

# MA-VAE: Multi-Head Attention-Based Variational Autoencoder Approach for Anomaly Detection in Multivariate Time-Series Applied to Automotive Endurance Powertrain Testing

Lucas Correia<sup>1,2</sup><sup>a</sup>, Jan-Christoph Goos<sup>1</sup><sup>b</sup>, Philipp Klein<sup>1</sup>, Thomas Bäck<sup>2</sup><sup>c</sup>  
and Anna V. Kononova<sup>2</sup><sup>d</sup>

<sup>1</sup>Mercedes-Benz AG, Stuttgart, Germany

<sup>2</sup>Leiden University, Leiden, The Netherlands

**Keywords:** Anomaly Detection, Multivariate, Time Series, Automotive, Test Bench, Variational Autoencoder, Bypass Phenomenon.


**Abstract:** A clear need for automatic anomaly detection applied to automotive testing has emerged as more and more attention is paid to the data recorded and manual evaluation by humans reaches its capacity. Such real-world data is massive, diverse, multivariate and temporal in nature, therefore requiring modelling of the testee behaviour. We propose a variational autoencoder with multi-head attention (MA-VAE), which, when trained on unlabelled data, not only provides very few false positives but also manages to detect the majority of the anomalies presented. In addition to that, the approach offers a novel way to avoid the bypass phenomenon, an undesirable behaviour investigated in literature. Lastly, the approach also introduces a new method to remap individual windows to a continuous time series. The results are presented in the context of a real-world industrial data set and several experiments are undertaken to further investigate certain aspects of the proposed model. When configured properly, it is 9% of the time wrong when an anomaly is flagged and discovers 67% of the anomalies present. Also, MA-VAE has the potential to perform well with only a fraction of the training and validation subset, however, to extract it, a more sophisticated threshold estimation method is required.


## 1 INTRODUCTION


Powertrain testing is an integral part of the wider automotive powertrain development and is undertaken at different stages of development. Each of these stages is composed of many integration levels. These integration levels range from powertrain sub-component testing, such as the electric drive unit (EDU) controller or high-voltage battery (HVB) management system, to whole vehicle powertrain testing. Each of these has its special type of controlled environment, called a test bench. The use-case in this paper is on an endurance powertrain test bench, where the EDU and HVB on their own are tested under different conditions and loads for longer periods to simulate wear over time. Given the costly maintenance and upkeep costs of such test benches, it is desirable to keep


downtime at a minimum and to avoid faulty measurements. Also, it is desirable to detect problems early to prevent damage to the testee. Given that evaluation is done manually by inspection, it is not feasible to analyse every single measurement, also evaluation tends to be delayed, only being undertaken days after the measurement is recorded, hence there is a clear need for automatic, fast and unsupervised evaluation methodology which can flag anomalous measurements before the next measurement is started.

To achieve this, we propose a multi-head attention variational autoencoder (MA-VAE). MA-VAE consists of a bidirectional long short-term memory (BiLSTM) variational autoencoder architecture that maps a time-series window into a temporal latent distribution (Park et al., 2018) (Su et al., 2019). Also, a multi-head attention (MA) mechanism is added to further enhance the sampled latent matrix before it is passed on to the decoder. As shown in the ablation study, this approach avoids the so-called bypassed phenomenon (Bahuleyan et al., 2018), which is the first contribu-

<sup>a</sup> <https://orcid.org/0000-0002-3653-5934>

<sup>b</sup> <https://orcid.org/0009-0002-7063-3615>

<sup>c</sup> <https://orcid.org/0000-0001-6768-1478>

<sup>d</sup> <https://orcid.org/0000-0002-4138-7024>

tion. Furthermore, this paper offers a unique methodology for the reverse-window process. It is used for remapping the fixed-length windows the model is trained on to continuous variable-length sequences.

This paper is structured as follows: First, a short background is provided in Section 2 on the powertrain testing methodology specific to this use case, as well as the theory behind VAE and MA mechanisms. Then, related work in variational autoencoder-based time-series anomaly detection is presented in Section 3, followed by an in-depth introduction of the real-world data set and the approach we propose in Section 4. Then, several experiments testing different aspects of the proposed method are conducted and discussed in Section 5, along with the final results. Finally, conclusions from this work are drawn and an outlook into future work is provided in Section 6. The source code for the data pre-processing, model training as well as evaluation can be found under <https://github.com/lcs-crr/MA-VAE>.

## 2 BACKGROUND

### 2.1 Real-World Application

During endurance testing a portfolio of different driving cycles is run, where a cycle is a standardised driving pattern, which enables repeatability of measurements. For this type of testing the portfolio consists exclusively of proprietary cycles, which differ from the public cycles used, for example, for vehicle fuel/energy consumption certification like the New European Driving Cycle (NEDC) or the Worldwide Harmonised Light Vehicles Test Cycle (WLTC). The reason why proprietary cycles are used for endurance runs is that they allow for more extensive loading of the powertrain.

Given the presence of a battery in the testee, some time has to be dedicated to battery soaking (sitting idle) and charging. These procedures are also standardised using cycles, although, for the intents and purposes of this paper, they are omitted. What is left are the eight dynamic driving cycles representing short, long, fast, slow and dynamic trips ranging from 5 to 30 minutes. There are multiple versions of the same cycle, which mostly differ in starting conditions such as state-of-charge (SoC) and temperature of the battery.

On powertrain test benches, there are several control methods to ensure the testee maintains the given driving cycle. In this particular test bench, the regulation is done by the acceleration pedal and the EDU revolutions-per-minute (rpm), which is nothing more

than a non-SI version of the angular velocity.

