

Neural network programming: variational autoencoders

Erik Spence

SciNet HPC Consortium

11 May 2023

You can get the slides and code for today's class at the SciNet Education web page.

`https://scinet.courses/1271`

Click on the link for the class, and look under "Lectures", click on "Variational autoencoders".

This class will cover the following topics:

- Generative networks.
- Regular autoencoders.
- Variational autoencoders.
- KL divergence loss function.
- Example.

Please ask questions if something isn't clear.

Let's examine a distinction we haven't yet made: discriminative versus generative networks.

- A discriminative network is trained to detect whether some input data is a member of a given class. Examples include the standard networks we've come to know and love, such as fully-connected and CNNs.
- In probabilistic terms, given the input data \mathbf{x} , and a desired label \mathbf{y} , the discriminative network calculates the conditional probability $p(\mathbf{y}|\mathbf{x})$.
- In contrast, a generative network is trained to calculate the probability distribution of the data itself, $p(\mathbf{x})$.
- Once we have $p(\mathbf{x})$, we can sample from this distribution to create new data points.

The ability to generate fake, authentic looking data has a number of applications.

There are several types of generative networks you may run into.


- PixelCNN: an auto-regressive model, the conditional distribution of each pixel is modelled given the left and above pixels.
- Variational Autoencoders: one network (the encoder) casts the input data into a lower-dimensional representation; a second network (the decoder) reconstructs the input from the low-D representation.
- Generative Adversarial Networks: two networks are trained simultaneously, one to generate fake data, and one to identify the fake data, when compared to real data.
- Diffusion networks: a network is trained to 'denoise' a noise sample.
- Boltzmann Machines, Fully Visible Belief Networks, Generative Stochastic Networks, and others.


Today we'll be discussing variational autoencoders.

A standard autoencoder network is actually just a pair of connected networks.

- These are an 'encoder' and a 'decoder'.
- The encoder network takes an input from a data set, \mathbf{x} , and converts it into a lower-dimensional form, \mathbf{z} (it is compressed into a 'latent space'). In a sense the encoding is a compression technique, based on the data itself.
- We've seen encoders before, we just didn't call them that.
- CNNs are encoders, at least the convolution/pooling sections of the networks.
- These convert the image into a dense form that is then used for classification, for example.
- In a standard autoencoder, the encoder is similar, in that it 'encodes' the input into a form which can be used as input into a another neural network.

Encoders are common starting points of many neural networks.





Standard autoencoders take this one step further: add a second network, the decoder, whose job is to take the encoded output (the latent variable, usually called z) and decode it back into the original input data. The goal is to make the generated data the same as the input data.

Some notes about autoencoders.

- The entire network is trained as a whole.
- The loss function is usually mean-squared error; it's known as the "reconstruction loss".
- Unfortunately, the encoded data is lossy, meaning that information is lost (thrown away).
- This is because the input data is much larger in size than what comes out of the encoder.
- But the encoder keeps enough of the important information that the essentials of the input can be rebuilt by the decoder.

Together these two networks form an autoencoder.

Why is this interesting?

- Autoencoders have a number of properties which are useful.
- Standard autoencoders are used in filtering and denoising applications, or other situations where we need to encode a definition of "normal".
- This is also a form of unsupervised learning, so the data can be unlabelled.
- Once the data is encoded by the encoder, we have a compressed "representative" form of the input data (latent space representation).
 - This can be used in clustering analyses.
 - This can be used as a dimensionality reduction technique.
 - The latent space vector can be manipulated and then put through the decoder, to useful (or amusing) effect.

Though it's not immediately obvious, this family of networks has uses.

But standard autoencoders have some shortcomings.

- Standard autoencoders can be trained to generate encoded representations of the input data, but other than a few niche applications, they aren't that commonly used.
- Why? Because the encodings which are generated by the encoder, which are 1D vectors, are often not continuous between one input and another.
- Why is that a problem? Because we don't want to just replicate the original data. We want to be able to create variations on the original data, continuously.
- If the encoded output is clustered, meaning that similar inputs end up getting encoded "close together", and the clusters are not close together, the decoder is not going to know what to do if I give it, as input, an encoding vector which is a combination of two inputs.


