20FALL Introduction to Deep Learning Recitation 10

Variational Autoencoders

By Akshat Gupta, Jiachen Lian 11/13/2020

Auto-Encoder

- Output is the input itself.

$$min\ E(||\hat{X} - X||^2)$$

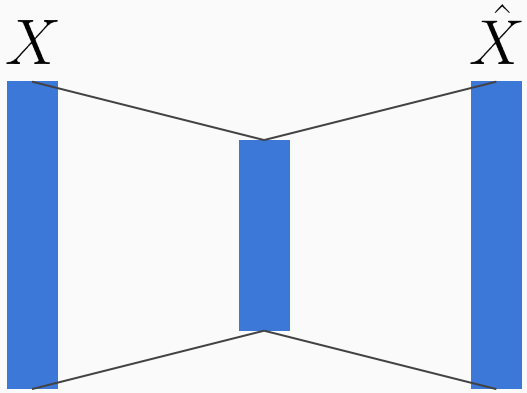

Auto-Encoder

- Output is the input itself.

  $min\ E(||\hat{X} - X||^2)$

- Compressed Latent Representation.

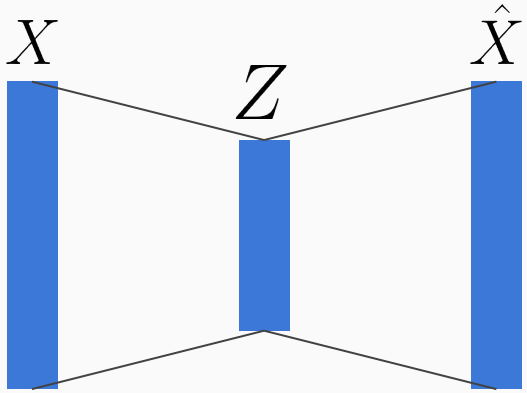

Auto-Encoder

- Output is the input itself.

$$min\ E(||\hat{X} - X||^2)$$

- Compressed Latent Representation.
- Linear single-layer AE performs PCA

$$\hat{X} = WW^T X$$

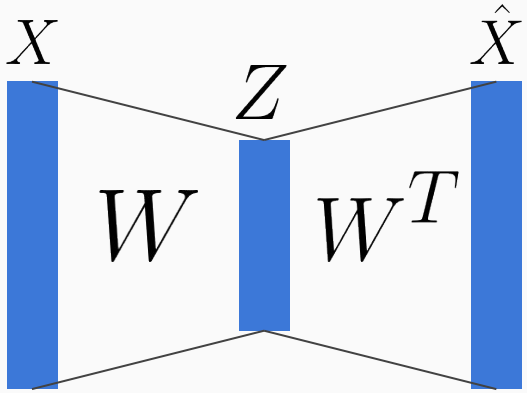

Auto-Encoder

- Output is the input itself.

$$min\ E(||\hat{X} - X||^2)$$

- Compressed Latent Representation.
- Linear single-layer AE performs PCA

$$\hat{X} = WW^T X$$

  - It is not PCA.

$$WW^T \neq I$$

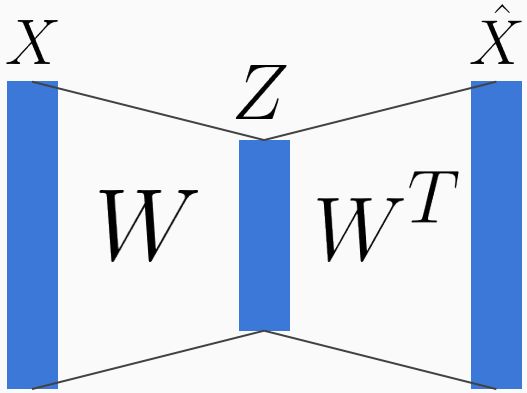

Auto-Encoder

- Output is the input itself.

$$min\ E(||\hat{X} - X||^2)$$

- Compressed Latent Representation.
- Linear single-layer AE performs PCA

$$\hat{X} = WW^T X$$

  - It is not PCA.

$$WW^T \neq I$$

- Deep non-linear AE generates powerful representation

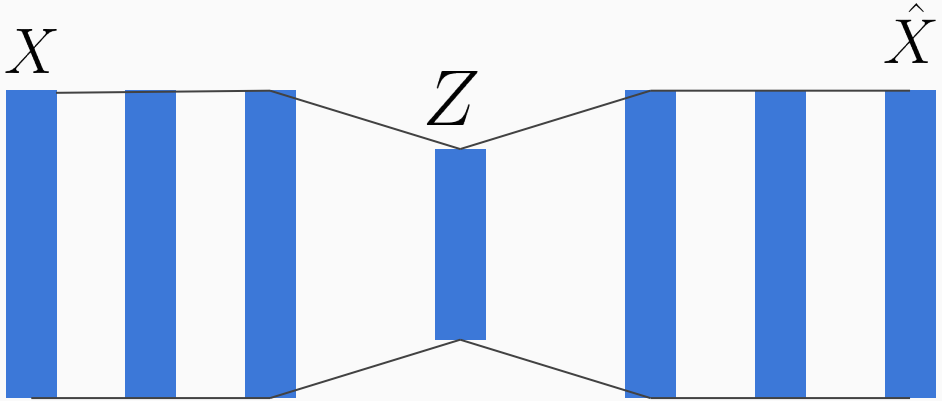

Auto-Encoder

- Output is the input itself.

$$min\ E(||\hat{X} - X||^2)$$

- Compressed Latent Representation.
- Linear single-layer AE performs PCA

$$\hat{X} = WW^T X$$

  - It is not PCA.

$$WW^T \neq I$$

- Deep non-linear AE generates powerful representation

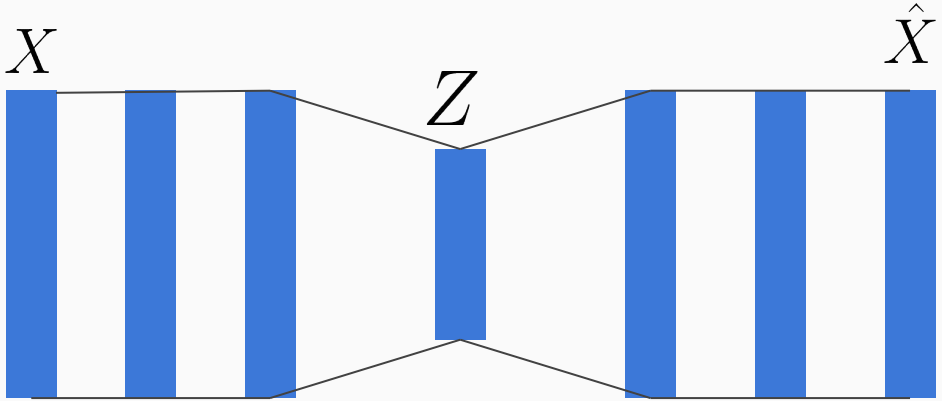

Auto-Encoder

- Why AE?

  Compressed Latent Representation.

  Promising Applications. E.g. Image denoising(Super Resolution), Neural Machine Translation

- Why not AE?

  Zero Reconstruction loss on a limited dataset

  It is difficult to generate a new datapoint

Why Variational Auto-Encoder?

Generate an image from a distribution

$$P(X) \rightarrow X$$

Why Variational Auto-Encoder?

Generate an image without the image itself

$$P(X) \rightarrow X$$

X_1

X_2

X_3

X_4

Why Variational Auto-Encoder?