### 2.2 Data Set

This real-world data set consists of 3385 normal measurement files, each of which contains hundreds of (mostly redundant or empty) channels. A measurement is considered normal when the testee behaviour conforms to the norm. For this work, a list of  $d_X = 13$  channels was hand-picked in consultation with the test bench engineers to choose a reasonable and representative number of channels. This list includes the vehicle speed, EDU torque, current, voltage, rotor temperature and stator temperature, left and right wheel shaft torque, HVB current, voltage, temperature and SoC and inverter temperature. Given that some channels (such as torque) are sampled much faster than others (like temperature and SoC), a common sampling rate of 2Hz is chosen. Channels sampled slower than 2Hz are linearly interpolated, which is seen as permissible due to the lower amplitude resolution of those channels. Channels sampled faster than 2Hz are passed through a low-pass filter with a cut-off frequency of 1Hz and then resampled to 2Hz, as is consistent with the Whittaker–Nyquist–Shannon theorem (Shannon, 1949). Then the driving cycles are z-score normalised, i.e. transformed such that the mean for each channel lies at 0 and the standard deviation at 1. Lastly, the driving cycles (generally referred to as *sequences* in this paper) are windowed to create a set of fixed-length sub-sequences, or *windows*. First, each channel is auto-correlated to obtain the number of lags of the slowest dynamic process present in the signal. Then, the window size  $W$  is set as the smallest power of two larger than the longest lag, in this case,  $W = 256$  time steps or 128 seconds. Each window overlaps their preceding and succeeding windows by half a window, i.e. the shift between windows is  $W/2 = 128$  time steps, in order to reduce computational load compared to a shift of one time step.

Due to the absence of labelled anomalies in the dataset, realistic anomalous events are intentionally simulated and recorded following the advice of test bench engineers. To this end, five anomaly types were recorded. In the first type, the virtual wheel diameter is changed, such that the resulting vehicle speed deviates from the norm. The wheel diameter is a parameter as resistances are connected to the shafts rather than actual wheels. The second type of anomaly involves changing the driving mode from comfort to sport, which leads to a higher HVB SoC drop over the cycle and a different torque response. In the third anomaly, the recuperation level is turned from maxi-

mum to zero, hence the minimum EDU torque is always non-negative and the HVB SoC experiences a higher drop in SoC. In the case of the fourth anomaly, the HVB is swapped for a battery simulator, where the HVB voltage behaviour deviates from a real battery. The inverter and EDU share a cooling loop, whose cooling capacity is reduced at the beginning or middle of the cycle, leading to higher EDU rotor, EDU stator and inverter temperatures than normal. Every anomaly type is recorded during every cycle at least once, leading to 60 anomalous driving cycles that are all used as the anomalous subset of the test set.

A plot of one normal and one wheel-diameter anomalous cycle is shown in Figure 1. Due to the long channel names, the plot only shows the channel indices, a table containing the legend is shown in Table 1 for context. Visual inspection may suggest that the red plot is anomalous, since the EDU and HVB voltage, temperature and state of charge deviate from the black plot. This deviation is to be expected because they depend on how charged the battery is and on how much the battery is used previous to the current cycle. In the case of this anomaly, the only channel that demonstrates anomalous behaviour is the vehicle speed, since:

$$v_{\text{vehicle}} = r \times \omega \quad (1)$$

where  $r$  is the wheel radius and  $\omega$  the angular velocity. Evidently, the anomalous behaviour is most visible at higher speeds.

In an operative environment, it is desirable to find out whether the previously recorded sequence had any problems to analyse before the next measurement is recorded. Also, a model that performs as well as possible with as little data as possible translates to faster deployment. *Good* performance is indicated by a model that can detect as many anomalies as possible and rarely labels normal measurements wrongly.

Table 1: Legend for the channel names in Figure 1.

No.	Name
1	Vehicle Speed
2	EDU Torque
3	Left Axle Torque
4	Right Axle Torque
5	EDU Current
6	EDU Voltage
7	HVB Current
8	HVB Voltage
9	HVB Temperature
10	HVB State of Charge
11	EDU Rotor Temperature
12	EDU Stator Temperature
13	Inverter Temperature




Figure 1: Features of a normal (black) and an anomalous (red) cycle plotted with respect to time. The anomalous cycle plotted represents a scenario where the wheel diameter has not been set correctly. The amplitude axis is z-score normalised to comply with confidentiality guidelines.

To investigate the required training subset size of the model, it is trained with 1h, 8h, 64h, and 512h worth of dynamic testing data, which corresponds to the first 6, 44, 348, and 2785 driving cycles, respectively. The results are also presented in Section 5. In each of the above-mentioned cases, the training subset is further split into a training (80%) and a validation (20%) subsets. Both the training and validation subsets are batched to sets of 512 windows. Given the anomalous subset size of 60 driving cycles, 600 normal driving cycles recorded after the ones in the training subset are chosen to make up the normal test subset. This would imply that 9% of measurements at the test bench are anomalous, however in reality this value is estimated to be much lower. This amount of anomalous data in relation to normal data is used as

it approximately matches the anomaly ratio in public data sets and because the data set is not large enough to create a larger normal test subset.

