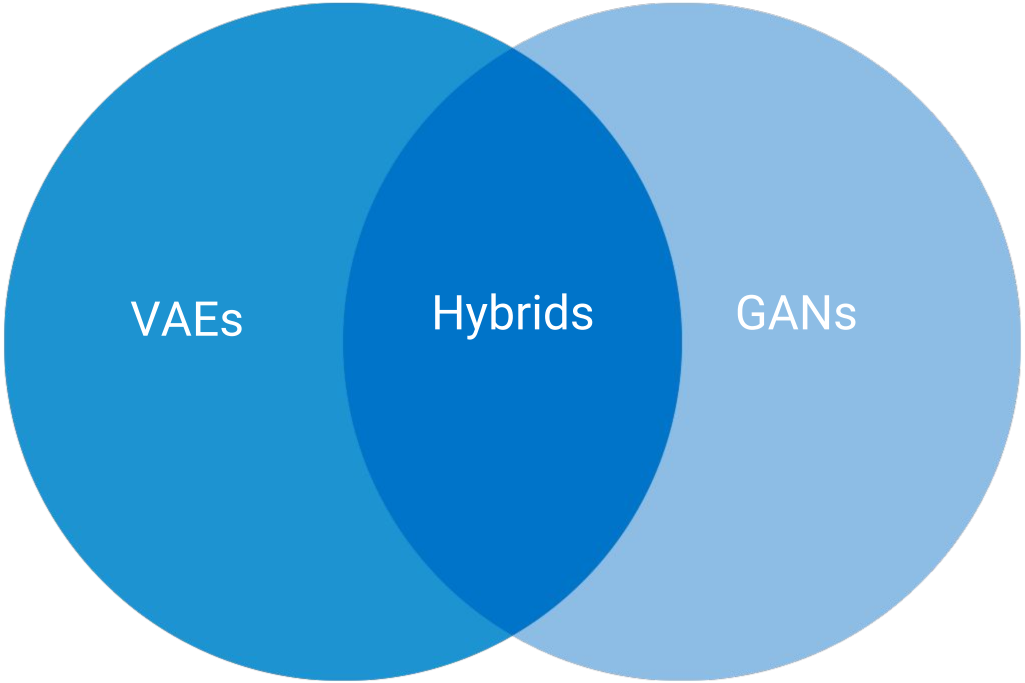

The promise of VAE-GAN hybrids

Improve **sample quality**

Improve **representation learning**

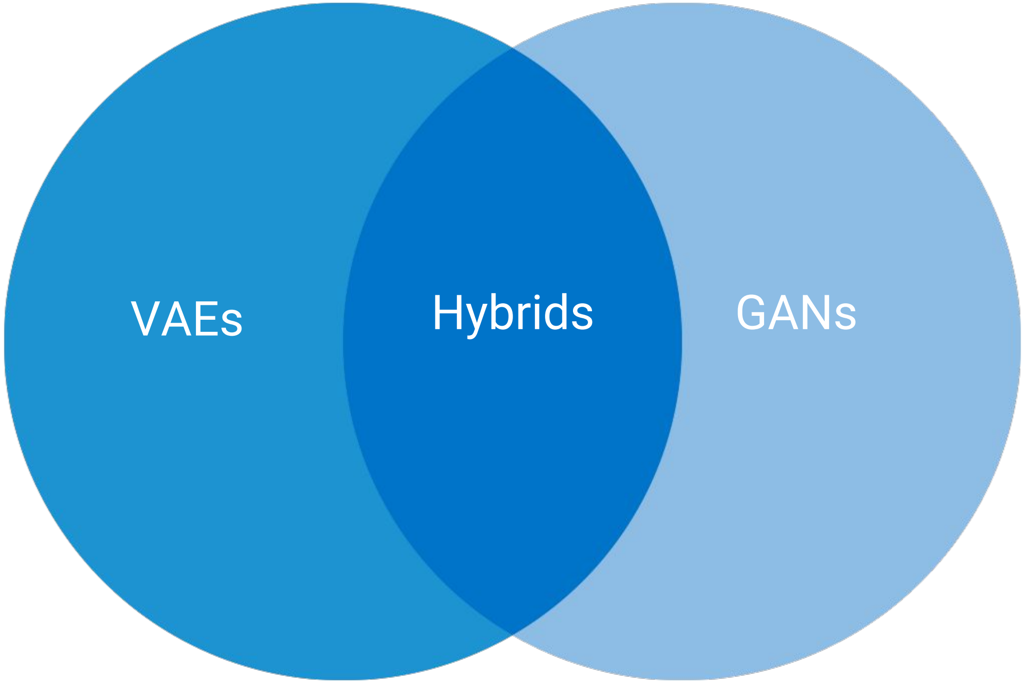

Improve **stability**

Improve **diversity**

VAE-GAN hybrids

- Adversarial Autoencoder
- Adversarial Variational Bayes
- VEEGAN
- ALI/BiGAN
- AlphaGAN
- ...

VAE-GAN hybrids via density ratios

Estimate the ratio of two distributions only from samples, by building a binary **classifier** to distinguish between them.

$$\frac{p(\mathbf{x})}{p_{\boldsymbol{\theta}}(\mathbf{x})} = \frac{\mathcal{D}(x)}{1 - \mathcal{D}(x)}$$

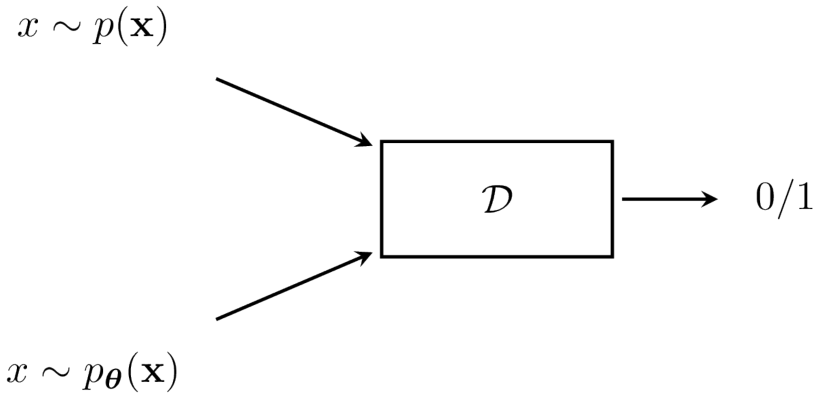

Do VAE-GAN hybrids improve inference?