Generate an image without the image itself

$$P(X) \rightarrow X$$

X_1

X_2

X_3

X_4

$$X_i = [X_{i1}, X_{i2}, ..., X_{ik}]^T$$

$$X \sim N(\mu, \Sigma)$$

$$\mu = \frac{X_1 + X_2 + X_3 + X_4}{4}$$

$$\Sigma = \Sigma_{i=1}^{4} \frac{(X_i - \mu)(X_i - \mu)^T}{4}$$

Why Variational Auto-Encoder?

Generate an image without the image itself

$$P(X) \rightarrow X$$

X_1

X_2

X_3

X_4

$$X_i = [X_{i1}, X_{i2}, ..., X_{ik}]^T$$

$$X \sim N(\mu, \Sigma)$$

$$\mu = \frac{X_1 + X_2 + X_3 + X_4}{4}$$

$$\Sigma = \Sigma_{i=1}^{4} \frac{(X_i - \mu)(X_i - \mu)^T}{4}$$

Why Variational Auto-Encoder?

Generate an image without the image itself

$$P(X) \rightarrow X$$

X_1

X_2

X_3

X_4

$$X_i = [X_{i1}, X_{i2}, ..., X_{ik}]^T$$

$$X \sim N(\mu, \Sigma)$$

$$\mu = \frac{X_1 + X_2 + X_3 + X_4}{4}$$

$$\Sigma = \Sigma_{i=1}^{4} \frac{(X_i - \mu)(X_i - \mu)^T}{4}$$

This is incorrect

Why Variational Auto-Encoder?

Generate an image without the image itself

$$P(X) \not\rightarrow X$$

$$P(X) = \int P(Z)P(X|Z)\,dZ$$

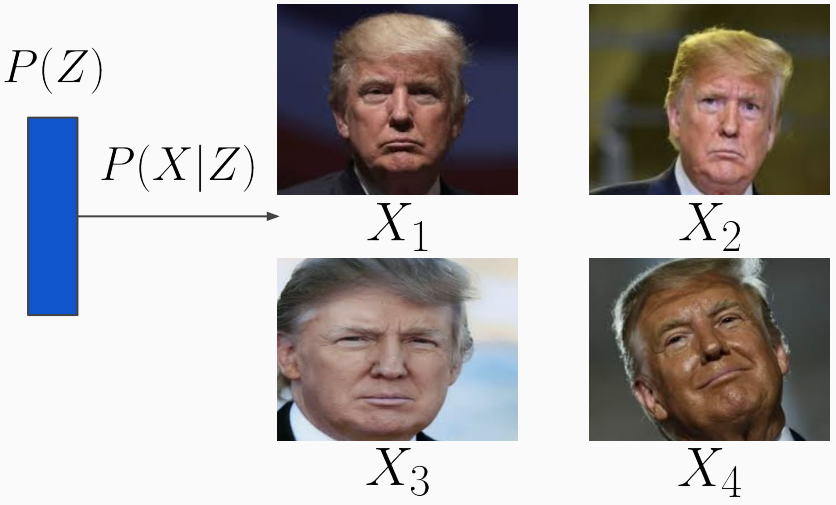

Why Variational Auto-Encoder?

Generate an image without the image itself

$$P(X) \rightarrow X$$

$$P(X) = \int P(Z)P(X|Z)\,dZ$$

Variational Inference

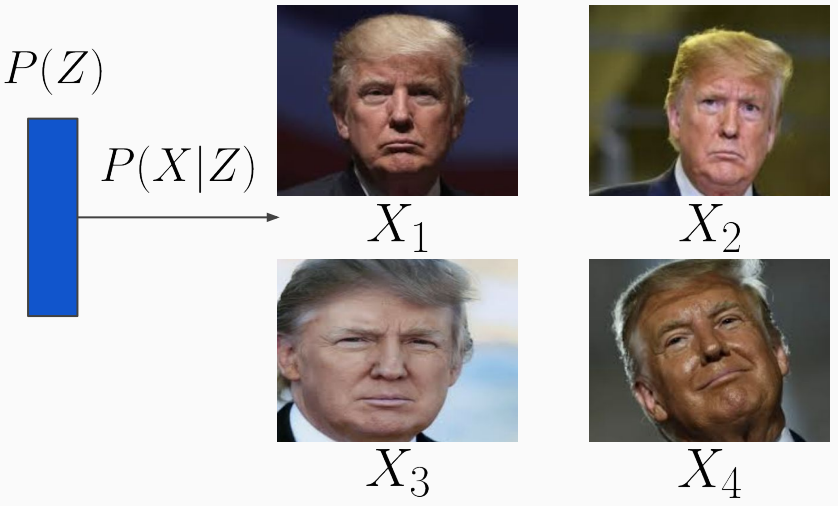

Why Variational Auto-Encoder?

Generate an image without the image itself

$$P(X) \rightarrow X$$

$$P(X) = \int P(Z)P(X|Z)\,dZ$$

Variational Inference

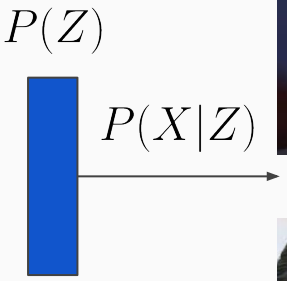

X_1

X_2

X_3

X_4

$$P(X,Z) = P(X)P(Z|X) = P(Z)P(X|Z)$$

Why Variational Auto-Encoder?

Generate an image without the image itself

$$P(X) \rightarrow X$$

$$P(X) = \int P(Z)P(X|Z)\,dZ$$

Variational Inference

$$P(X,Z) = P(X)P(Z|X) = P(Z)P(X|Z)$$

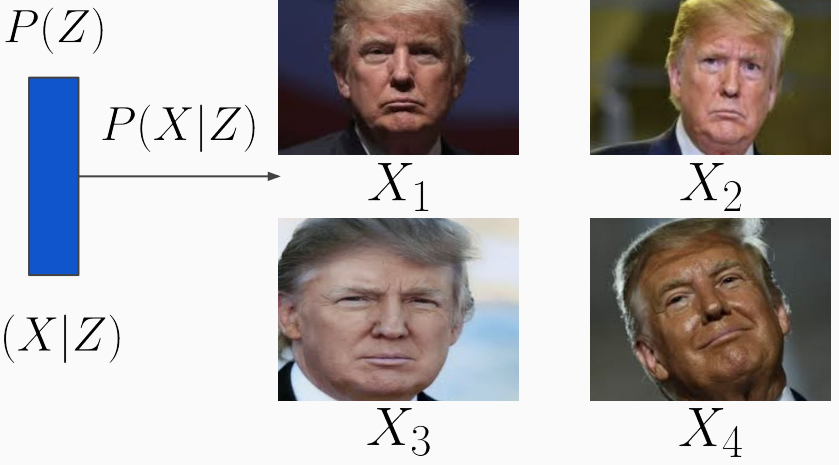

Why Variational Auto-Encoder?

Generate an image without the image itself

$$P(X) \rightarrow X$$

$$P(X) = \int P(Z)P(X|Z)\,dZ$$

Variational Inference

$$P(X,Z) = P(X)P(Z|X) = P(Z)P(X|Z)$$

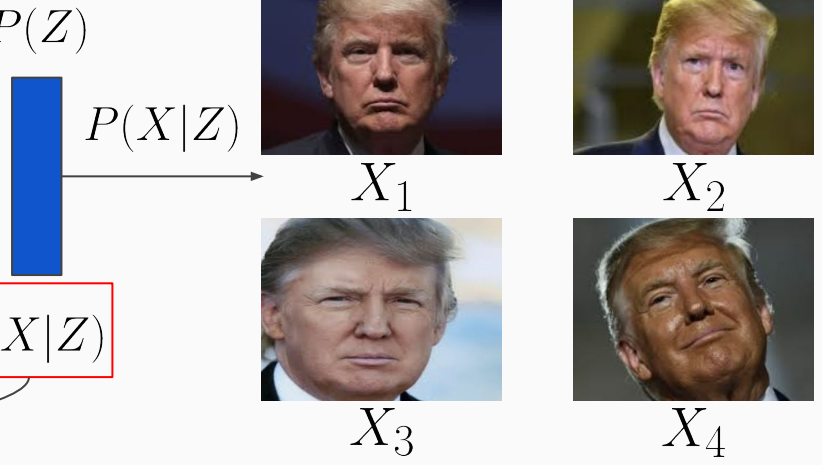

Why Variational Auto-Encoder?

