
Simon
Walser

Student	Simon Walser
Examiner	Prof. Dr. Guido Schuster
Subject Area	Software and Systems

Signal modelling using unsupervised learning

Variational Autoencoder applied to ECG signals for feature recognition and signal generation

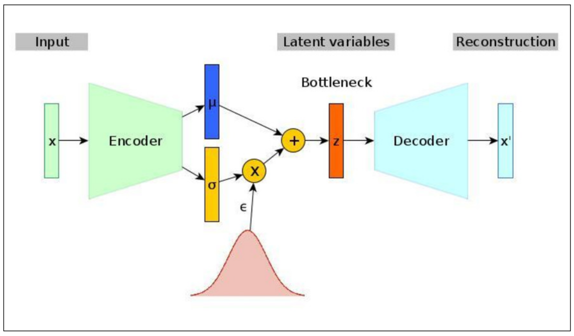

Illustration of a VAE. Encoder reduces input data to the essence while Decoder reconstructs the input from this info.
Own presentment

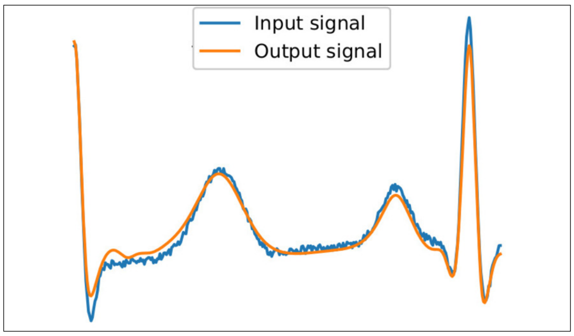

Input ECG signal and artificially generated output signal.
Own presentment

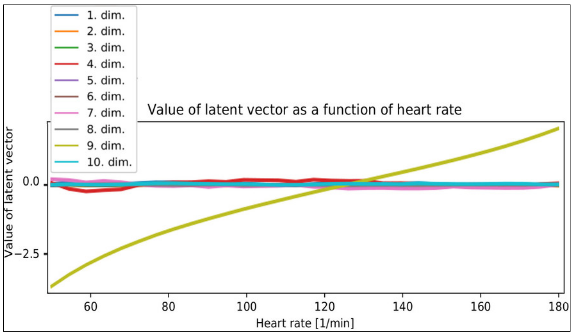

Response of bottleneck variables to change of heart rate. Info is found automatically and encoded in the latent vector.
Own presentment

Initial Situation: One of the biggest challenges associated with training artificial neural networks is collecting sufficient data to enable the generalization of the neural network. For a neural network to work well when fed with data it has never seen before, it needs to know the underlying distribution of the data as well as possible. For more demanding tasks, this requires a very large amount of data. It would therefore be highly desirable to determine the most important variables, which describe the data generating distribution and the relationship between these variables and the samples of the distribution. This would allow to generate artificial data, resulting from a certain constellation of latent variables.

Approach: Under the title "Beta-VAE: learning basic visual concepts with a constrained variational framework" a paper has been published which introduces a new framework, which is called Beta Variational Autoencoder. According to the authors of this paper, the Beta-VAE is used for automated discovery of interpretable factorized latent representations from raw image data in a completely unsupervised manner. This is achieved by connecting an encoder and decoder network. The encoder network maps the high-dimensional input data into a low-dimensional latent representation, whereas the decoder network reconstructs the input data from the low-dimensional latent representation.

The results presented in the paper look promising. However, the tests were made exclusively with images as input data. The aim of this thesis is to test the application of a Beta-VAE to sequential data such as ECG signals.

Result: In a first step, the basic functionality of a Beta-VAE built with Gated Recurrent Units (GRU's) was tested. These experiments were carried out with sinusoidal waves whose amplitude and frequency were the only variables. The result was surprisingly good. Both the unsupervised discovery of the two variables and the reconstruction of the sinusoidal signals based on the latent variables worked very well under reasonable conditions. In a second step, the model was fed with ECG signals to test its performance in a more useful and less deterministic environment. The ECG signals used for the tests were artificially generated, to control the number and type of variables which describe the input signal. Also for ECG signals, the Beta-VAE with GRU's works very well. An optimally tuned autoencoder could learn the signal generative factors and reconstruct input signals under reasonable conditions. In conclusion, it was found that the number of generative factors does not need to be known for a successful disentanglement of the latent variables. This is learned automatically and correctly by the Beta-VAE.

The tests show that the Recurrent Neural Network (RNN) based Beta-VAE can be a valid model to find uncorrelated, interpretable variables describing ECG data. This can be a strong support to train networks based on ECG signals. Not only provides the autoencoder a generator of training signals but it can also make signal labelling much easier.