### 2.3 Variational Autoencoders

The variational autoencoder (Kingma and Welling, 2014)(Rezende et al., 2014) is a generative model that structurally resembles an autoencoder, but is theoretically derived from variational Bayesian statistics. As opposed to the regular deterministic autoencoder, the VAE uses the evidence lower bound (ELBO), which is a lower bound approximation of the so-called log evidence  $\log p_\theta(\mathbf{X})$ , as its objective function. The ELBO, Equation 2, can be expressed as the reconstruction log-likelihood and the negative Kullback-Leibler Divergence ( $D_{KL}$ ) between the approximate posterior  $q_\phi(\mathbf{Z}|\mathbf{X})$  and the prior  $p_\theta(\mathbf{Z})$ , which is typically assumed to be a Gaussian distribution (Goodfellow et al., 2016).

$$\mathcal{L}_{\theta,\phi}(\mathbf{X}) = \mathbb{E}_{\mathbf{Z} \sim q_\phi(\mathbf{Z}|\mathbf{X})} [\log p_\theta(\mathbf{X}|\mathbf{Z})] - D_{KL}(q_\phi(\mathbf{Z}|\mathbf{X})||p_\theta(\mathbf{Z})) \quad (2)$$

where  $\mathbf{Z} \in \mathbb{R}^{W \times d_Z}$  is the sampled latent matrix and  $\mathbf{X} \in \mathbb{R}^{W \times d_X}$  is the input window.  $W$  refers to the window length, whereas  $d_X$  and  $d_Z$  refer to the input window and latent matrix dimensionality, respectively. Gradient-based optimisation minimises an objective function and the goal is the maximisation of the ELBO, hence the final loss function is defined as the negative of Equation 2, shown in Equation 3.

$$\mathcal{L}_{VAE} = -\mathcal{L}_{\theta,\phi}(\mathbf{X}) \quad (3)$$

Finally, to enable the backpropagation through the otherwise intractable gradient of the ELBO, the *reparametrisation trick* (Kingma and Welling, 2014) is applied, shown in Equation 4.

$$\mathbf{Z} = \mu_{\mathbf{Z}} + \varepsilon \cdot \sigma_{\mathbf{Z}} \quad (4)$$

where  $\varepsilon \sim \mathcal{N}(0, 1)$  and  $(\mu_{\mathbf{Z}}, \log \sigma_{\mathbf{Z}}^2) = q_\phi(\mathbf{X})$ .

### 2.4 Multi-Head Attention Mechanism

To simplify the explanation of MA as employed in this work, multi-head self-attention (MS) will be explained instead with the small difference between MA and MS being pointed out at the end.

MS consists of two different concepts: self-attention and its multi-head extension. Self-attention is nothing more than scaled dot-product attention (Vaswani et al., 2017) where the key, query and value are the same. The scaled dot-product attention score is the softmax (Bridle, 1990) of the product between

query matrix  $\mathbf{Q}$  and key matrix  $\mathbf{K}$  which is scaled by  $\sqrt{d_K}$ . The product between the attention score and the value matrix  $\mathbf{V}$  yields the context matrix  $\mathbf{C}$ , as shown in Equation 5.

$$\mathbf{C} = \text{Softmax} \left( \frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{d_K}} \right) \mathbf{V} \quad (5)$$

Compared to recurrent or convolutional layers, self-attention offers a variety of benefits, such as the reduction of computational complexity, as well as an increased amount of operations that can be parallelised. (Vaswani et al., 2017). Also, self-attention inherits an advantage over Bahdanau-style attention (Bahdanau et al., 2015) from the underlying scaled dot-product attention mechanism: it can run efficiently in matrix multiplication manner (Vaswani et al., 2017).

Multi-head self-attention then allows the attention model to attend to different representation subspaces (Vaswani et al., 2017), in addition to learning useful projections rather than it being a stateless transformation (Chollet, 2021). This is achieved using weight matrices  $\mathbf{W}_i^Q$ ,  $\mathbf{W}_i^K$ ,  $\mathbf{W}_i^V$ , which contain trainable parameters and are unique for each head  $i$ , as shown in Equation 6.

$$\mathbf{Q}_i = \mathbf{Q}\mathbf{W}_i^Q \quad \mathbf{K}_i = \mathbf{K}\mathbf{W}_i^K \quad \mathbf{V}_i = \mathbf{V}\mathbf{W}_i^V \quad (6)$$

Once the query, key and value matrices are linearly transformed via the weight matrices, the context matrix  $\mathbf{C}_i$  for each head  $i$  is computed using Equation 7.

$$\mathbf{C}_i = \text{Softmax} \left( \frac{\mathbf{Q}_i\mathbf{K}_i^T}{\sqrt{d_K}} \right) \mathbf{V}_i \quad (7)$$

Then, for  $h$  heads, the different context matrices are concatenated and linearly transformed again via the weight matrix  $\mathbf{W}^O$ , resulting in the multi-head context matrix  $\mathbf{C} \in \mathbb{R}^{W \times d_Z}$ , Equation 8.

$$\mathbf{C} = [\mathbf{C}_1, \dots, \mathbf{C}_h]\mathbf{W}^O \quad (8)$$

The underlying mechanism of MA is identical to MS, with the only difference being that  $\mathbf{K} = \mathbf{Q} \neq \mathbf{V}$ . The benefit of this alteration is discussed in Section 4.

## 3 RELATED WORK

MA-VAE belongs to the so-called generative model class, which encompasses both variational autoencoders, as well as generative adversarial networks. This section focuses solely on the work on VAE proposed in the context of time-series anomaly detection.

In time-series anomaly detection literature, the only other model that uses the combination of a VAE and an attention mechanism is by (Pereira and Silveira, 2018). For the purpose of our paper, it is named

VS-VAE. Their approach consists of a BiLSTM encoder and decoder, where, for an input window of length  $W$ , the  $t = W$  encoder hidden states of each direction are passed on to the variational self-attention (VS) mechanism (Bahuleyan et al., 2018). The resulting context vector is then concatenated with the sampled latent vector and then passed on to the decoder. The author claims that applying VS to the VAE model solves the bypass phenomenon, however, no evidence for this claim is provided.