Adversarial autoencoders

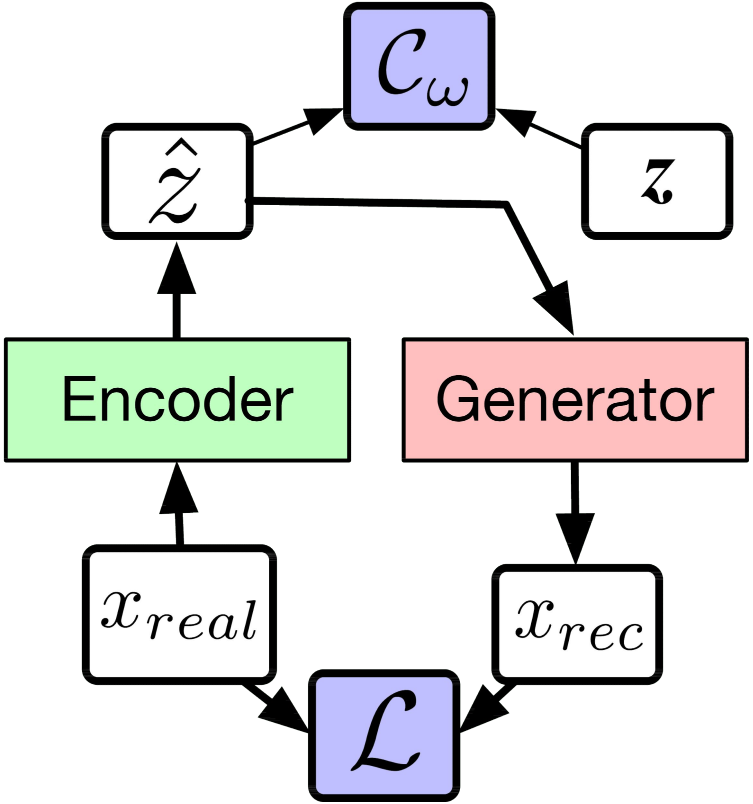

Replace KL with a discriminator matching marginal distributions

Marginal distribution matching in latent space. Implicit encoder distribution.

The effect of adversarial training on bounds

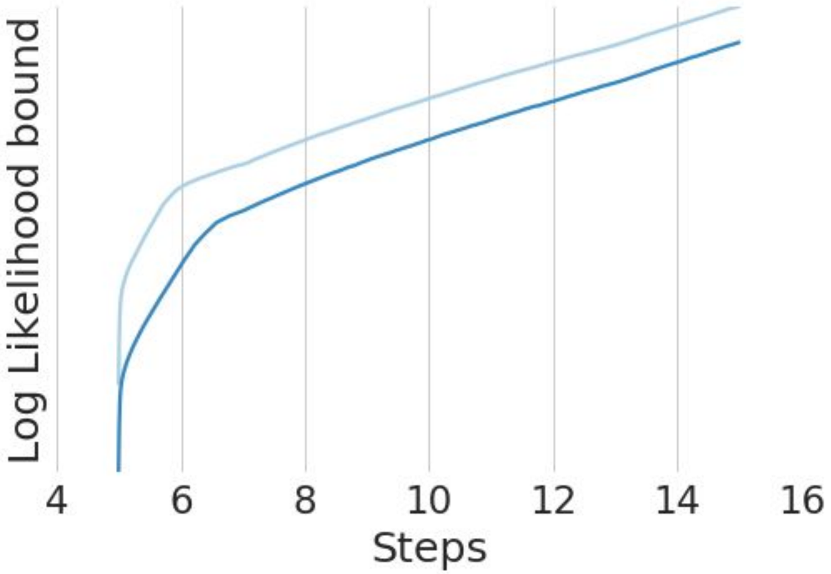

VAE

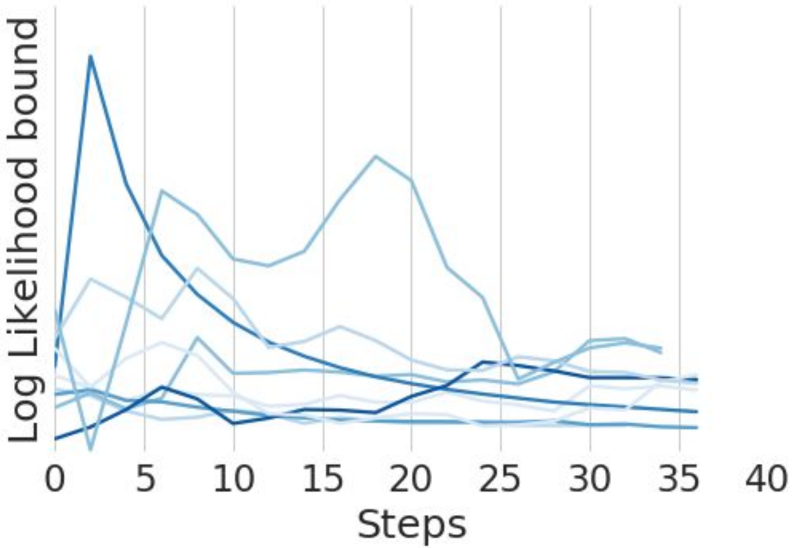

AAE

Classifier probabilities can be used for learning, but not for estimation.