Generate an image without the image itself

$$P(X) \rightarrow X$$

$$P(X) = \int P(Z)P(X|Z)\,dZ$$

Variational Inference

$$P(X,Z) = P(X)P(Z|X) = P(Z)P(X|Z)$$

Variational Auto-Encoder

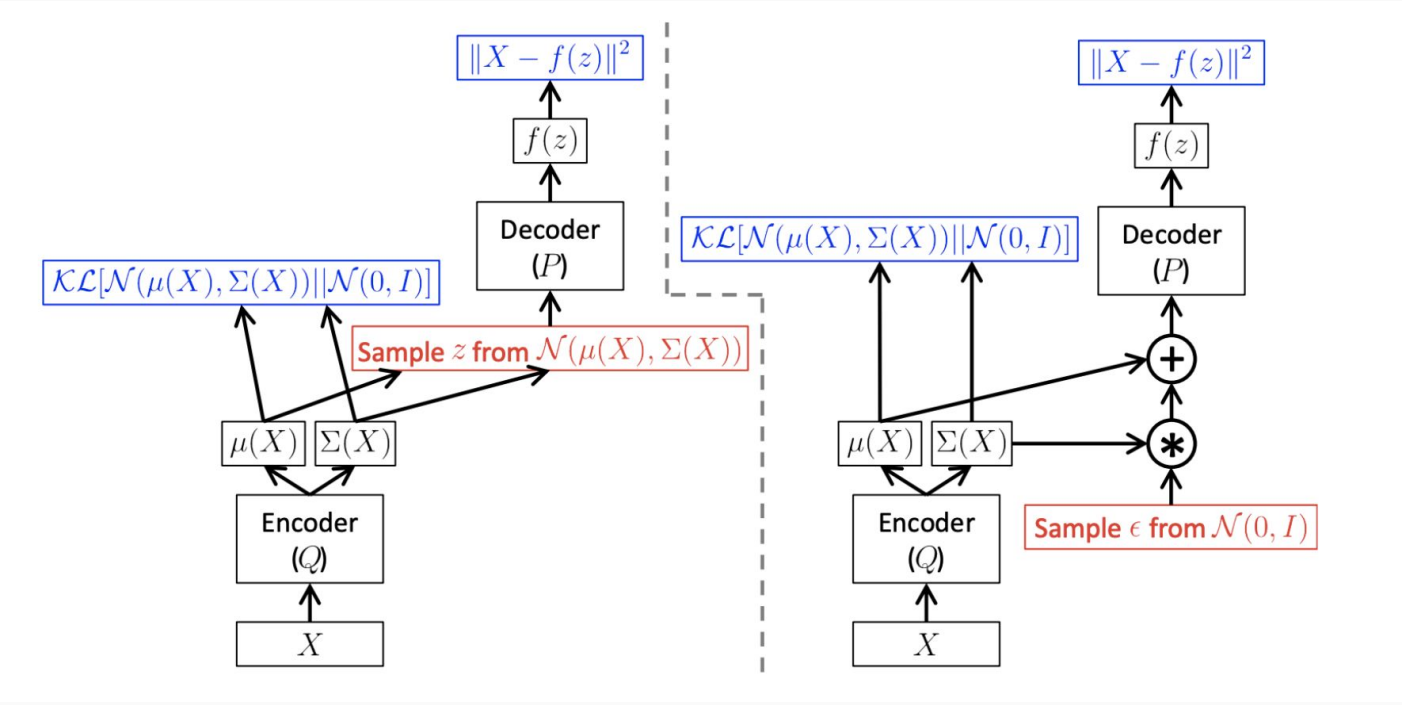

Variational Auto-Encoder

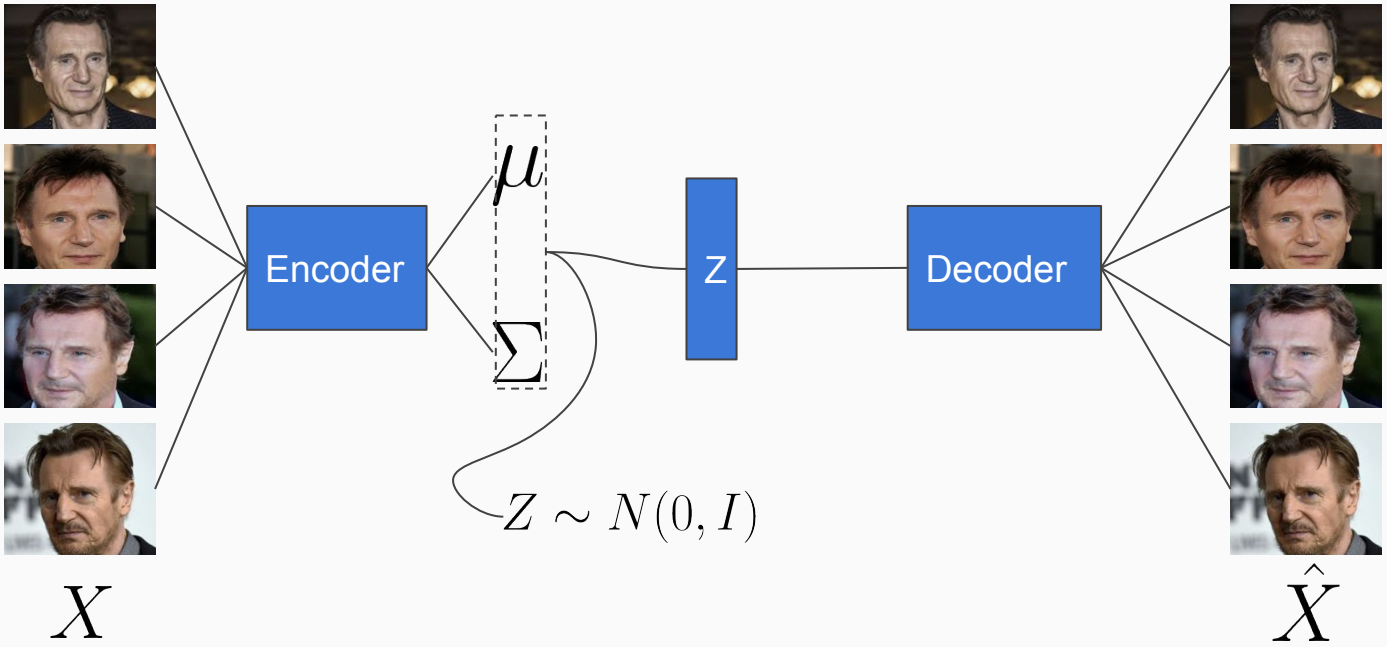

Variational Auto-Encoder

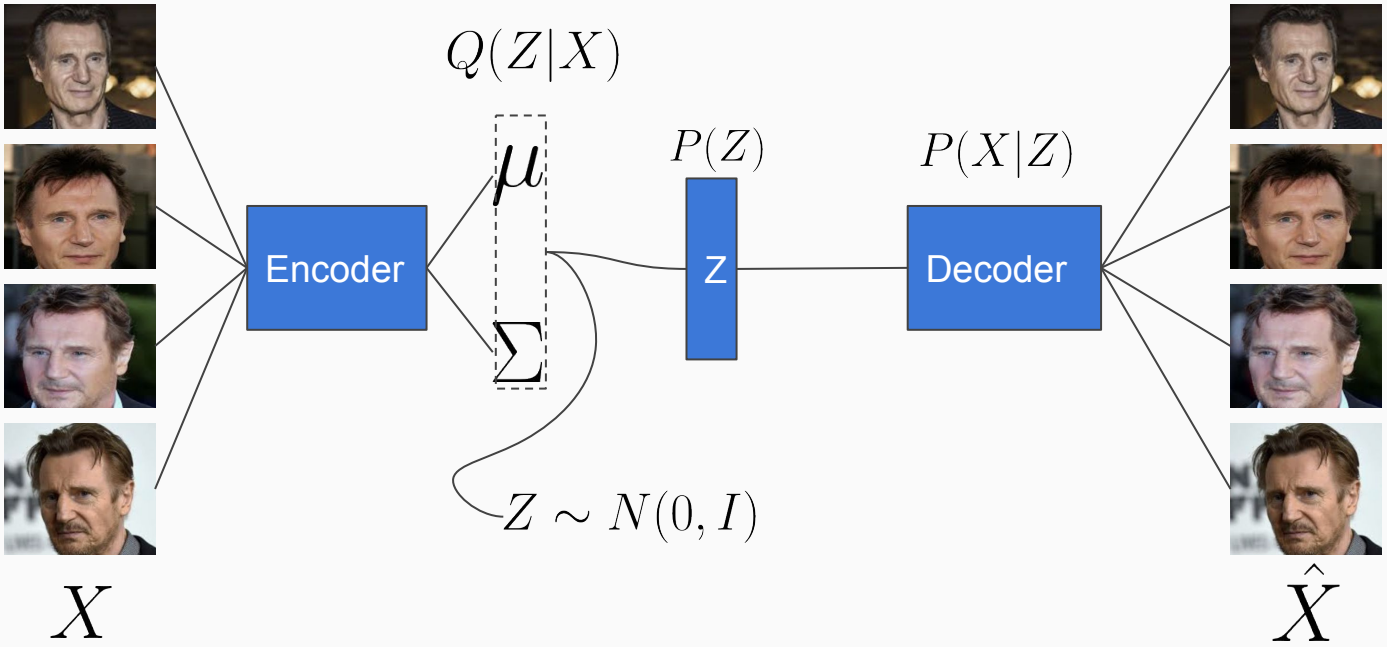

Variational Auto-Encoder

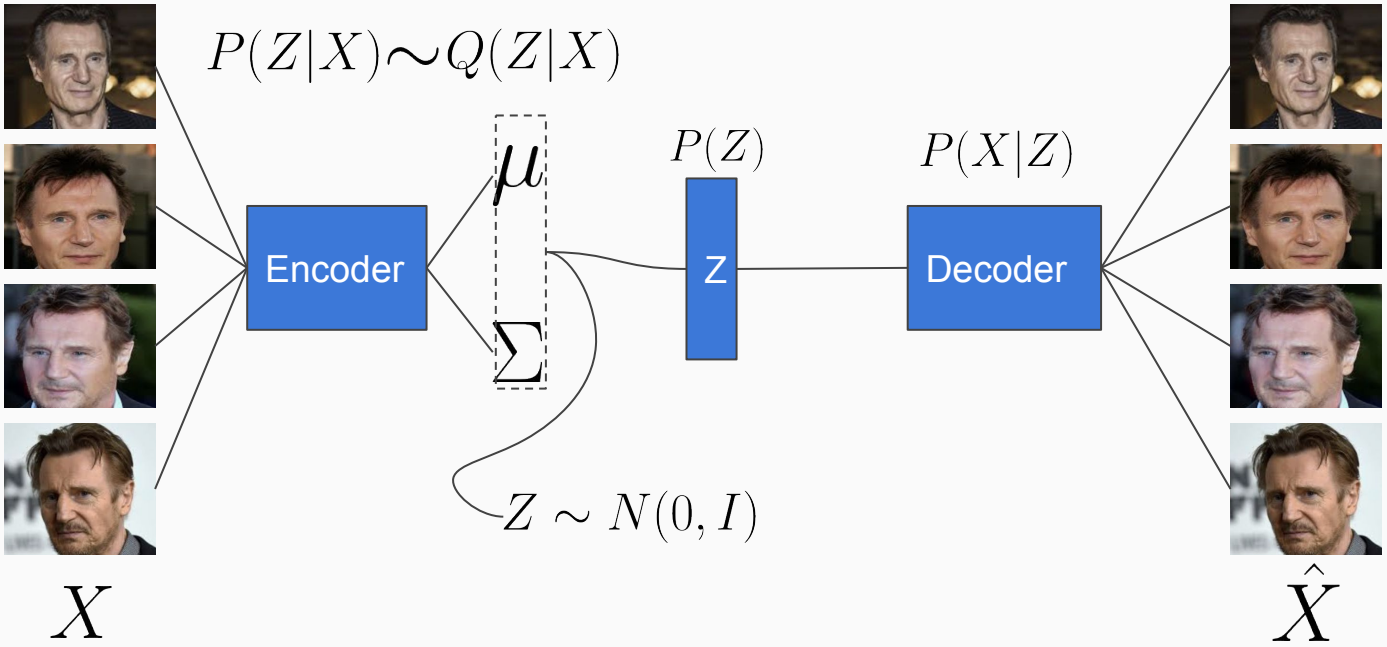