Creating new variations on the original data set is not possible in this scenario. The decoder just won't know what to do.

Create an autoencoder:

- make the latent space 2D
- train the encoder on MNIST
- run the test data through just the encoder
- plot the latent space output

The red circles are areas of latent space where there are no data. Moving into these areas will confuse the decoder.

This is an even bigger problem in higher dimensions.



Variational autoencoders are similar to standard autoencoders, but with a twist.

- By design, the space into which the encoder encodes is continuous.
- This allows the decoder to be able to generate output which combines different features of the data from the original data set.
- We can do this by feeding the decoder combinations of the encodings of different inputs.
- How is the encoder forced to encode continuously? By not outputting just a single encoding vector, but rather by outputting two vectors:
 - a vector of mean values
 - a vector of variances
- We now sample from the distributions represented by these means and variances to create an input vector, which is fed to the decoder.

This means that, even if we use the same input to the encoder, the decoder will see a different input due to the random sampling.

Variational autoencoders add randomness to the input of the decoder.

- For a given input to the encoder, the decoder ends up sampling a range of values which get associated with the decoder's output.
- This local variation of the input to the decoder removes some of the discontinuity in the decoder's input space, at least for a class of inputs.
- But this still doesn't get us continuity between different classes of input.
- We need the inputs to all be as close as possible to each other, while still being distinguishable.

How is that accomplished?

To force this, we need to visit a new topic: KL divergence.

- KL divergence (Kullback-Leibler divergence) is a measure of the "distance" between two probability distributions.
- To minimize this distance, in this case, means to optimize the probability distribution parameters which are output by the encoder to closely resemble that of the target distribution.
- To do this we need to calculate a new cost function.

This will require some math to explain properly.

Let:

- \mathbf{x} be our data set,
- \mathbf{z} be our latent variable (the output of the encoder),
- $p(\mathbf{x})$ be the probability distribution of the data set,
- $p(\mathbf{x}|\mathbf{z})$ be the distribution of the data given the latent variable (decoder output).
- $q_{\theta}(\mathbf{z}|\mathbf{x})$ be a variational distribution that we will try to match to $p(\mathbf{z}|\mathbf{x})$.

The goal is to get $p(\mathbf{x}) = \int p(\mathbf{x}|\mathbf{z})p(\mathbf{z})d\mathbf{z}$, but this integral is intractable.

To get around this, let us call the output of our encoder network $q_{\theta}(\mathbf{z}|\mathbf{x})$, which has parameters θ , and use this to approximate $p(\mathbf{z}|\mathbf{x})$ using a Bayesian technique called "Variational Inference".

$$\begin{aligned}\log(p(\mathbf{x})) &= \log(p(\mathbf{x})) \int q_{\theta}(\mathbf{z}|\mathbf{x}) d\mathbf{z} \\ &= \int \log(p(\mathbf{x})) q_{\theta}(\mathbf{z}|\mathbf{x}) d\mathbf{z} \\ &= E_{q_{\theta}(\mathbf{z}|\mathbf{x})} [\log(p(\mathbf{x}))] \\ &= E_{q_{\theta}(\mathbf{z}|\mathbf{x})} \left[\log \left(\frac{p(\mathbf{x}|\mathbf{z})p(\mathbf{z})}{p(\mathbf{z}|\mathbf{x})} \frac{q_{\theta}(\mathbf{z}|\mathbf{x})}{q_{\theta}(\mathbf{z}|\mathbf{x})} \right) \right]\end{aligned}$$

Where we have used Bayes' rule, and the $E_{q_{\theta}(\mathbf{z}|\mathbf{x})}$ is the expectation value over $q_{\theta}(\mathbf{z}|\mathbf{x})$. More formally, the right-hand side is known as the Evidence Lower Bound (ELBO).