The effect of adversarial training on representations

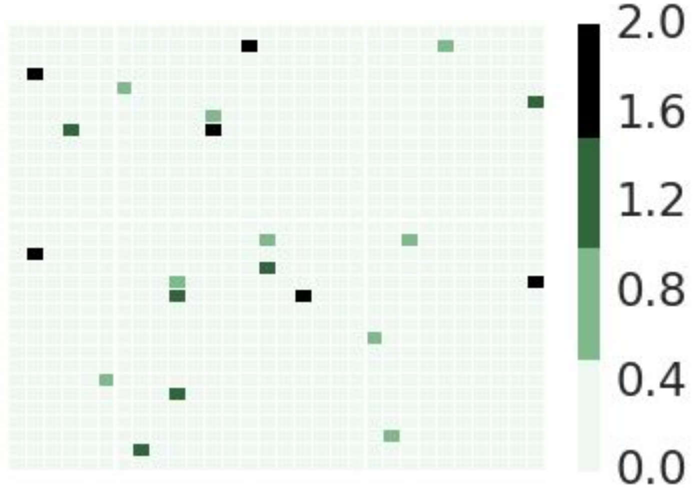

Learned VAE representations are sparse.

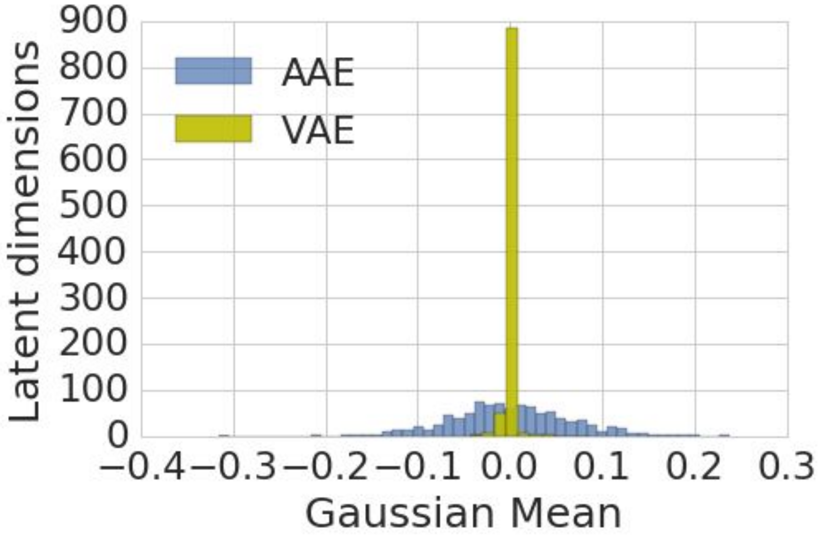

Learned AAE representations are not sparse.

Large latent sizes

VAE

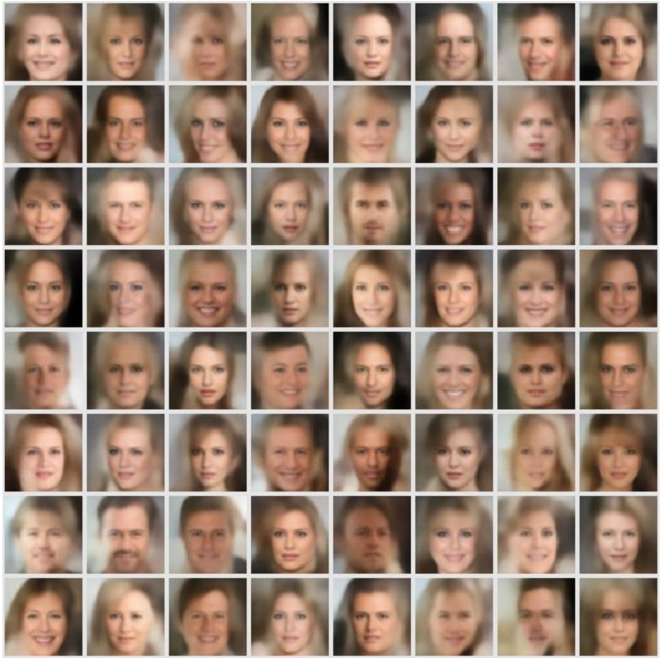

AAE

Do VAE-GAN hybrids improve generation?

VAE - GAN Hybrid (VGH)

Marginal matching and implicit distributions using GANs both in **latent and visible** space.

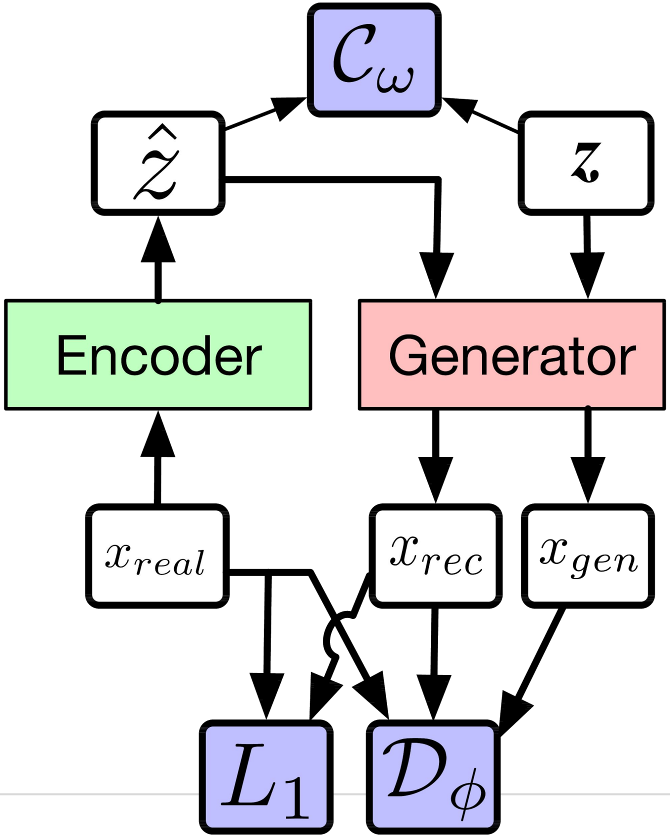

Joint space hybrids - VEEGAN

Directly match in **joint** space.