The first published time-series anomaly detection approach based on VAE was LSTM-VAE (Park et al., 2018). One of the contributions is its use of a dynamic prior  $\mathcal{N}(\mu_p, 1)$ , rather than a static one  $\mathcal{N}(0, 1)$ . In addition to that, they introduce a state-based threshold estimation method consisting of a support-vector regressor (SVR), which maps the latent distribution parameters  $(\mu_z, \sigma_z)$  to the resulting anomaly score using the validation data. Hence, the dynamic threshold can be obtained through Equation 9.

$$\eta_t = \text{SVR}(\mu_{z,t}, \sigma_{z,t}) + c \quad (9)$$

where  $c$  is a pre-defined constant to control sensitivity.

OmniAnomaly (Su et al., 2019) attempts to create a temporal connection between latent distributions by applying a linear Gaussian state space model to them. For the purpose of this paper, it is called (OmniA). Also, it concatenates the last gated recurrent unit (GRU) hidden state with the latent vector sampled in the previous time step. In addition to that, it uses planar normalising flow (Rezende and Mohamed, 2015) by applying  $K$  transformations to the latent vector in order to approximate a non-Gaussian posterior, as shown in Equation 10.

$$f^k(\mathbf{z}_t^{k-1}) = \mathbf{u} \tanh(\mathbf{w} \mathbf{z}_t^{k-1}) + \mathbf{b} \quad (10)$$

where  $\mathbf{u}$ ,  $\mathbf{w}$  and  $\mathbf{b}$  are trainable parameters.

A simplified VAE architecture (Pereira and Silveira, 2019) based on BiLSTM layers is also proposed. For the purpose of our paper, it is called W-VAE. Unlike its predecessor (Pereira and Silveira, 2018), it drops the attention mechanism but provides contributions elsewhere. It offers two strategies to detect anomalies based on the VAE outputs. The first involves clustering the space characterised by the mean parameter of the latent distribution into two clusters and labelling the larger one as normal. This strategy has a few weaknesses: it cannot be used in an operative environment as it requires some sort of history of test windows to form the clusters and it assumes that there are always anomalous samples present. The second strategy finds the Wasserstein similarity measure (hence the  $W$  in the name) between the latent mean space mapping of the test window in question and the respective mapping  $i$  resulting from a representative

data subset, such as the validation subset. Equation 11 shows how the Wasserstein similarity measure is computed

$$W_i(\mathbf{z}_{\text{test}}, \mathbf{z}_i) = \|\mu_{\mathbf{z}_{\text{test}}} - \mu_{\mathbf{z}_i}\|_2^2 + \|\Sigma_{\mathbf{z}_{\text{test}}}^{1/2} - \Sigma_{\mathbf{z}_i}^{1/2}\|_F^2 \quad (11)$$

where the first term represents the  $L2$ -Norm between the mean distribution parameters resulting from the test window and each point of the representative subset. The second term represents the Frobenius norm between the covariance matrix resulting from the test window and each point of the representative subset.

SWCVAE (Chen et al., 2020) is the first that applies convolutional neural networks (CNN) to VAE for multivariate time-series anomaly detection. Particularly, 2D CNN layers are used with the justification of being able to process the input both spatially and temporally. We, however, doubt the ability of the model to properly detect anomalies through spatial processing, as a kernel moving along the feature axis can only capture features adjacent to each other. To create a continuous anomaly score from windows they append the last value of each window to the previous one. For the purpose of this paper, this process is referred to as *last-type reverse-windowing*.

SISVAE (Li et al., 2021) tries to improve the modelling robustness by the addition of a smoothing term in the loss function which contributes to the reduction of sudden changes in the reconstructed signal, making it less sensitive to noisy time steps.

As part of the VASP framework (von Schleinitz et al., 2021), a variational autoencoder architecture is proposed to increase the robustness of time-series prediction when faced with anomalies. While the main contribution is attributed to the framework itself, not the VAE, it should be noted that during inference only the mean parameter of the latent distribution is passed to the decoder.

## 4 PROPOSED APPROACH

### 4.1 Overview

To detect anomalies in multivariate time-series data, we propose a variational autoencoder architecture consisting of BiLSTM layers. The model architecture is illustrated in Figure 2. During training, the encoder  $q_\phi$  maps multivariate input window  $\mathbf{X}$  to a temporal distribution with parameters  $\mu_z$  and  $\log \sigma_z^2$  in the forward pass, Equation 12.

$$(\mu_z, \log \sigma_z^2) = q_\phi(\mathbf{X}) \quad (12)$$

Given the latent distribution parameters  $\mu_z$  and  $\log \sigma_z^2$ , the latent matrix is sampled from the result-

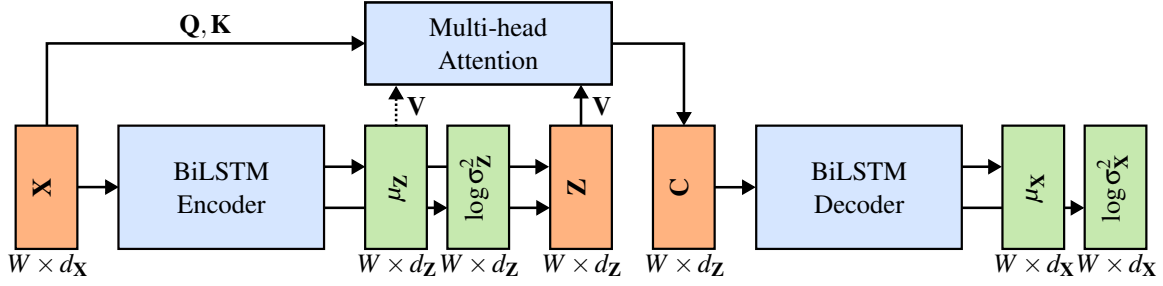


Figure 2: An illustration of the proposed MA-VAE model. Blue shapes designate trainable models, orange deterministic tensors and green distribution parameters. The shape of each tensor is designated below it. During training  $\mathbf{Z}$  is used as the value matrix, denoted by the solid arrow, whereas during inference  $\mu_{\mathbf{Z}}$  is used as the value matrix, denoted by the traced arrow.