$$\begin{aligned}\log(p(\mathbf{x})) &= E_{q_{\theta}(\mathbf{z}|\mathbf{x})} \left[\log \left(\frac{p(\mathbf{x}|\mathbf{z})p(\mathbf{z})}{p(\mathbf{z}|\mathbf{x})} \frac{q_{\theta}(\mathbf{z}|\mathbf{x})}{q_{\theta}(\mathbf{z}|\mathbf{x})} \right) \right] \\ &= E_{q_{\theta}(\mathbf{z}|\mathbf{x})} \left[\log(p(\mathbf{x}|\mathbf{z})) - \log \left(\frac{q_{\theta}(\mathbf{z}|\mathbf{x})}{p(\mathbf{z})} \right) + \log \left(\frac{q_{\theta}(\mathbf{z}|\mathbf{x})}{p(\mathbf{z}|\mathbf{x})} \right) \right] \\ &= E_{q_{\theta}(\mathbf{z}|\mathbf{x})} [\log(p(\mathbf{x}|\mathbf{z}))] - D_{KL} [q_{\theta}(\mathbf{z}|\mathbf{x})||p(\mathbf{z})] + \\ &\quad D_{KL} [q_{\theta}(\mathbf{z}|\mathbf{x})||p(\mathbf{z}|\mathbf{x})] \\ &\geq E_{q_{\theta}(\mathbf{z}|\mathbf{x})} [\log(p(\mathbf{x}|\mathbf{z}))] - D_{KL} [q_{\theta}(\mathbf{z}|\mathbf{x})||p(\mathbf{z})]\end{aligned}$$

Where the D_{KL} terms are the KL divergence. We now drop one of the terms, since KL divergences are always positive. The $D_{KL} [q_{\theta}(\mathbf{z}|\mathbf{x})||p(\mathbf{z}|\mathbf{x})]$ term represents the uncertainty in the final model $q_{\theta}(\mathbf{z}|\mathbf{x})$.

And there it is, the VAE loss function!

Formally, the right-hand side is known as the Evidence Lower Bound (ELBO).

$$\log(p(\mathbf{x})) \geq E_{q_{\theta}(\mathbf{z}|\mathbf{x})} [\log(p(\mathbf{x}|\mathbf{z}))] - D_{KL}[q_{\theta}(\mathbf{z}|\mathbf{x})||p(\mathbf{z})]$$

How do we interpret this?

- $\log(p(\mathbf{x}))$ represents the data set.
- We should maximize $E_{q_{\theta}(\mathbf{z}|\mathbf{x})} [\log(p(\mathbf{x}|\mathbf{z}))]$.
- We should minimize the difference between $q_{\theta}(\mathbf{z}|\mathbf{x})$ and $p(\mathbf{z})$.
- We can approximate $p(\mathbf{x}|\mathbf{z})$ using the decoder half of the network.
- We can approximate $q_{\theta}(\mathbf{z}|\mathbf{x})$ using the encoder half of the network.

We can use the ELBO to craft a loss function.

How do we minimize the difference between $q_{\theta}(\mathbf{z}|\mathbf{x})$ and $p(\mathbf{z})$?

- $p(\mathbf{z})$ is the latent variable distribution. We can model this using $N(\mathbf{0}, \mathbf{1})$ (Gaussian with $\mu = \mathbf{0}$ and $\sigma^2 = 1$).
- We model $q_{\theta}(\mathbf{z}|\mathbf{x})$ as a Gaussian with parameters $\mu(\mathbf{x})$ and $\Sigma(\mathbf{x})$.
- Then the KL divergence between those two pieces can be calculated exactly.

$$D_{KL} [N(\mu(\mathbf{x}), \Sigma(\mathbf{x})) || N(\mathbf{0}, \mathbf{1})] = \frac{1}{2} \sum_k (\Sigma(\mathbf{x}) + \mu^2(\mathbf{x}) - 1 - \log(\Sigma(\mathbf{x})))$$

Where k is the number of dimensions in \mathbf{z} , and $\Sigma(\mathbf{x})$ is the diagonal of the usual covariance matrix. We use this, combined with the regular loss function, to train the network.

So what is the approach? Create a network with the following architecture.

- Build an encoder.
- As we mentioned earlier, the encoder will output two vectors, a vector of mean values ($\mu(\mathbf{x})$), and variances ($\Sigma(\mathbf{x})$).
- To create the latent variable, \mathbf{z} , to input into the decoder, we sample randomly from the Gaussians with means of $\mu(\mathbf{x})$ and variances $\Sigma(\mathbf{x})$.
- This \mathbf{z} is fed into the decoder, which generates an output.
- To train, we use two loss functions:
 - $L_{\text{reconstruction}} = \frac{1}{2} (\mathbf{x}_{\text{in}} - \mathbf{x}_{\text{out}})^2$
 - $L_{\text{KL}} = \frac{1}{2} \sum_k (\Sigma(\mathbf{x}) + \mu^2(\mathbf{x}) - 1 - \log(\Sigma(\mathbf{x})))$
- Train the network in the usual way.