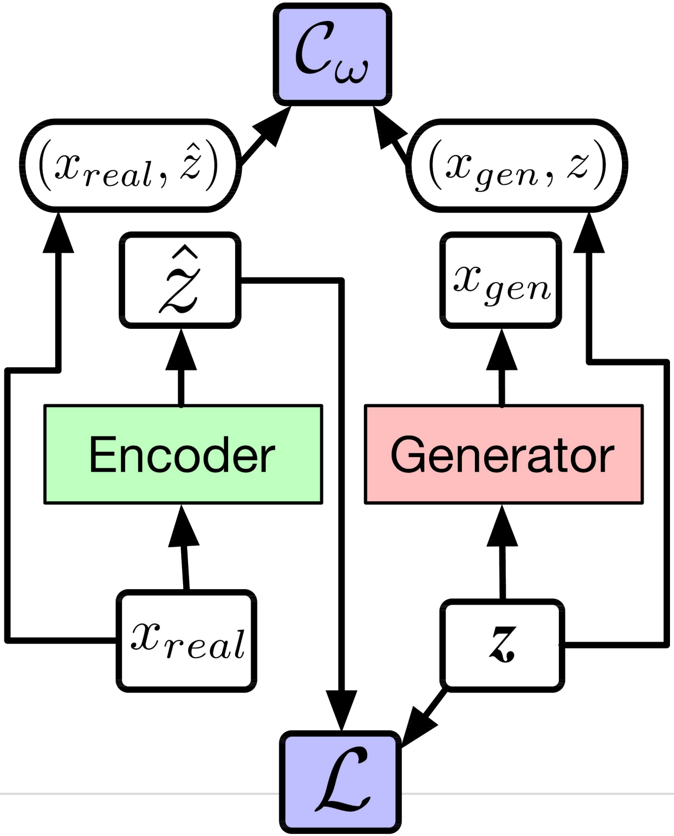

Srivastava et all: VEEGAN: Reducing Mode Collapse in GANs using Implicit Variational Learning

