

# A comparison of classical and variational autoencoders for anomaly detection

Technical Report IDSIA-2020-10.1

Fabrizio Patuzzo

IDSIA / USI-SUPSI

Manno, Switzerland

September 2020

## Abstract

This paper analyzes and compares a classical and a variational autoencoder in the context of anomaly detection. To better understand their architecture and functioning, describe their properties and compare their performance, it explores how they address a simple problem: reconstructing a line with a slope.

## Introduction

An autoencoder is a neural network that receives an input (usually an image), compresses it into a short code (the ‘bottleneck’) and outputs a reconstruction of the image. They are used to allow computers to produce works of art (by modifying the values in the bottleneck, the network creates new, unseen pictures), and to detect anomalies: an autoencoder can be trained to reproduce properly only clean images, but not defective ones, making it possible to detect the defect. There exist two types of autoencoders: classical and variational. Variational autoencoders have been particularly successful to generate art.

The aim of this short paper is to better understand how classical and variational autoencoders work in the anomaly detection domain, and to compare their performance. To do this, we will look at how they tackle a simple problem: reconstructing a line with a slope.

## 1 Training Set

The training set is composed of 1000 black and white, 32x32 pixel images, representing a straight line with a given slope that divides the image into white

and black halves. Figure 1 shows a few images from the training set (plotted using the colormap ‘viridis’).



Figure 1: A few samples from the training set

## 2 Classical Autoencoder

The classical autoencoder (CA) is a Multi Layer Perceptron (MLP) with the following architecture:

LAYER	type	neurons	activation
0	Input	32x32	
1	Dense	2	relu
2 (bottleneck)	Dense	1	
3	Dense	2	relu
4	Dense	32x32	sigmoid

The input is a 32x32 px black and white image, and all layers are Dense. L1 and L3 contain two neurons and are followed by a *relu* activation function, defined as  $\max(0, x)$ . L2, the bottleneck, contains only one neuron. The final layer is followed by a *sigmoid* function:  $\frac{e^x}{e^x+1}$ .

We trained the CA to recompose the input images in the output. Figure 2 shows the results. The reconstructions are not precise: there is some blur along the dividing lines, and some sort of strange kaleidoscopic effects in the first two images.




Figure 2: Results using A1

To try to pinpoint the cause of this blur, we have created an even simpler dataset composed of four 2x2 images, where the slope can either be horizontal or vertical, shown in Figure 3.

We trained the CA again (adapting its input and output dimensions) and, after several failed attempts (more on this later), it finally succeeded in reconstructing



Figure 3: Simplified training set

the inputs: ‘0011’, ‘1010’, ‘1100’ and ‘0101’<sup>1</sup>.

Let us see which weights the network has learned<sup>2</sup>.

1. Figure 4 shows the weights leading to the first output neuron.




Figure 4: weights leading to output neuron 1

The output of the first neuron in L4 as a function of the bottleneck is:

$$y = \sigma(\text{relu}(10x) \cdot (-1) + \text{relu}(-10x) \cdot (-1) + b). \quad (1)$$

where  $\sigma$  is the sigmoid,  $x$  is the bottleneck value and  $b$  the bias.

Let us approximate the sigmoid with a Heaviside function, which returns 0 if  $x < 0$ , and 1 if  $x \geq 0$ .

Then, when the bottleneck equals  $-2$ , the output is  $H(-5) = 0$ ; when it equals 0 the output is  $H(15) = 1$ ; and a bottleneck of 2 generates an output of  $H(-5) = 0$ .

We can express this relation in tabular form.

BN	$-\infty$	$-1.5$	$1.5$	$+\infty$
output 1	0	1	0	

The thresholds,  $-1.5$  and  $1.5$ , can be obtained by computing  $t_1 = -\frac{b}{10}$  and  $t_2 = \frac{b}{10}$ .

<sup>1</sup>the four digits correspond to the pixel values in the upper left, upper right, bottom left and bottom right corners.

<sup>2</sup>To simplify the discussion, we will assume that all biases in L1, L2 and L3 equal zero.

2. Figure 5 illustrates the weights leading to the second output neuron.




Figure 5: weights leading to output neuron 2

This time, any strictly negative bottleneck will output a 0, and any positive bottleneck will output a 1.

The relation between the bottleneck and output 2 is:

BN	$-\infty$	0	$+\infty$
output 2	0	1	

3. In Figure 6, one can see the connections from the bottleneck to output neuron 3.




Figure 6: weights leading to output neuron 3

This is exactly the reversed of output 2. Notice that in all the inputs, the second and the third digit are complementary.