But, as you might expect, there is a slight problem with what's been described here.

How do you do backpropagation through a function which is sampled?

- Sampling doesn't have a gradient.
- There is a way around this, called the "reparameterization trick".

Suppose that x is sampled from $N(\mu(x), \Sigma(x))$, and suppose that we standardize it, so that $\mu = 0$ and $\Sigma = 1$.

$$x_{\text{std}}^2 = (x - \mu) \Sigma^{-1} (x - \mu)^T$$

This leaves $x = \mu + \Sigma^{\frac{1}{2}} x_{\text{std}}$. We do the same to sample our latent variable:

$$z = \mu(x) + \Sigma^{\frac{1}{2}}(x)\epsilon$$


where ϵ is sampled from $N(0, 1)$. This moves the non-differentiable part of the operation out of the network, so that the network can still be trained in the usual way.

Let's do an example. We'll use our old friend, MNIST.


- We will use convolutional and transpose convolutional layers.
- We will have a latent space size of 100.
- We will use mean squared error as our reconstruction cost function, combined with the KL divergence.
- Rather than have the encoder calculate the variances of the mean values it outputs, we will instead have it calculate the log of the variances (it is more numerically stable to calculate an exponential than a logarithm).
- We will build the encoder and decoder as two separate networks, so that they can be saved, and later used, separately. They will be trained together, obviously.
- We normalize the input data, and then use sigmoid on the output of the decoder.

Our decoder is going to sample from the latent space to create an image. As you might imagine, to do this we need to deconvolve the data. How does that work?

- Suppose we start with a 1D vector of latent values.
- Reshape the vector into a square.
- Convolution involves taking a, say, (5×5) square of an image and processing it into a (1×1) square (a single point).
- Deconvolution involves taking a say, (2×2) square of an image and processing it into something larger, (4×4) for example.
- There are several ways to do this. The "transpose convolution" technique involves either
 - padding the image with zeros around the outside,
 - or 'upsampling' the image to double its size in each dimension,and then doing the usual convolution.



The green padding consists of zeros. The deconvolution is the convolution of a padded version of the input data. It is equivalent to the transpose of the convolution of a (6×6) layer using a (3×3) filter, resulting in a (4×4) layer. The weights and biases of each deconvolution feature map are the same.