Improving GAN stability

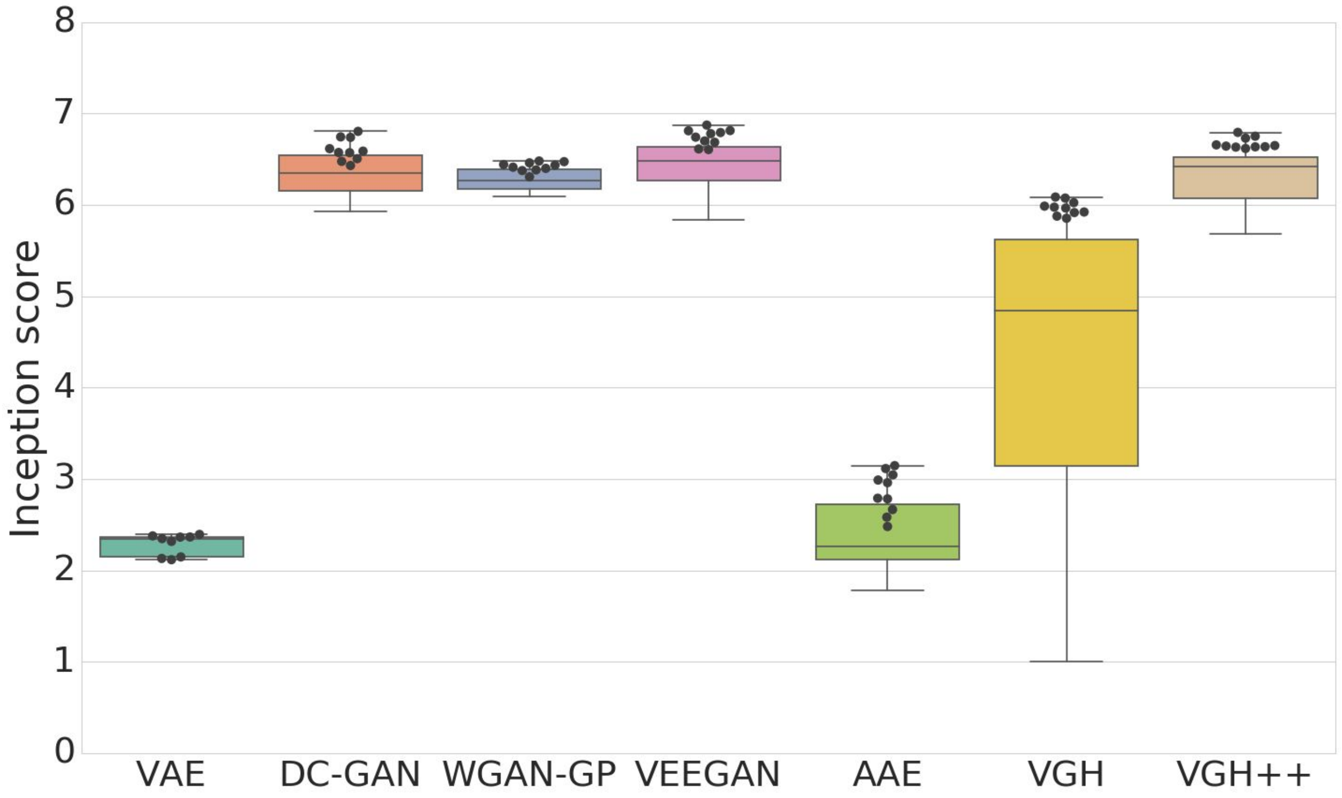

Improving GAN stability

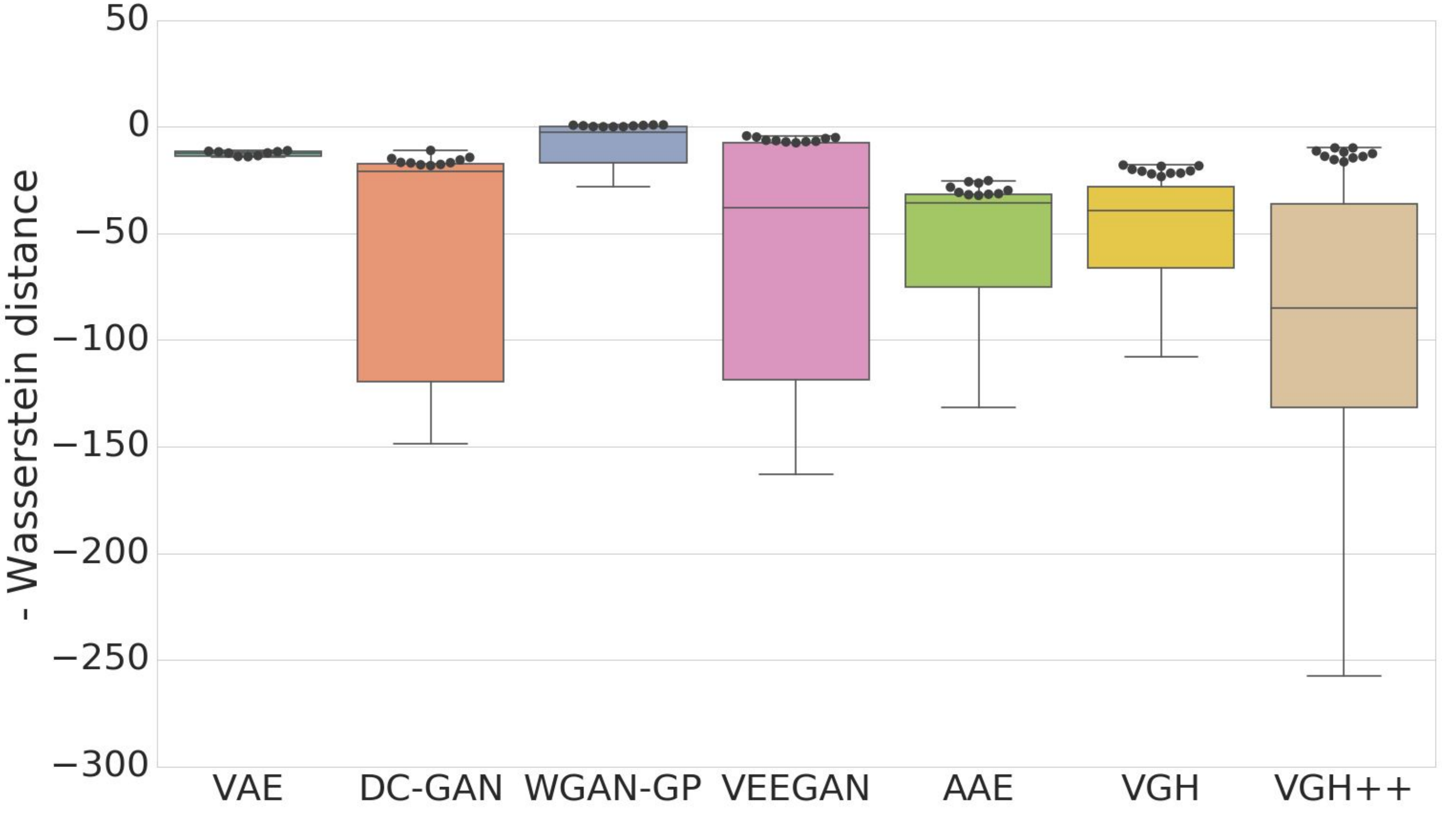

DeepMind

Jiwoong Im et all.: Quantitatively Evaluating GANs With Divergences Proposed for Training