The relation between BN and output 3 is:

BN	$-\infty$	0	$+\infty$
output 3	1	0	

4. Figure 7 shows the weights that lead from the bottleneck to output neuron 4.




Figure 7: connections leading to output neuron 4

This time, output 4 is the opposite of output 1 and the relation between BN and output 4 is:

BN	$-\infty$	-1.5	1.5	$+\infty$
output 4	1	0	1	

If we increase BN evenly from  $-\infty$  to  $+\infty$ , the network will output a copy of the inputs: ‘0011’, ‘1010’, ‘1100’ and ‘0101’. The output will change when BN crosses the thresholds -1.5, 0 and 1.5.

But let us remember, we have approximated the sigmoid with a Heaviside function. The only difference if we use a true sigmoid is that as the bottleneck approaches a threshold, some output neurons will become temporarily gray, as shown in Figure 8.



Figure 8: outputs as BN crosses the critical value 1.5

Hence, some ‘blur’ can occur when BN is very close to a threshold. To remove the blur, we can replace the sigmoid with a Heaviside function.

Back to our initial dataset. The logic of the network is exactly the same, only there are more output neurons, each one with its own bias  $b$  that will determine at which thresholds  $t_1$  and  $t_2$  the neuron value changes. As we increase the bottleneck value, some neurons will change value and become temporarily gray, causing the blur in the figure.

Let us try to apply the same trick to remove the blur.

By replacing the sigmoid with a Heaviside function after the training, the blur disappears, as well as the strange kaleidoscopic effects, as one can see in Figure 9.



Figure 9: outputs after the postprocessing

### 3 Variational Autoencoder

Let us see now how a variational autoencoder (VA) would perform.

The architecture of a VA is similar to the architecture of a CA, but with the little modification shown in Figure 10<sup>3</sup>.




Figure 10: The variational architecture

A new layer is added just before the bottleneck, composed of two neurons,  $a$  and  $b$ . The bottleneck value is created by picking it from a random normal distribution with mean  $a$  and variance  $e^{b/2}$ .

Not only the architecture, but also the *loss* function used to train the network is different. In a CA it is just the mean squared error (*mse*), but in a VA it includes a new term<sup>5</sup>:

$$loss = mse - \frac{1}{2n} \cdot \sum_1^n (1 + b - a^2 - e^b) \quad (2)$$

where  $n$  is the number of training samples.

The two changes in the architecture and in the *loss* have an impact on the characteristics of the autoencoder:

<sup>3</sup>Our implementation of the variational autoencoder is inspired by Chollet [2019].

<sup>4</sup>Notice that in Chollet [2016] the implementation is slightly different, and  $e^b$  represents the standard deviation rather than the variance.

<sup>5</sup>The mathematical derivation of the new term, obtained computing the Kullback-Leibler divergence between two normal distributions, is explained in Kingma and Welling [2014], the paper that set the mathematical foundations of the variational architecture.

1. Every input image now produces a bottleneck value picked from a normal distribution with  $\mu = a$  and  $\sigma^2 = e^b$ .
2. The new *loss* ensures that the bottleneck values generated by the training set will follow a normal distribution with  $\mu = 0$  and  $\sigma^2 = 1$ , and, as a consequence, they will remain within the range  $[-3, 3]$  with a probability of 0.97.

After training VA with our dataset, we could verify that the bottleneck values approach a standard normal distribution, as shown in Figure 11a. In Figure 11b one can see the standard deviations ( $\sigma$ ). They are all around 0.01.

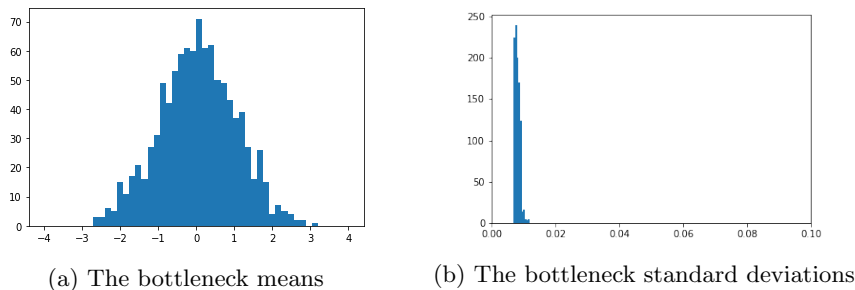


Figure 11

From the bottleneck on, the CA and VA are identical. Hence, we might expect that a VA will find the same kind of weights as a CA. However, since the bottleneck values are packed together in the interval  $[-3,3]$ , we might expect that they will be close to the thresholds, and thus cause more blur. Plus, they are generated randomly, so our first guess is that the outputs can only be more blurry.