Variational Auto-Encoder

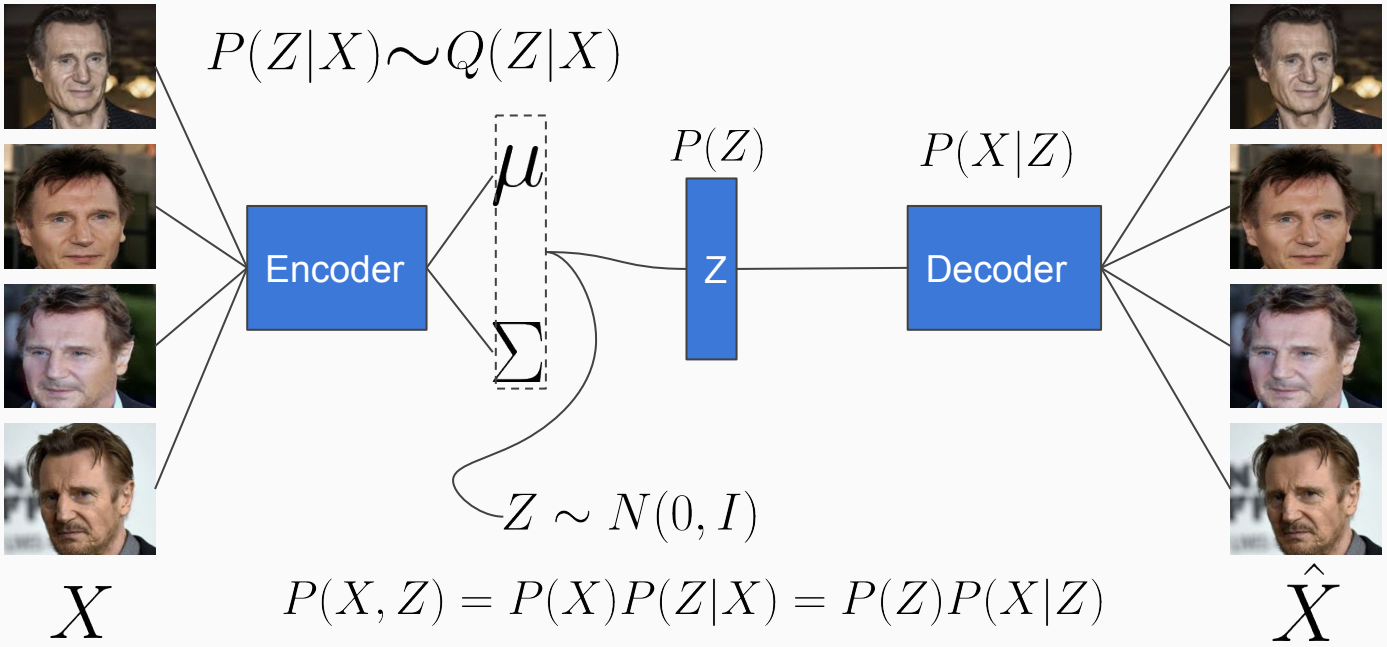

$$P(X, Z) = P(X)P(Z|X) = P(Z)P(X|Z)$$

Variational Auto-Encoder

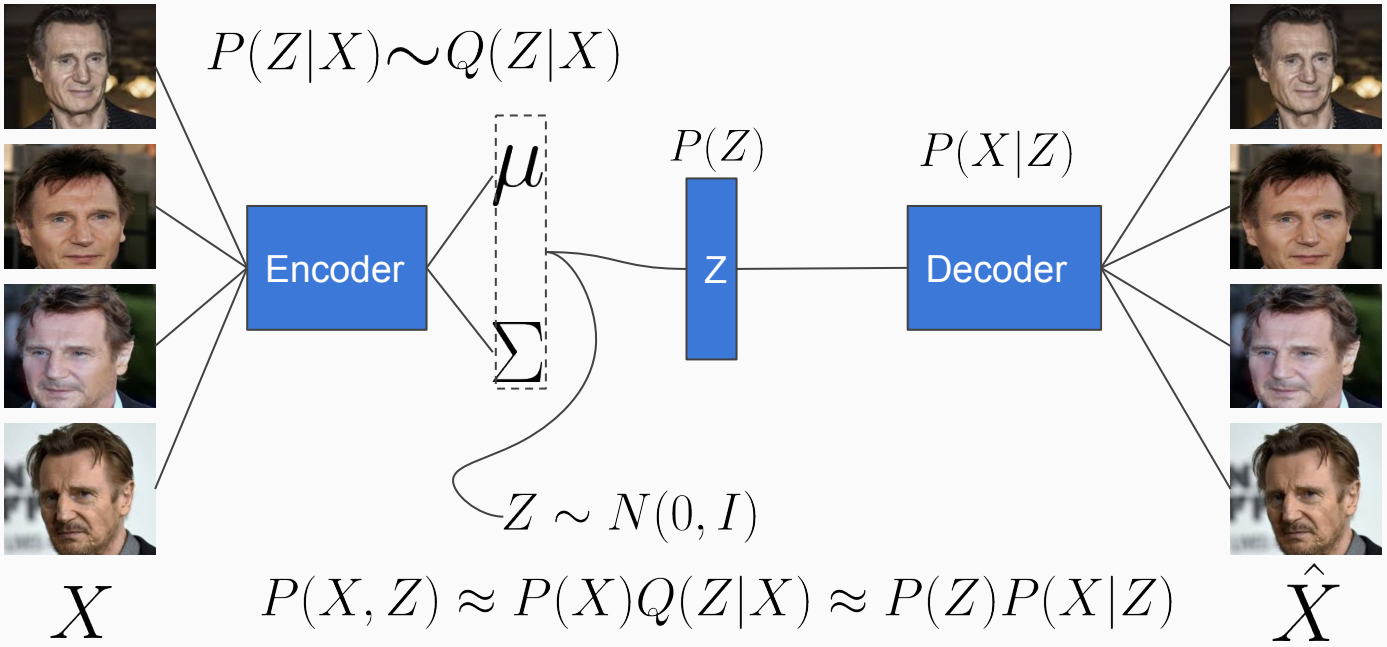

Variational Auto-Encoder

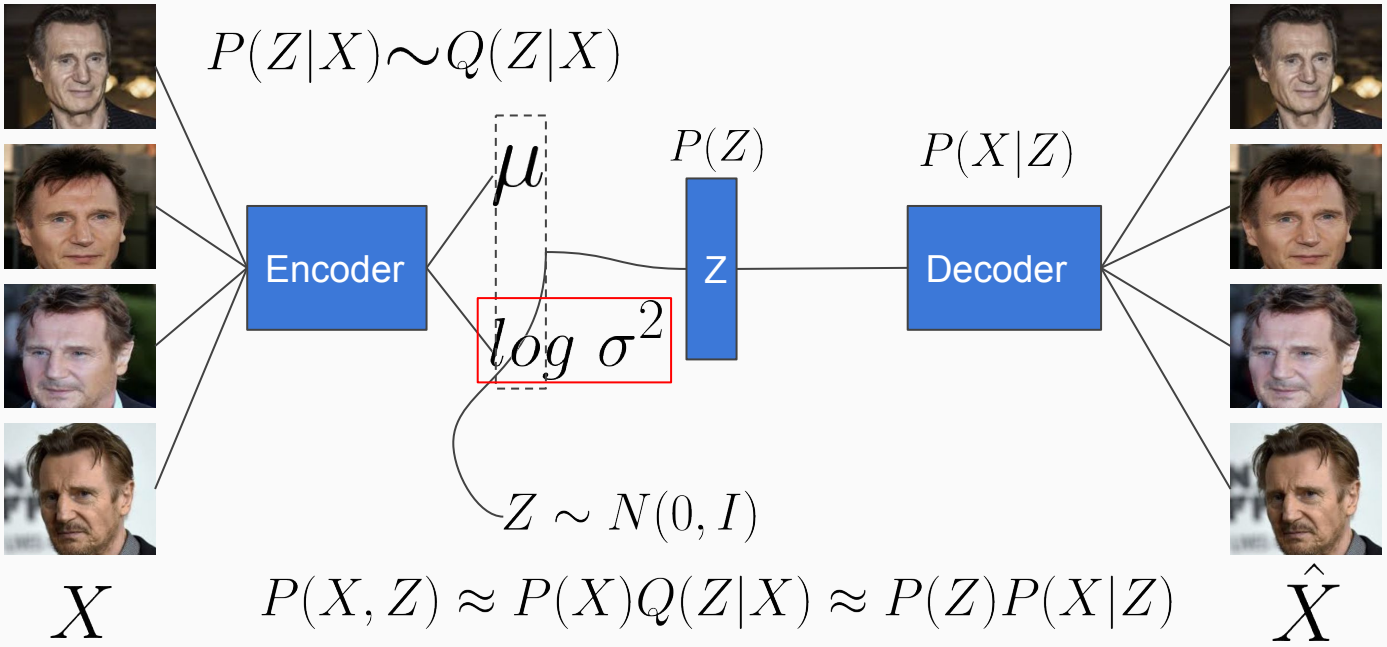

Variational Auto-Encoder

Statistical Distance between Distributions P(X) and Q(X)

Variational Auto-Encoder

Statistical Distance between Distributions P(X) and Q(X)

KL-Divergence

Variational Auto-Encoder

Statistical Distance between Distributions P(X) and Q(X)

KL-Divergence

Information: $I(p(X)) = -log\ p(X)$

Variational Auto-Encoder

Statistical Distance between Distributions P(X) and Q(X)

KL-Divergence

Information: $I(p(X)) = -log\ p(X)$

Entropy: $H = \Sigma_i p(X_i) I(p(X_i)) = \Sigma_i -p(X_i) log\ p(X_i)$

Variational Auto-Encoder

Statistical Distance between Distributions P(X) and Q(X)