Improving GAN diversity

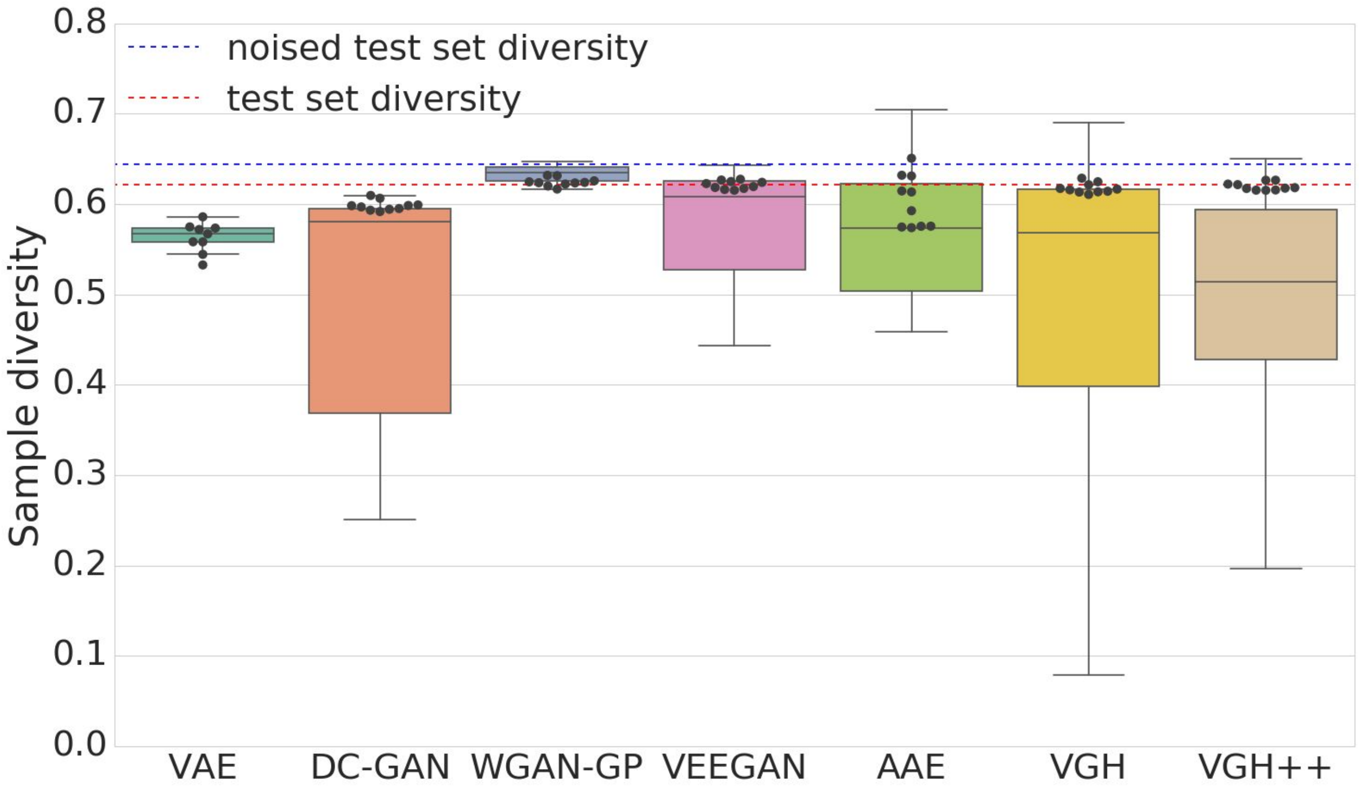

Improving VAE sample quality

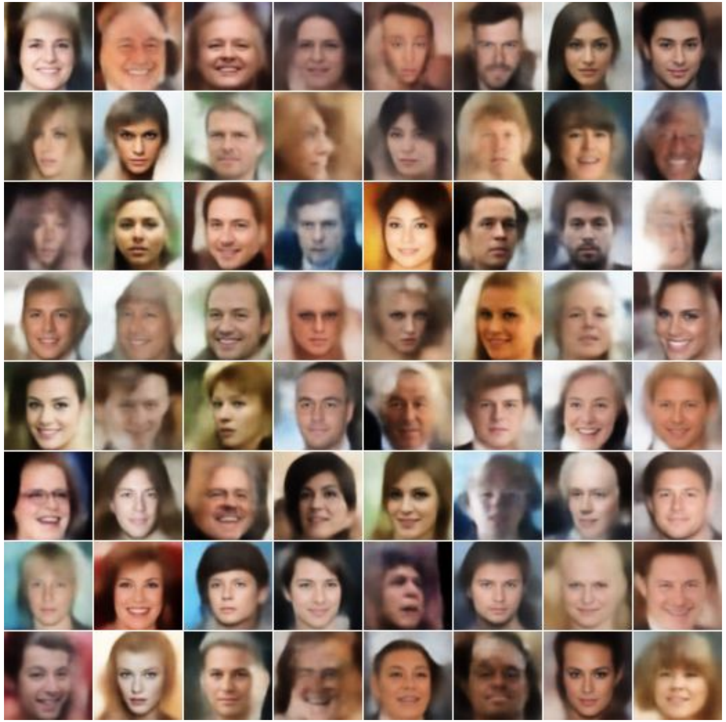

VAE

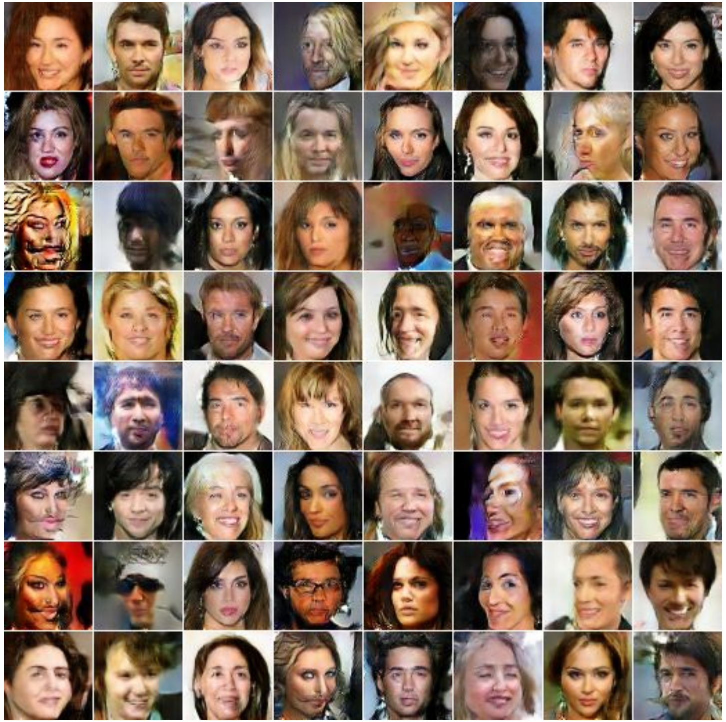

VEEGAN

Improving GAN sample quality