ing distribution, as shown in Equation 13.

$$\mathbf{Z} \sim \mathcal{N}(\mu_{\mathbf{Z}}, \log \sigma_{\mathbf{Z}}^2) \quad (13)$$

Then, the input window  $\mathbf{X}$  is linearly transformed to obtain the query matrices  $\mathbf{Q}_i$  and key matrices  $\mathbf{K}_i$  for each head  $i$ . Likewise, the sampled latent matrix  $\mathbf{Z}$  is also transformed to the value matrix  $\mathbf{V}_i$ , as shown in Equation 14.

$$\mathbf{Q}_i = \mathbf{X}\mathbf{W}_i^Q \quad \mathbf{K}_i = \mathbf{X}\mathbf{W}_i^K \quad \mathbf{V}_i = \mathbf{Z}\mathbf{W}_i^V \quad (14)$$

To output the context matrix  $\mathbf{C}_i$  for each head  $i$ , the softmax of the through  $\sqrt{d_{\mathbf{K}}}$  normalised query and key product is multiplied with the value matrix, Equation 15.

$$\mathbf{C}_i = \text{Softmax} \left( \frac{\mathbf{Q}_i \mathbf{K}_i^T}{\sqrt{d_{\mathbf{K}}}} \right) \mathbf{V}_i \quad (15)$$

The final context matrix  $\mathbf{C}$  is the result of the linearly-transformed concatenation of each head-specific context matrix  $\mathbf{C}_i$ , as expressed in Equation 16.

$$\mathbf{C} = [\mathbf{C}_1, \dots, \mathbf{C}_h] \mathbf{W}^O \quad (16)$$

The decoder  $p_{\theta}$  then maps the context matrix  $\mathbf{C}$  to an output distribution with parameters  $\mu_{\mathbf{X}}$  and  $\log \sigma_{\mathbf{X}}^2$ , as shown in Equation 17.

$$(\mu_{\mathbf{X}}, \log \sigma_{\mathbf{X}}^2) = p_{\theta}(\mathbf{C}) \quad (17)$$

## 4.2 Inference Mode

Despite the generative capabilities of VAE, MA-VAE does not leverage generation for anomaly detection. Rather than sampling a latent matrix as shown in Equation 13 during inference, sampling is disabled and only  $\mu_{\mathbf{Z}}$  is taken as the input for the multi-head attention mechanism, like in (von Schleinitz et al., 2021). Equation 13 in the forward pass, therefore, is replaced by Equation 18.

$$\mathbf{Z} = \mu_{\mathbf{Z}} \quad (18)$$

This not only accelerates inference by eliminating the sampling process but is also empirically found to be a good approximation of an averaged latent matrix if it were sampled several times like in (Pereira and Silveira, 2018). The MA-VAE layout during inference is shown in Figure 2, where the traced arrow designates the information flow from the encoder to the MA mechanism.

## 4.3 Threshold Estimation Method

Anomalies are by definition very rare events, hence an ideal anomaly detector only flags measurements very rarely but accurately. Test bench engineers prefer an algorithm that only flags a sequence it is sure is an anomaly, in other words, an algorithm that outputs very few to no false positives. A high false positive count would lead to a lot of stoppages and therefore lost testing time and additional cost. Of course, the vast majority of measurements evaluated will be normal and hence it is paramount to classify them correctly, naturally leading to a high precision value. Also, there is no automatic evaluation methodology currently running at test benches, other than rudimentary rule-based methods, therefore a solution that plugs into the existing system that automatically detects *some* or *most* anomalies undetectable by rules is already a gain. To achieve this, the threshold  $\tau$  is set as the maximum log probability observed when the model is fed with validation data.

## 4.4 Bypass Phenomenon

VAE, when combined with an attention mechanism, can exhibit a behaviour called the bypass phenomenon (Bahuleyan et al., 2018). When the bypass phenomenon happens the latent path between encoder and decoder is ignored and information flow occurs mostly or exclusively through the attention mechanism, as it has deterministic access to the encoder hid-

den states and therefore avoids regularisation through the  $D_{KL}$  term. In an attempt to avoid this, (Bahuleyan et al., 2018) propose variational attention, which, like the VAE, maps the input to a distribution rather than a deterministic vector. Applied to natural language processing, (Bahuleyan et al., 2018) demonstrate that this leads to a diversified generated portfolio of sentences, indicating alleviation of the bypassing phenomenon. As previously mentioned, only (Pereira and Silveira, 2018) applies this insight in the anomaly detection domain, however, they do not present any proof that it alleviates the bypass phenomenon in their work. MA-VAE on the other hand, cannot suffer from the bypass phenomenon in the sense that information flow ignores the latent variational path between encoder and decoder since the MA mechanism requires the value matrix  $\mathbf{V}$  from the encoder to output the context matrix. Assuming the bypass phenomenon also applies to a case where information flow ignores the attention mechanism, one could claim that MA-VAE is not immune. To disprove this claim, the attention mechanism is removed from the model in an ablation study to see if anomaly detection performance remains the same. In this case,  $\mathbf{V}$  is instead directly input into the decoder. If it drops, it is evidence of the contribution of the attention mechanism to the model performance and hence is not bypassed. The results for this ablation study are shown and discussed in Section 5.

