Variational Auto-Encoders (VAEs)

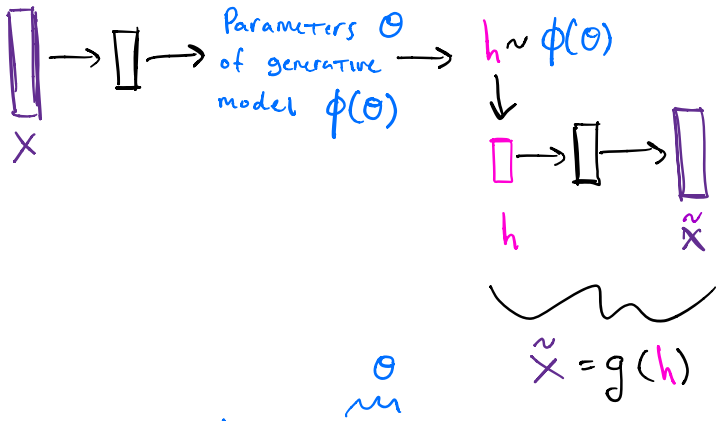

In practice $\phi = \underset{\theta}{\sim} N(\mu, \sigma)$

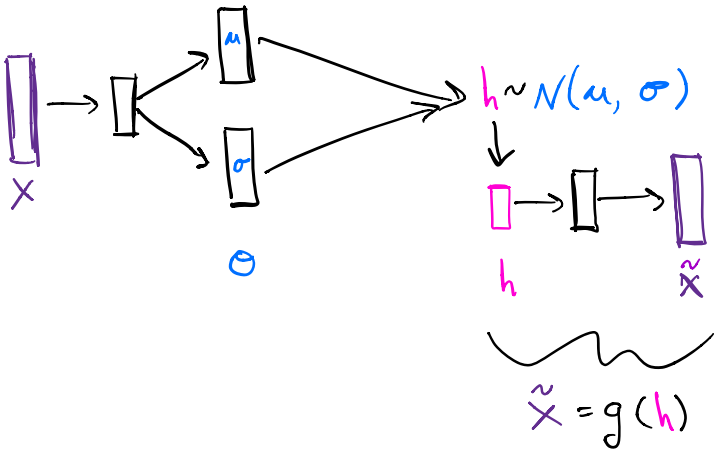

Note that we predict a vector <1 x hidden> of μ and σ, which means we assume dimensions in h are uncorrelated

$[5, -3, 0]$

μ

$$\begin{bmatrix} 1 & 0 & 0 \\ 0 & .5 & 0 \\ 0 & 0 & 2 \end{bmatrix}$$

σ

So a draw from this (e.g. via `np.random.normal`) might be

$h \approx [4.6, -2.7, 1.3]$

NOTE that for Normal distributions, we can equivalently sample $h = \mu + \sigma N(0, I)$ "reparameterization trick"

This is important because we cannot backprop through a stochastic tensor.

From the auto-diff view, $N(0, I)$ is just a constant during any particular forward pass

$$e \sim N(0, I)$$

$$h = \mu + \sigma e$$

Torch cares about these

Let's see some code.