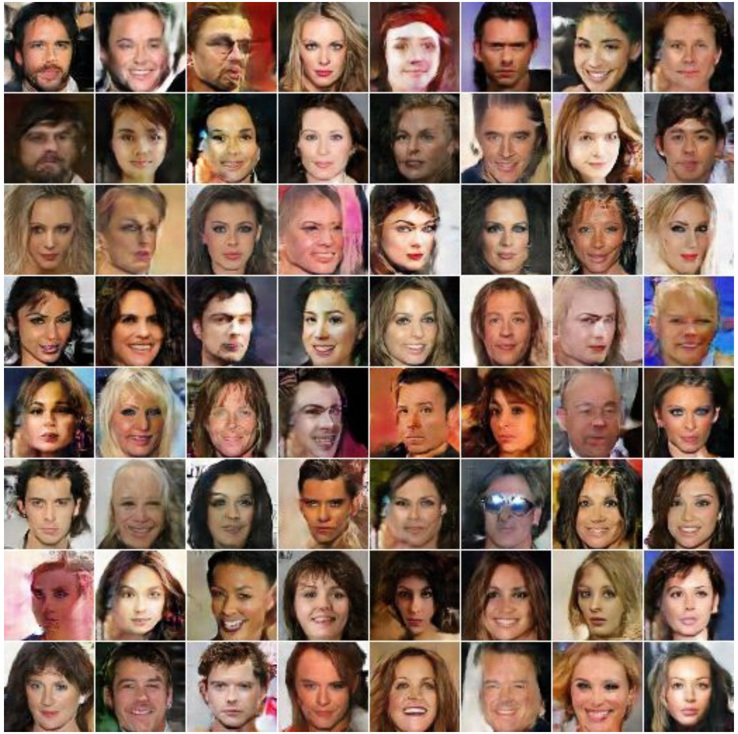

DCGAN

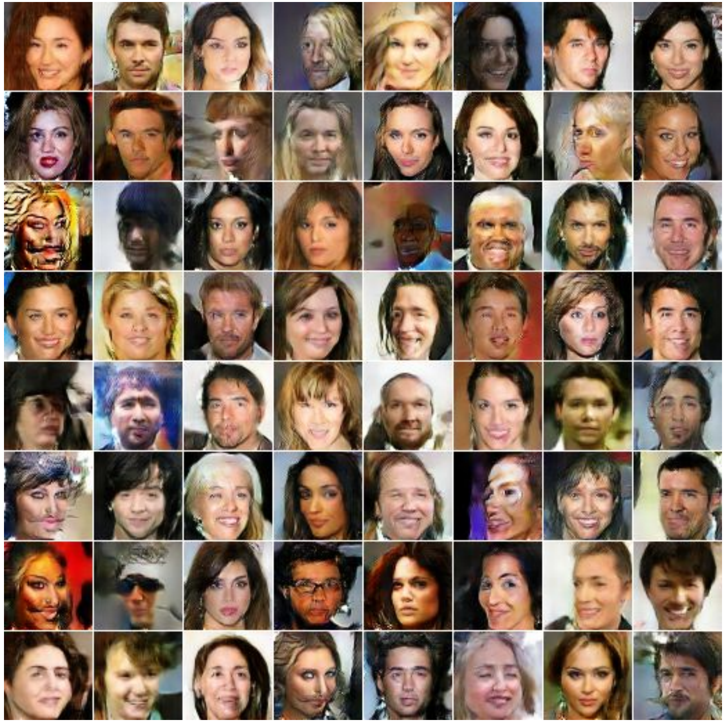

VEEGAN

At present, VAE-GAN hybrids

do not improve

distribution matching in latent and visible space.

Wait - how about CycleGAN?