The default Keras upsampling layer doubles all data in all dimensions. Upsampling is often used instead of just padding with zeros, since there's more information available.

```
# mnist_vae.py
import tensorflow.keras.models as km
import tensorflow.keras.layers as kl
import tensorflow.keras.backend as K
from tensorflow.keras.datasets import mnist

def sampling(args):
    z_mean, z_log_var = args
    batch_size = K.shape(z_mean)[0]
    dim = K.int_shape(z_mean)[1]
    eps = K.random_normal(shape = (batch_size, dim))
    return z_mean + K.exp(0.5 * z_log_var) * eps

(x_train, _), (x_test, _) = mnist.load_data()
x_train = x_train.astype('float32')
x_test = x_test.astype('float32')
x_train = x_train.reshape(60000, 28, 28, 1) / 255.
x_test = x_test.reshape(10000, 28, 28, 1) / 255.
```

```
# mnist_vae.py, continued
latent_dim = 100
input_img = kl.Input(shape = (28, 28, 1))
x = kl.Conv2D(16, kernel_size = (3, 3),
    activation = "relu",
    padding = "same")(input_img)
x = kl.MaxPooling2D(pool_size = (2, 2),
    stride = 2)(x)
x = kl.Conv2D(32, kernel_size = (3, 3),
    activation = "relu", padding = "same")(x)
x = kl.MaxPooling2D(pool_size = (2, 2),
    stride = 2)(x)
x = kl.Flatten()(x)

z_mean = kl.Dense(latent_dim,
    activation = "linear")(x)
z_log_var = kl.Dense(latent_dim,
    activation = "linear")(x)
```

```
# mnist.vae.py, continued
```

```
z = kl.Lambda(sampling)([z_mean, z_log_var])
```

```
# Our custom loss function.
```

```
def vae_loss(input_img, output_img):
```

```
    kl_loss = 1 + z_log_var
```

```
    - K.square(z_mean) - K.exp(z_log_var)
```

```
    kl_loss = -0.5 * K.sum(kl_loss, axis = -1)
```

```
    r_loss = 784 * K.mean(K.square(input_img -  
                                output_img))
```

```
    return r_loss + kl_loss
```

```
# mnist.vae.py, continued
```

```
decoder_input = kl.Input(shape = (latent_dim,))
```

```
x = kl.Dense(7 * 7 * 32,  
             activation = "relu")(decoder_input)
```

```
x = kl.Reshape((7, 7, 32))(x)
```

```
x = kl.Conv2DTranspose(32, kernel_size = (3, 3),  
                      activation = "relu", padding = "same")(x)
```

```
x = kl.UpSampling2D(size = (2, 2))(x)
```

```
x = kl.Conv2DTranspose(16, kernel_size = (3, 3),  
                      activation = "relu", padding = "same")(x)
```

```
x = kl.UpSampling2D(size = (2, 2))(x)
```

```
decoded = kl.Conv2DTranspose(1, kernel_size = (3, 3),  
                           activation = "sigmoid", padding = "same")(x)
```

```
# mnist_vae.py, continued
```

```
encoder = km.Model(inputs = input_img,  
                   outputs = [z_mean, z_log_var, z])
```

```
decoder = km.Model(inputs = decoder_input,  
                   outputs = decoded)
```

```
output_img = decoder(encoder(input_img)[2])  
vae = km.Model(inputs = input_img,  
               outputs = output_img)
```

```
vae.compile(loss = vae_loss,  
            optimizer = "rmsprop",  
            metrics = ["mean_squared_error"])
```

```
# mnist_vae.py, continued
```

```
fit = vae.fit(x_train, x_train, epochs = 50,  
             batch_size = 64, verbose = 2)
```

```
score = vae.evaluate(x_test, x_test)  
print("score is", score)
```

```
# Save the final models.
```

```
vae.save('mnist_vae.h5')  
encoder.save('mnist_encoder.h5')  
decoder.save('mnist_decoder.h5')
```

VAE example, running the code


```
ejspence@mycomp ~> python mnist_vae.py
Epoch 1/50
13s - loss: 43.2175 - mean_squared_error: 0.0438
Epoch 2/50
8s - loss: 35.0138 - mean_squared_error: 0.0307
:
Epoch 48/50
8s - loss: 28.3416 - mean_squared_error: 0.0203
Epoch 49/50
8s - loss: 28.3258 - mean_squared_error: 0.0203
Epoch 50/50
9s - loss: 28.3135 - mean_squared_error: 0.0203
9600/10000 [=====>..] - ETA: 0s
score is [28.21702071533203, 0.019792119041085243]
ejspence@mycomp ~>
```

What can we do with it?

Because we know the labels of the input data (the digits 0-9), we can play around with encoder and decoder to interesting effect. For example:

- Using the encoder, we can average over the encoder-output means of all the inputs of the same label (all the sixes, for example).
- We then have a set of 10, latent vectors which represent each digit.
- We can then linearly interpolate from one latent vector to another, and run those interpolations through the decoder.
- If the decoder is any good, as we transition from one latent vector to the next we should smoothly transition from one digit to the next.

In this way, we can manipulate data or images of different types, blending, subtracting and modifying images in a predictable way.



Variational autoencoders:

- <https://arxiv.org/abs/1312.6114>
- <https://arxiv.org/abs/1401.4082>
- <https://wiseodd.github.io/techblog/2016/12/10/variational-autoencoder>
- <https://blog.keras.io/building-autoencoders-in-keras.html>
- <http://kvfrans.com/variational-autoencoders-explained>
- <https://towardsdatascience.com/intuitively-understanding-variational-autoencoders-1bfe67eb5daf>
- <https://jaan.io/what-is-variational-autoencoder-vae-tutorial>

Deconvolution:

- <https://arxiv.org/abs/1603.07285v1>