Let us try, and apply a VA to reconstructing the four sample images of the training set. The results are shown in Figure 12. They are nearly perfect: no blurring, and no strange kaleidoscopic effects.



Figure 12: Outputs using VA

Our first guess was wrong. Why is the output not more blurry? And why is it even better?

The fact that the bottlenecks are close together can be easily compensated by finding big weights in L3 and L4, which the network did. And the standard deviations are small enough to ensure that a bottleneck will not cross a threshold.

But why did the VA perform even better? Because it found a set of optimal weights which the CA did not find, as it converged to a suboptimal solution.

## 4 Converging to a suboptimal solution

Even when trained on the simplified dataset of four 2x2 images, the CA did not always find an optimal solution. Figure 13 shows what an optimal solution might look like.




Figure 13: optimal solution

This set of weights entirely solves the problem, and the four inputs ‘0011’, ‘1010’, ‘1100’ and ‘0101’ will output copies of themselves, and generate bottleneck values of  $-10$ ,  $-1$ ,  $1$  and  $10$ .

But despite the fact that optimal weights exist, when we train the CA it rarely finds them. Let us see some of the solutions it converges to.

1. With an observed frequency of  $\frac{1}{8}$ , the CA did find a set of optimal weights.
2. About  $\frac{1}{3}$  of the times, it failed to reconstruct two images, producing an ‘average’ of the two, as shown in Figure 14.



Figure 14: suboptimal solution 1

What happened? After weight initialization, two inputs happened to generate negative values in all L1 neurons (which the relu converted to zeros). As a consequence, they produced the exact same output, and no weight update could improve the situation.<sup>6</sup>

<sup>6</sup>The probability that an input will generate a negative number in at least one neuron in L1 after a random weight initialization is  $\frac{1}{4}$ , and the probability that this happens to at least two neurons is  $1 - (3/4)^4 - 4 \cdot (3/4)^3 \cdot 1/4 \approx 0.26$



3. Occasionally, this happened with two pairs of images, as illustrated in Figure 15.



Figure 15: suboptimal solution 2

4. About  $\frac{1}{3}$  of the times, all the outputs were completely gray, as shown in Figure 16.



Figure 16: suboptimal solution 3

This can happen if all weights in L1 are negative. But there can be other reasons. For example, it can also happen if both weights in L2 are negative and both weights in L3 are positive (or viceversa)<sup>7</sup>.

These are some of the problems that can happen. In all these cases, two inputs generated the same values after some layer in the network.

One way to reduce the number of times this happens is to increase the number of neurons in L1, L2 and L3 (but too many neurons will cause other difficulties).

Another way, possibly, is to use a VA. For, due to its properties, a VA will never produce exactly the same bottleneck values, preventing the training from stalling. Our guess is that this is the reason why the VA output sharper images.

## 5 Conclusions

We have explored how a classical and a variational autoencoder deal with a simple problem, trying to describe their behaviour and characteristics.

The CA often converged to suboptimal solutions. It reconstructed the images with some blur and some kaleidoscopic effects, even though, in this case, we could remedy by replacing the sigmoid with a Heaviside function.

In contrast, the VA found much better solutions, and we suggested that this was possibly due to the random generation of bottlenecks.

If one gives an input to the VA that contains a defect or does not represent a

<sup>7</sup>This last scenario happens  $\frac{1}{8}$  of the times right after weight initialization

line, it will still output a line, and the reconstruction error will allow to detect the anomaly.

In conclusion, our initial doubts about the ability of a VA to perform better than a CA in anomaly detection were unfounded.

## Acknowledgments

We would like to thank our supervisor Alessandro Giusti, and our colleagues Jamal Saeedi and Gabriele Abbate for their useful comments and feedback.

## References

- F. Chollet. Building autoencoders in keras, 2016. URL <https://blog.keras.io/building-autoencoders-in-keras.html>.
- F. Chollet. Keras github repository, 2019. URL <https://keras.io/examples/generative/vae/>.
- D. P. Kingma and M. Welling. Auto-encoding variational bayes. *CoRR*, abs/1312.6114, 2014.