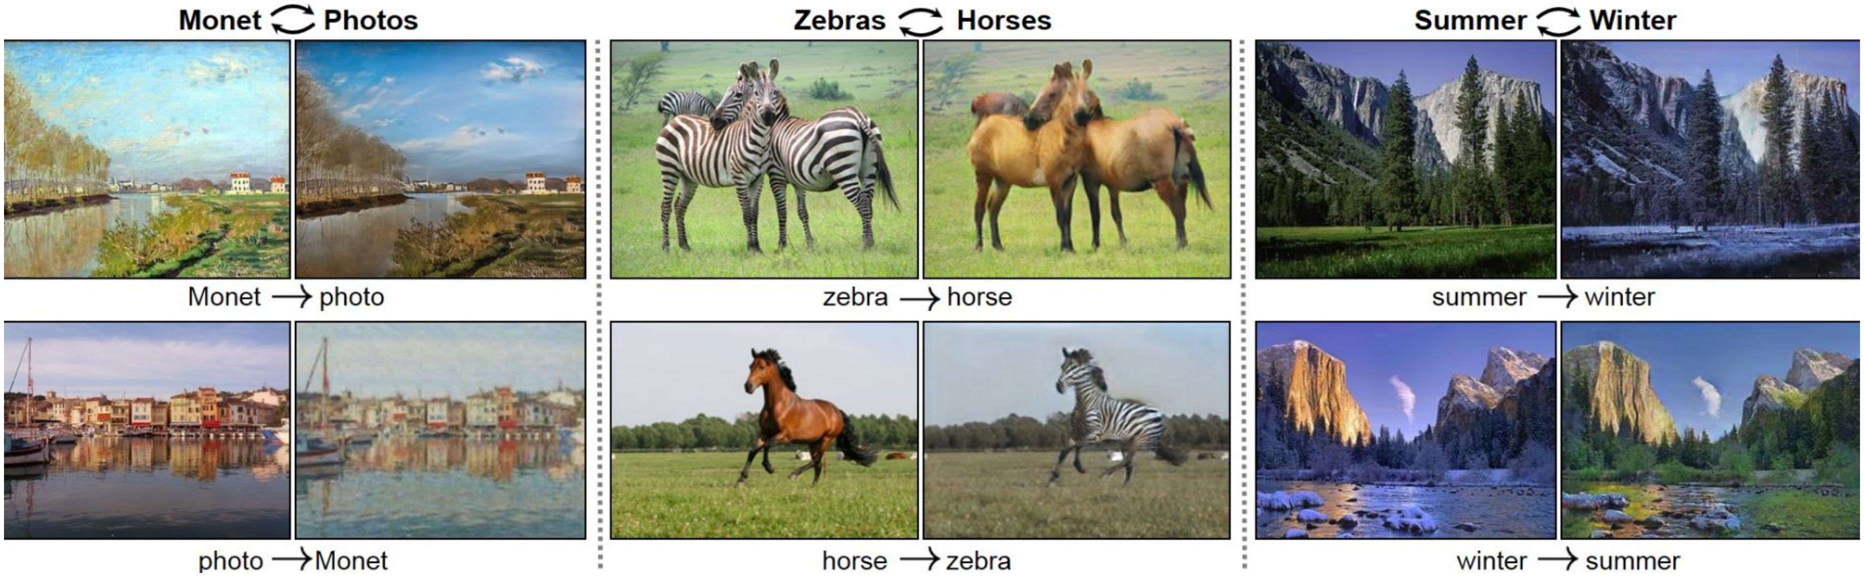

Zhu, Park, Isola, Efros: Unpaired Image-to-Image Translation using Cycle-Consistent Adversarial Networks

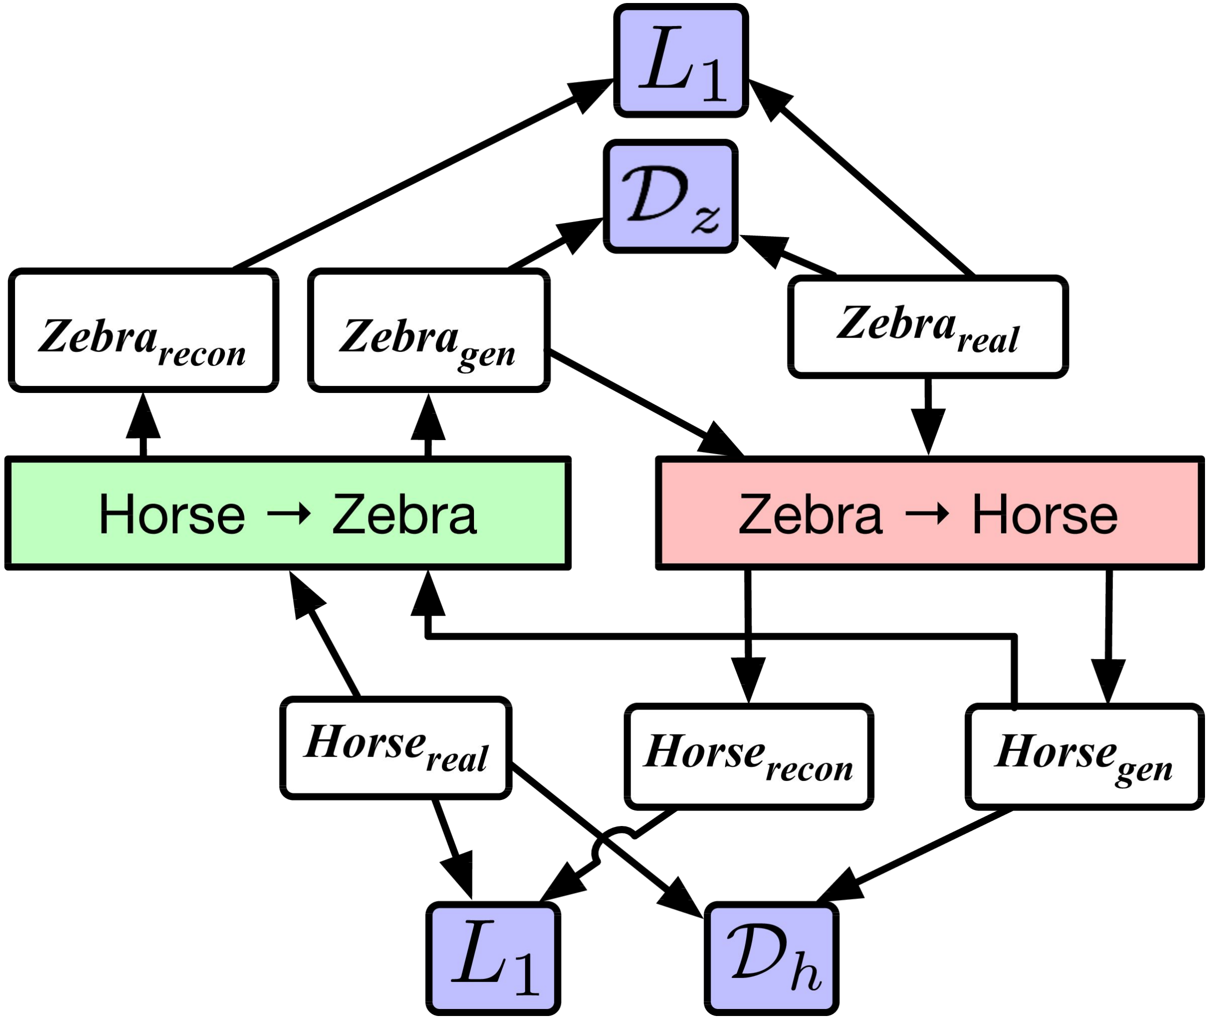
L_1
$\mathcal{D}_z$
$Zebra_{recon}$
$Zebra_{gen}$
$Zebra_{real}$
Horse → Zebra
Zebra → Horse
$Horse_{real}$
$Horse_{recon}$
$Horse_{gen}$
L_1
$\mathcal{D}_h$

But...

- Image to image translation versus image generation
  - No latent variables
  - More constrained
  - Easier to do architecture search

Currently, VAE -GANs do not deliver on their promise to stabilize GAN training or improve VAEs.

Currently, VAE -GANs do not deliver on their promise to stabilize GAN training or improve VAEs.

If you want good samples, use GANs.
If you care about representation learning, use VAEs.

THANK YOU

Credits

Shakir Mohamed, Balaji Lakshminarayanan

Additional Credits