#### 4.5 Impact of Seed Choice

Given the stochastic nature of the VAE, the chosen seed can impact the anomaly detection performance as it can lead to a different local minimum during training. To investigate the impact the seed choice has on model training, MA-VAE is trained on three different seeds, the respective results are also shown in Section 5.

#### 4.6 Reverse-Window Process

Since the model is trained to reconstruct fixed-length windows, the same applies during inference. However, to decide whether a given measurement sequence  $\mathcal{S} \in \mathbb{R}^{T \times d_{\mathbf{X}}}$  is anomalous, a continuous reconstruction of the measurement is required. The easiest way to do so would be to window the input measurement using a shift of 1, input the windows into the model and chain the last time step from each output window to obtain a continuous sequence (Chen et al., 2020). Considering the BiLSTM nature of the encoder and decoder, the first and last time steps of a window can only be computed given the states from one direction, making these values, in theory, less ac-

Algorithm 1: Anomaly Detection Process.

---

**Input:** Sequence  $\mathcal{S} \in \mathbb{R}^{T \times d_{\mathbf{X}}}$ , Threshold  $\tau$   
**Result:** Label  $l$

```

 $n_{\text{windows}} \leftarrow T - W + 1$ 
 $\mu_{\mathbf{X}, \text{temp}} \leftarrow \text{zeros}(n_{\text{windows}}, T, d_{\mathbf{X}}) + \text{NaN}$ 
 $\sigma_{\mathbf{X}, \text{temp}}^2 \leftarrow \text{zeros}(n_{\text{windows}}, T, d_{\mathbf{X}}) + \text{NaN}$ 
for  $i = 1 \rightarrow n_{\text{windows}}$  do
     $\mathbf{X} \leftarrow \mathcal{S}[i : W + i]$ 
     $(\mu_{\mathbf{Z}}, \log \sigma_{\mathbf{Z}}^2) \leftarrow q_{\phi}(\mathbf{X})$ 
     $\mathbf{C} \leftarrow \text{MA}(\mathbf{X}, \mathbf{X}, \mu_{\mathbf{Z}})$ 
     $(\mu_{\mathbf{X}}, \log \sigma_{\mathbf{X}}^2) \leftarrow p_{\theta}(\mathbf{C})$ 
     $\mu_{\mathbf{X}, \text{temp}}[i, i : i + W] \leftarrow \mu_{\mathbf{X}}$ 
     $\sigma_{\mathbf{X}, \text{temp}}^2[i, i : i + W] \leftarrow \sigma_{\mathbf{X}}^2$ 
end for
 $\mu_{\mathbf{X}, \text{seq}} \leftarrow \text{nanmean}(\mu_{\mathbf{X}, \text{temp}})$ 
 $\sigma_{\mathbf{X}, \text{seq}}^2 \leftarrow \text{nanmean}(\sigma_{\mathbf{X}, \text{temp}}^2)$ 
 $s \leftarrow -\log p(\mathbf{X} | \mu_{\mathbf{X}, \text{seq}}, \sigma_{\mathbf{X}, \text{seq}})$ 
 $l \leftarrow \max(s) > \tau$ 

```

---

curate, however. To overcome this, we propose averaging matching time steps in overlapping windows, which is called *mean-type* reverse-window method. This is done by pre-allocating an array with NaN values, filling it, and taking the mean for each time step while ignoring the NaN values. This process and the general anomaly detection process are described in Algorithm 1. This reverse-window process is done for the mean and *variance* parameters of the output distribution, then the variance is converted to *standard deviation* since two distributions cannot be combined by averaging the standard deviations. With a continuous mean and standard deviation, the continuous negative log probability, i.e. the anomaly score  $s$ , is computed for the respective measurement. A comparison between the mean, last and first reverse-window process is provided in Section 5.

## 5 RESULTS

### 5.1 Setup

The encoder and decoder both consist of two BiLSTM layers, with the outer ones having 512 hidden- and cell-state sizes and the inner ones 256. All other parameters are left as the default in the TensorFlow API.

During training only, input windows are corrupted using Gaussian noise using 0.01 standard deviation to increase robustness to noise.

Key factors that are investigated in Section 5 are given a default value which applies to all experiments unless otherwise specified. These factors are train-

ing/validation subset size, which is set to 512h, seed choice, which has been kept at 1, reverse-window method, where the mean-type is used, the latent dimension size, which is set to  $d_Z = 16$  and the MA mechanism, which is set up as proposed in (Vaswani et al., 2017) with a head count of  $h = 8$  and a key dimension size  $d_K = \lfloor d_X/h \rfloor = 1$ .

The optimiser used is the AMSGrad optimiser with the default parameters in the TensorFlow API.

Cyclical  $D_{KL}$  annealing (Fu et al., 2019) is applied to the training of MA-VAE, to avoid the  $D_{KL}$  vanishing problem. The  $D_{KL}$  vanishing problem occurs when regularisation is too strong at the beginning of training, i.e. the Kullback-Leibler divergence term has a larger magnitude in relation to the reconstruction term. Cyclical  $D_{KL}$  annealing allows the model to weigh the Kullback-Leibler divergence lower than the reconstruction term in a cyclical manner through a weight  $\beta$ . This callback is configured with a grace period of 25 epochs, where  $\beta$  is linearly increased from 0 to  $10^{-8}$ . After the grace period,  $\beta$  is set to  $10^{-8}$  and is gradually increased linearly to  $10^{-2}$  throughout the following 25 epochs, representing one loss cycle. This loss cycle is repeated until the training stops.

All priors in this work are set as standard Gaussian distributions, i.e.  $p = \mathcal{N}(0, 1)$ .

To prevent overfitting, early stopping is implemented. It works by monitoring the log probability component of the validation loss during training and stopping if it does not improve for 250 epochs. Logically, the model weights at the lowest log probability validation loss are saved.