KL-Divergence

Information: $I(p(X)) = -log\ p(X)$

Entropy: $H = \Sigma_i p(X_i) I(p(X_i)) = \Sigma_i -p(X_i) log\ p(X_i)$ $\quad - \int p\ log\ p\ dx$

Variational Auto-Encoder

Statistical Distance between Distributions P(X) and Q(X)

KL-Divergence

Information: $I(p(X)) = -log\ p(X)$

Entropy: $H = \Sigma_i p(X_i) I(p(X_i)) = \Sigma_i -p(X_i) log\ p(X_i)$ $\quad - \int p\ log\ p\ dx$

$KL(P||Q) = -\Sigma_i P(X_i) log \frac{Q(X_i)}{P(X_i)}$ $\quad - \int p\ log \frac{q}{p}\ dx$

Variational Auto-Encoder

Statistical Distance between Distributions P(X) and Q(X)

KL-Divergence

Information: $I(p(X)) = -log\ p(X)$

Entropy: $H = \Sigma_i p(X_i) I(p(X_i)) = \Sigma_i -p(X_i) log\ p(X_i)$ $\quad -\int plogp\ dp$

$KL(P||Q) = -\Sigma_i P(X_i) log \frac{Q(X_i)}{P(X_i)}$ $\quad -\int p\ log \frac{q}{p}\ dx$

$KL(Q||P) = -\Sigma_i Q(X_i) log \frac{P(X_i)}{Q(X_i)}$ $\quad -\int q\ log \frac{p}{q}\ dx$

Variational Auto-Encoder

Statistical Distance between Distributions P(X) and Q(X)

KL-Divergence

Information: $I(p(X)) = -log\ p(X)$

Entropy: $H = \Sigma_i p(X_i) I(p(X_i)) = \Sigma_i -p(X_i) log\ p(X_i)$ $\quad -\int plogp\ dp$

$$KL(P||Q) = -\Sigma_i P(X_i) log \frac{Q(X_i)}{P(X_i)}$$

$$KL(Q||P) = -\Sigma_i Q(X_i) log \frac{P(X_i)}{Q(X_i)}$$

$$KL(Q||P) \neq KL(P||Q)$$

You can pick either!

Variational Auto-Encoder

Statistical Distance between Distributions P(X) and Q(X)

KL-Divergence

$$lemma\ 1 : KL(P||Q) \geq 0$$

Variational Auto-Encoder

Statistical Distance between Distributions P(X) and Q(X)

KL-Divergence

$$lemma\ 1 : KL(P||Q) \geq 0$$

Proof: $$KL(P||Q) = -\Sigma P log\frac{Q}{P} \geq log(\Sigma P\frac{Q}{P}) = 0$$

Variational Auto-Encoder

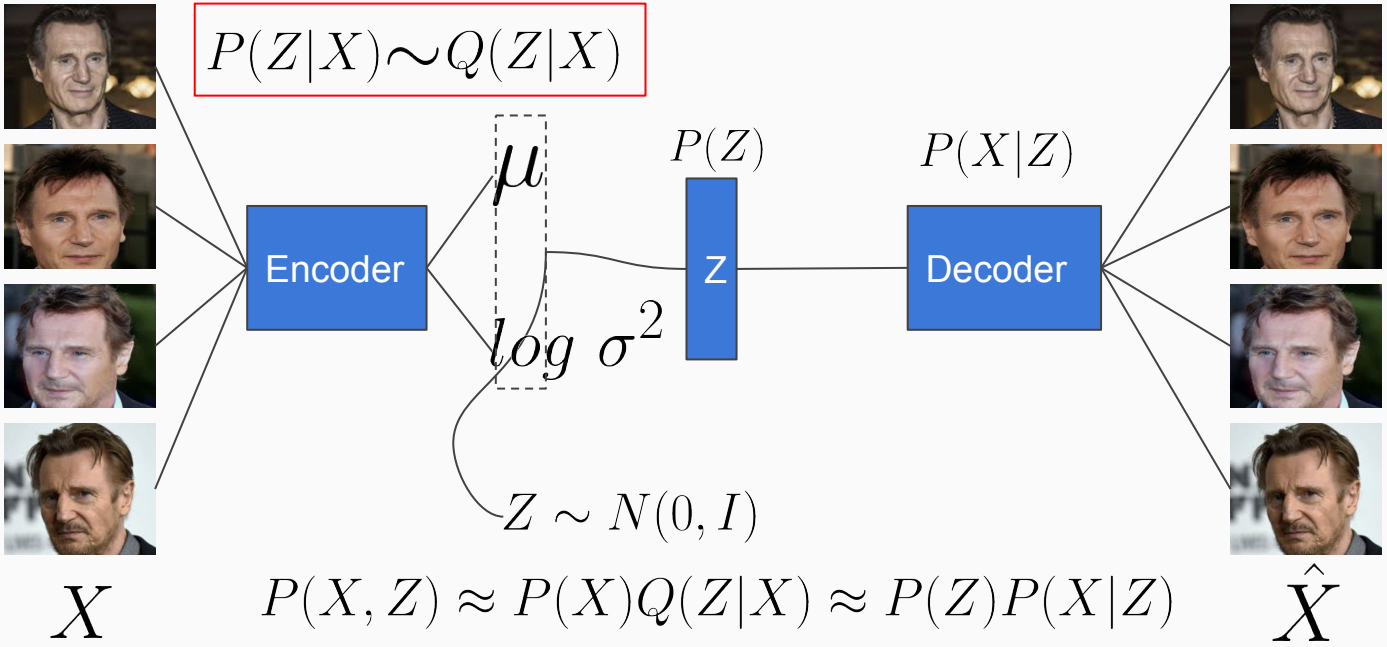

Variational Auto-Encoder

$$P(Z|X) \sim Q(Z|X)$$

Variational Auto-Encoder

$$P(Z|X) \sim Q(Z|X)$$

$$KL(Q(Z|X), P(Z|X)) = -\Sigma Q(Z|X) log \frac{P(Z|X)}{Q(Z|X))}$$

Variational Auto-Encoder

$$P(Z|X) \sim Q(Z|X)$$

$$KL(Q(Z|X), P(Z|X)) = -\Sigma Q(Z|X) log \frac{P(Z|X)}{Q(Z|X))}$$

$$= -\Sigma Q(Z|X) log \frac{P(X,Z)}{P(X)Q(Z|X)}$$

Variational Auto-Encoder

$$P(Z|X) \sim Q(Z|X)$$

$$KL(Q(Z|X), P(Z|X)) = -\Sigma Q(Z|X) log \frac{P(Z|X)}{Q(Z|X))}$$

$$= -\Sigma Q(Z|X) log \frac{P(X,Z)}{P(X)Q(Z|X)}$$

$$= -\Sigma Q(Z|X) log \frac{P(X,Z)}{Q(Z|X)} + log P(X)$$

Variational Auto-Encoder

$$P(Z|X) \sim Q(Z|X)$$

$$KL(Q(Z|X), P(Z|X)) = -\Sigma Q(Z|X) log \frac{P(Z|X)}{Q(Z|X))}$$

$$= -\Sigma Q(Z|X) log \frac{P(X,Z)}{P(X)Q(Z|X)}$$

$$= -\Sigma Q(Z|X) log \frac{P(X,Z)}{Q(Z|X)} + log P(X)$$

$$log P(X) = KL(Q(Z|X), P(Z|X)) + \Sigma Q(Z|X) log \frac{P(X,Z)}{Q(Z|X)}$$

Variational Auto-Encoder

$$logP(X) = KL(Q(Z|X), P(Z|X)) + \Sigma Q(Z|X) log \frac{P(X,Z)}{Q(Z|X)}$$

Variational Auto-Encoder

$$logP(X) = KL(Q(Z|X), P(Z|X)) + \Sigma Q(Z|X) log \frac{P(X,Z)}{Q(Z|X)}$$

Constant

Unknown

Evidence Lower Bound(ELBO)