Training is done on a workstation configured with an NVIDIA RTX A6000 GPU. The library used for model training is TensorFlow 2.10.1 on Python 3.10 on Windows 10 Enterprise LTSC version 21H2.

The results provided are given in the form of the calibrated and uncalibrated anomaly detection performance, i.e. with and without consideration of threshold  $\tau$ , respectively. Recall that the threshold used is the absolute maximum negative log probability obtained from the validation set. Calibrated metrics are the precision, recall and  $F_1$  score. Precision  $P$  represents the ratio between correctly identified anomalies (true positives) and all positives (true and false), shown in Equation 19, recall  $R$  represents the ratio between true positives and all anomalies, shown in Equation 19, and  $F_1$  score represents the harmonic mean of the precision and recall, shown in Equation 20. The underlying metrics used to calculate all of the below are the true positives ( $TP$ ), false negatives ( $FN$ ) and false positives ( $FP$ ).

$$P = \frac{TP}{TP + FP} \quad R = \frac{TP}{TP + FN} \quad (19)$$

$$F_1 = \frac{TP}{TP + 0.5 * (FP + FN)} = 2 * \frac{P * R}{P + R} \quad (20)$$

The theoretical maximum  $F_1$  score,  $F_{1,best}$ , is also provided to aid discussion. This represents the best possible score achievable by the approach if the ideal threshold were known, i.e. the point on the precision-recall curve that comes closest to the  $P = R = 1$  point, though, in reality, this value is not observable and hence *cannot* be obtained in an unsupervised manner.

The uncalibrated anomaly detection performance, i.e. the performance for a range of thresholds, each 0.1 apart, is represented by the area under the continuous precision-recall curve  $A_{PRC}^{cont}$ , Equation 21.

$$A_{PRC}^{cont} = \int_0^1 P dR \quad (21)$$

As the integral cannot be computed for the continuous function, the area under the discrete precision-recall curve  $A_{PRC}^{disc}$  is used which is done using the trapezoidal rule, Equation 22.

$$A_{PRC}^{disc} = \sum_{k=1}^N \frac{f(R_{k-1}) + f(R_k)}{2} \Delta R_k \quad (22)$$

where  $N$  is the number of discrete sub-intervals,  $k$  the index of sub-intervals and  $\Delta R_k$  the sub-intervals length at index  $k$ . Precision is a function of recall, i.e.  $P = f(R)$ .

## 5.2 Ablation Study

MA-VAE is tested without the MA mechanism and with a direct connection from the encoder to the decoder to observe whether it impacts results.

The anomaly detection performance of MA-VAE and its counterpart without MA, henceforth referred to as *No MA* model, are shown in Table 2. While the precision value of the No MA model is slightly higher than the MA-VAE, the recall value on the other hand is much lower. Overall, MA-VAE has a higher  $F_1$  score, as well as a higher theoretical maximum  $F_1$  score, although both values are so close enough to each other that one could claim the threshold is near ideal. The uncalibrated performance is also higher in the case of the MA-VAE, as evident in the precision-recall plot in Figure 3. Interestingly, MA-VAE may feature a lower precision value for the chosen unsupervised threshold but has the potential to have a higher maximum recall at  $P = 1$ .

The results hence point towards an improvement brought about by the addition of the MA mechanism and therefore the bypass phenomenon can be ruled out.



Table 2: Precision  $P$ , recall  $R$ ,  $F_1$  score, theoretical best  $F_1$  score  $F_{1,\text{best}}$  and area under the precision-recall curve  $A_{\text{PRC}}$  results for the model variant without the MA mechanism and MA-VAE. The best values for each metric are given in **bold**.

Model	$P$	$R$	$F_1$	$F_{1,\text{best}}$	$A_{\text{PRC}}$
No MA	<b>1.00</b>	0.35	0.52	0.54	0.52
MA-VAE	0.92	<b>0.55</b>	<b>0.69</b>	<b>0.70</b>	<b>0.66</b>




Figure 3: Precision-recall curves for the model variation without MA and MA-VAE.

### 5.3 Data Set Size Requirements

To evaluate how much data is required to train MA-VAE to a point of adequate anomaly detection performance, it has been trained with 1h, 8h, 64h, and 512h worth of dynamic testing data.

The results for this experiment are presented in Table 3. On the one hand, as the training/validation subset increases in size, the precision value improves, with the largest jump occurring when the dynamic testing time goes from 1h to 8h. The recall value on the other hand decreases as the subset grows. This can be attributed to the fact that smaller subset sizes lead to a small validation set and therefore less data to obtain a threshold from. With a limited amount of data to obtain a threshold from, it is more difficult to get a representative error distribution, leading to a threshold that is very small and hence marks most anomalies correctly but also leads to a lot of false positives.  $F_1$  score reaches a point of diminishing returns with the 8h subset onwards, this can also be observed in the case of the theoretical maximum  $F_1$  score,  $F_{1,\text{best}}$ , as well as in the  $A_{\text{PRC}}$  value, further supported by the precision-recall plot in Figure 4. Lastly, the  $F_1$  score seems to approach the  $F_{1,\text{best}}$  score as the subset grows, also backing the fact that with a small subset

size, a good threshold cannot easily be obtained.

Table 3: Precision  $P$ , recall  $R$ ,  $F_1$  score, theoretical best  $F_1$  score  $F_{1,\text{best}}$  and area under the precision-recall curve  $A_{\text{PRC}}$  results for the different training/validation subset sizes. The best values for each metric are given in **bold**.