Variational Auto-Encoder

$$logP(X) = KL(Q(Z|X), P(Z|X)) + \Sigma Q(Z|X) log \frac{P(X,Z)}{Q(Z|X)}$$

Constant | Unknown | Evidence Lower Bound(ELBO)

$$Max\ L = \Sigma Q(Z|X) log \frac{P(X,Z)}{Q(Z|X)}$$

Variational Auto-Encoder

$$P(X,Z) = P(X)P(Z|X) = P(Z)P(X|Z)$$

$$P(Z|X) \sim Q(Z|X)$$

$$logP(X) = KL(Q(Z|X), P(Z|X)) + \Sigma Q(Z|X) log \frac{P(X,Z)}{Q(Z|X)}$$

$$Max\ L = \Sigma Q(Z|X) log \frac{P(X,Z)}{Q(Z|X)}$$

Variational Auto-Encoder

$$Max\ L = \Sigma Q(Z|X) log \frac{P(X,Z)}{Q(Z|X)}$$

Variational Auto-Encoder

$$Max\ L = \Sigma Q(Z|X) log \frac{P(X,Z)}{Q(Z|X)}$$

$$= \Sigma Q(Z|X) log \frac{P(X|Z)P(Z)}{Q(Z|X)}$$

Variational Auto-Encoder

$$Max\ L = \Sigma Q(Z|X) log \frac{P(X,Z)}{Q(Z|X)}$$

$$= \Sigma Q(Z|X) log \frac{P(X|Z)P(Z)}{Q(Z|X)}$$

$$= \Sigma Q(Z|X) (logP(X|Z) + log \frac{P(Z)}{Q(Z|X)})$$

Variational Auto-Encoder

$$
\begin{aligned}
Max\ L &= \Sigma Q(Z|X) log \frac{P(X,Z)}{Q(Z|X)} \\
&= \Sigma Q(Z|X) log \frac{P(X|Z)P(Z)}{Q(Z|X)} \\
&= \Sigma Q(Z|X) (logP(X|Z) + log \frac{P(Z)}{Q(Z|X)}) \\
&= \Sigma Q(Z|X) logP(X|Z) - KL(Q(Z|X)||P(Z))
\end{aligned}
$$

Variational Auto-Encoder

$$
\begin{aligned}
Max\ L &= \Sigma Q(Z|X) log \frac{P(X,Z)}{Q(Z|X)} \\
&= \Sigma Q(Z|X) log \frac{P(X|Z)P(Z)}{Q(Z|X)} \\
&= \Sigma Q(Z|X) (logP(X|Z) + log \frac{P(Z)}{Q(Z|X)}) \\
&= \Sigma Q(Z|X) logP(X|Z) - KL(Q(Z|X)||P(Z)) \\
&= E_{Q(Z|X)} logP(X|Z) - KL(Q(Z|X)||P(Z))
\end{aligned}
$$

Variational Auto-Encoder

$$L = E_{Q(Z|X)} logP(X|Z) - KL(Q(Z|X)||P(Z))$$

Reconstruction Loss

Tractable

$$Loss = -E_{Q(Z|X)} logP(X|Z) + KL(Q(Z|X)||P(Z))$$

Variational Auto-Encoder

$$P(X,Z) = P(X)P(Z|X) = P(Z)P(X|Z)$$

$$P(Z|X) \sim Q(Z|X)$$

$$logP(X) = KL(Q(Z|X), P(Z|X)) + \Sigma Q(Z|X) log \frac{P(X,Z)}{Q(Z|X)}$$

$$Max\ L = \Sigma Q(Z|X) log \frac{P(X,Z)}{Q(Z|X)}$$

$$L = E_{Q(Z|X)} logP(X|Z) - KL(Q(Z|X)||P(Z))$$

$$Loss = -E_{Q(Z|X)} logP(X|Z) + KL(Q(Z|X)||P(Z))$$

Variational Auto-Encoder

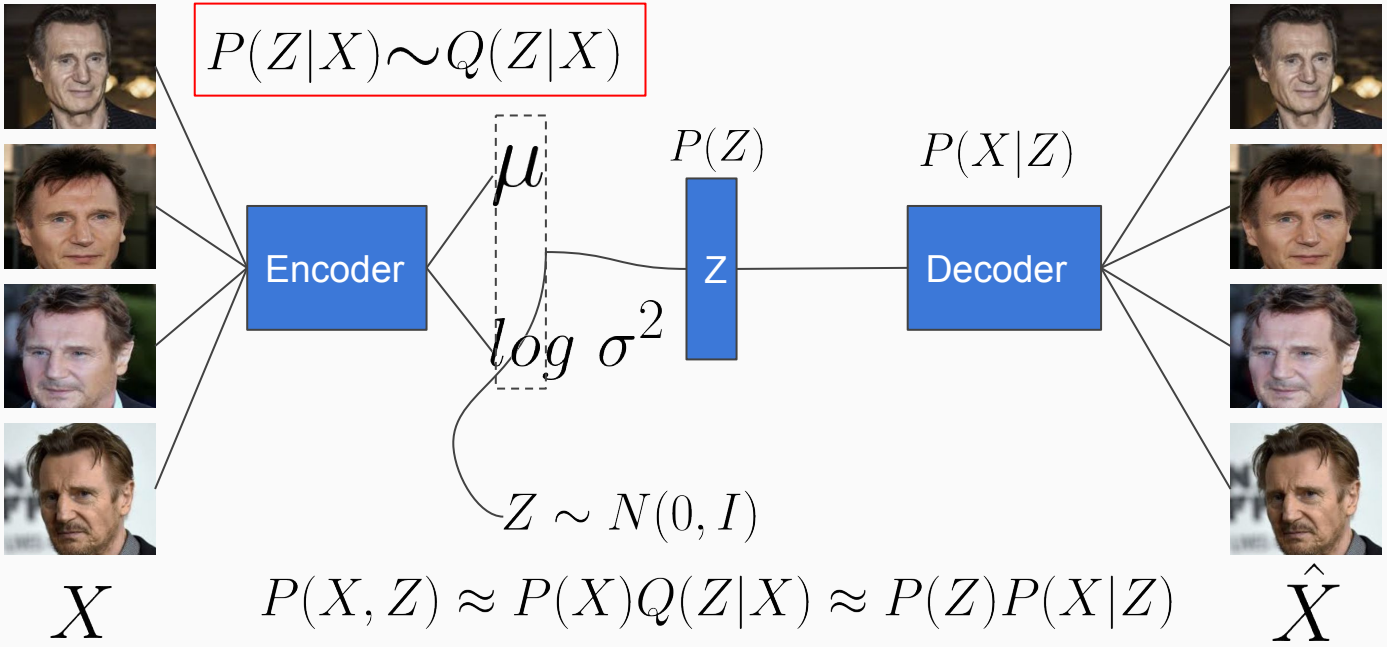

Variational Auto-Encoder

- Discussion 1

$$Q(Z|X) \sim P(Z) \sim P(Z|X)$$

Variational Auto-Encoder

- Discussion 1

$$Q(Z|X) \sim P(Z) \sim P(Z|X)$$

$$KL(N(\mu, \sigma^2)||N(0,1)) = \frac{1}{2}(-log\sigma^2 + \mu^2 + \sigma^2 - 1)$$

$$KL(p||q) = -\int p\, log\frac{q}{p}\, dx$$

Variational Auto-Encoder

- Discussion 1

$$Q(Z|X) \sim P(Z) \sim P(Z|X)$$

$$KL(N(\mu, \sigma^2)||N(0,1)) = \frac{1}{2}(-log\sigma^2 + \mu^2 + \sigma^2 - 1)$$

$$KL(p||q) = -\int p\, log\frac{q}{p}\, dx$$

$$Loss = E(||X - \hat{X}||^2) + \frac{1}{2}(-log\sigma^2 + \mu^2 + \sigma^2 - 1)$$

Variational Auto-Encoder

- Discussion 2

$$Q(Z|X) \sim P(Z) \sim P(Z|X)$$

$$Loss = -E_{Q(Z|X)} logP(X|Z) + KL(Q(Z|X)||P(Z))$$

Variational Auto-Encoder

- Discussion 2

$$Q(Z|X) \sim P(Z) \sim P(Z|X)$$

$$Loss = -E_{Q(Z|X)} logP(X|Z) + KL(Q(Z|X)||P(Z))$$

$$P(Z) = P(Z|X) \rightarrow P(Z,X) = P(Z)P(X)$$

Variational Auto-Encoder

- Discussion 2

$$Q(Z|X) \sim P(Z) \sim P(Z|X)$$

$$Loss = -E_{Q(Z|X)} logP(X|Z) + KL(Q(Z|X)||P(Z))$$

$$P(Z) = P(Z|X) \rightarrow P(Z,X) = P(Z)P(X)$$

Undesired!

Variational Auto-Encoder

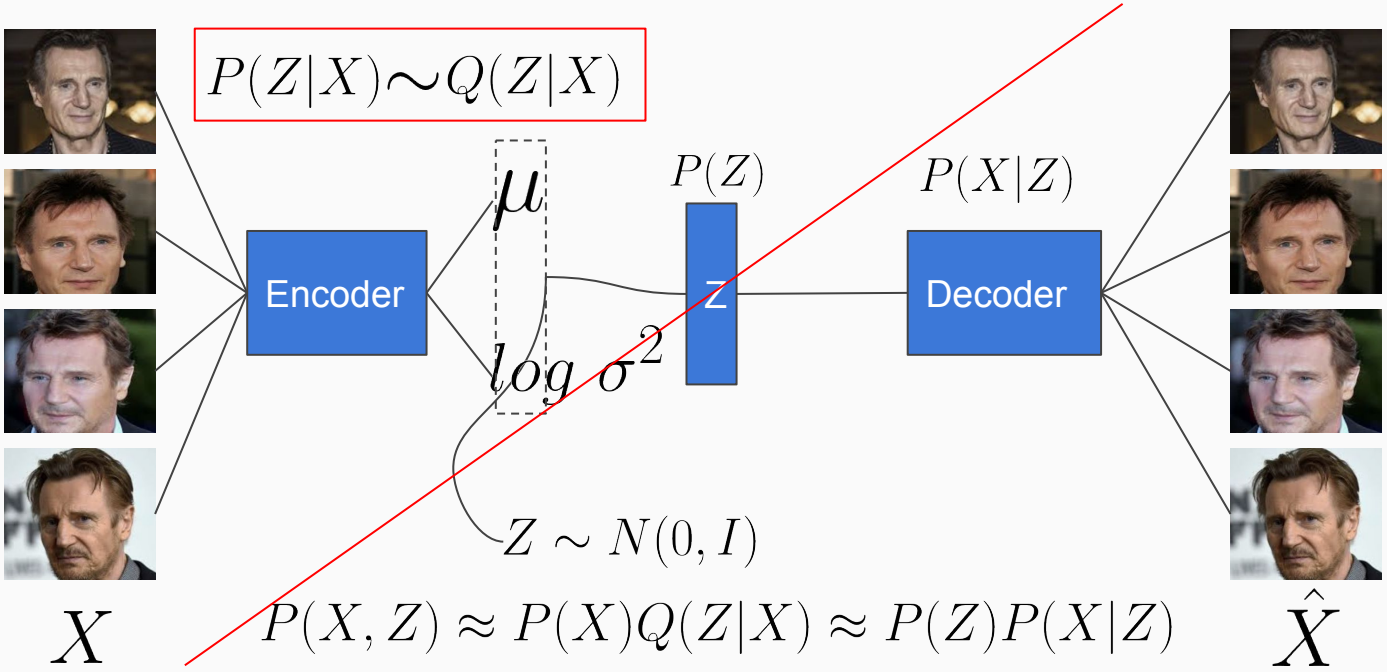

Variational Auto-Encoder

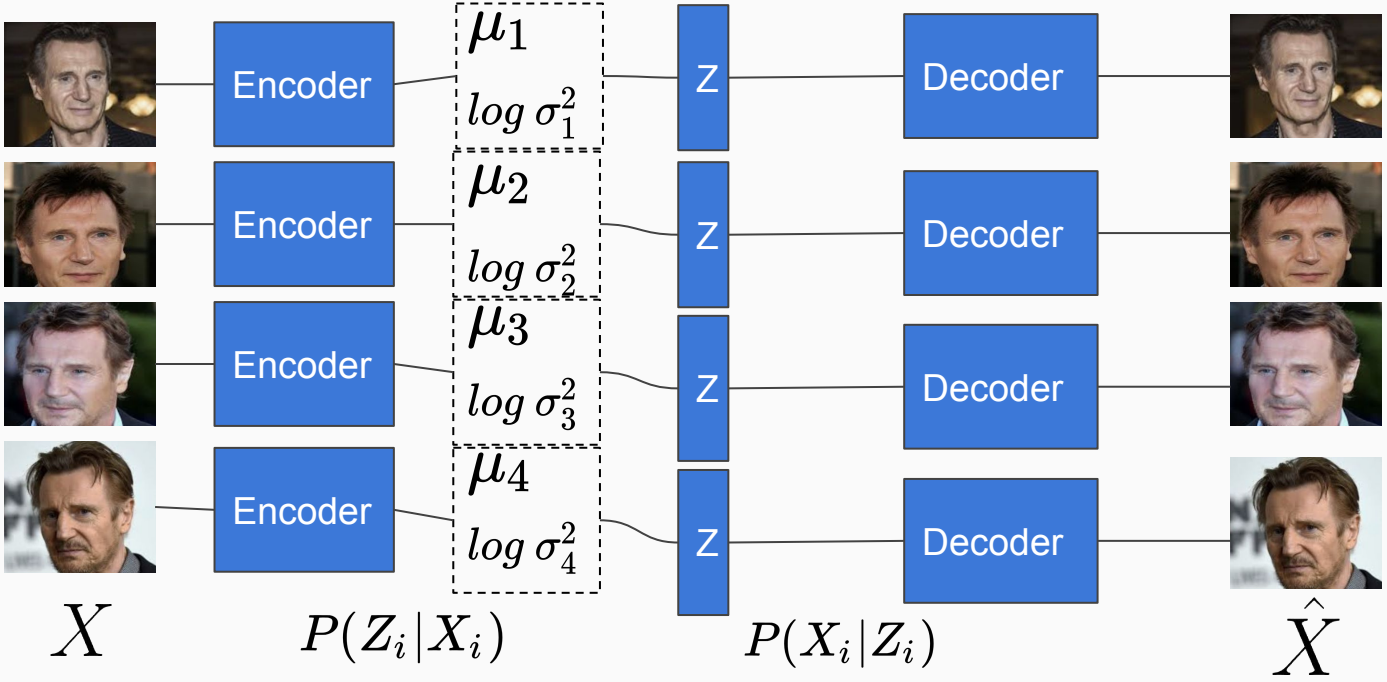

Variational Auto-Encoder

- Discussion 3

Reparameterization

$$Z = \mu + \epsilon\sigma \sim N(\mu, \sigma^2)$$

$$\epsilon \sim N(0, 1)$$

Variational Auto-Encoder