Size	$P$	$R$	$F_1$	$F_{1,\text{best}}$	$A_{\text{PRC}}$
1h	0.09	<b>0.88</b>	0.17	0.55	0.49
8h	0.66	0.63	0.64	<b>0.72</b>	<b>0.69</b>
64h	0.71	0.57	0.63	0.69	0.68
512h	<b>0.92</b>	0.55	<b>0.69</b>	0.70	0.66

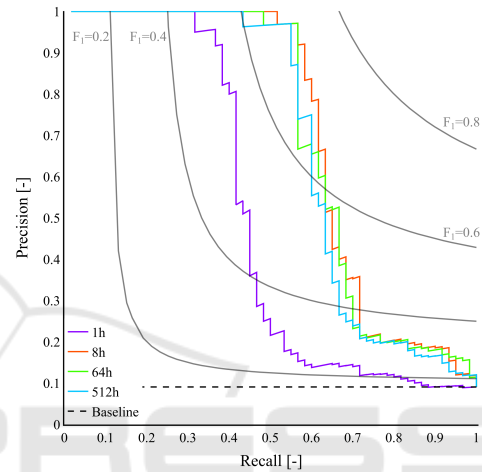


Figure 4: Precision-recall curves for the model trained on different training/validation subset sizes.

Therefore, for application at the test bench, the largest subset size is desirable due to the higher precision value and a *closer-to-ideal* threshold value.

### 5.4 Impact of Seed Choice

To illustrate the impact it has on the performance metrics, they are presented using three different seeds.

Table 4 shows that while the precision values are roughly in the same range for all seeds, the recall values vary more significantly, which also reflects on the  $F_1$  score. However, by inspecting the  $F_{1,\text{best}}$  and  $A_{\text{PRC}}$  values it becomes clear that the seeds are not as far apart as the recall value suggests and that the issue may lie with the threshold choice. Figure 5 further supports this, as all lines have roughly the same path, with the exception of seed 3 at very high precision values. The plot clearly shows that a more suitable (lower) threshold would lead to seed 3 having a comparable recall value to the other seeds while maintaining high precision.

Some differences can be observed between the seeds, especially in the recall values, however, this

can be attributed to the unsupervised threshold choice.

Table 4: Precision  $P$ , recall  $R$ ,  $F_1$  score, theoretical best  $F_1$  score  $F_{1,best}$  and area under the precision-recall curve  $A_{PRC}$  results for the different seeds. The best values for each metric are given in **bold**.

Seed	$P$	$R$	$F_1$	$F_{1,best}$	$A_{PRC}$
1	0.92	0.55	0.69	0.70	0.66
2	0.90	<b>0.60</b>	<b>0.72</b>	<b>0.73</b>	<b>0.67</b>
3	<b>0.96</b>	0.40	0.56	0.70	0.64

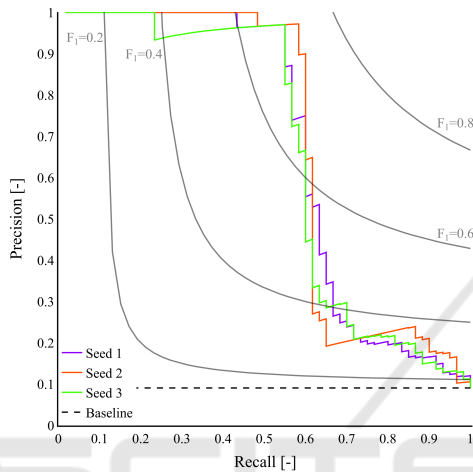


Figure 5: Precision-recall curves for the model trained on different seeds.

### 5.5 Reverse-Window Process

To investigate the effect of the mean-type reverse-window method, it is compared with the first-type and last-type methods where the first and last values of each window are carried over, respectively.

The results in this subsection, Table 5 and Figure 6, tell a similar story to the previous subsection. The metrics independent of the chosen threshold are very similar regardless of the reverse-window method, implying that they are comparable and that any differences in the calibrated metrics can be attributed to the chosen threshold. The mean-type reverse-window method results in a higher computational load, though negligible. For a rather long sequence of 4000 time steps, i.e. around 33 minutes long, the mean-type method only takes around 2 seconds longer. One source of delay that can appear, however, is during *online* anomaly detection. An online anomaly detection algorithm is defined as an algorithm which evaluates the sequence as it is being recorded. To obtain time step  $t$  using the mean-type (or the first-type) you have to wait for time step  $t + W$  while  $t < W$ . This

Table 5: Precision  $P$ , recall  $R$ ,  $F_1$  score, theoretical best  $F_1$  score  $F_{1,best}$  and area under the precision-recall curve  $A_{PRC}$  results for the different reverse-window types. The best values for each metric are given in **bold**.

Type	$P$	$R$	$F_1$	$F_{1,best}$	$A_{PRC}$
first	<b>0.97</b>	0.48	0.64	0.69	0.64
last	0.88	<b>0.58</b>	<b>0.71</b>	<b>0.71</b>	<b>0.67</b>
mean	0.92	0.55	0.69	0.70	0.66




Figure 6: Precision-recall curves for different reverse-window methods.

translates to a delay of around 2 minutes in the real world, given the chosen window size. If the evaluation is done offline, i.e. when  $t = W$ , then this delay is eliminated since the last value does not have other overlapping values to compute the mean.

### 5.6 Hyperparameter Optimisation

As part of the hyperparameter optimisation of MA-VAE, a list of latent dimension sizes  $d_Z$  in combination with a list of key dimension sizes  $d_K$  is tested. Despite the larger learning capacity associated with a higher  $d_K$ , the concatenation is always transformed to a matrix of size  $d_O = d_Z$ . For the two variables, values of 1, 4, 16, and 64 are tested.

The best result is achieved with  $d_Z = d_K = 64$ . Given that they are the respective highest values of  $d_Z$  and  $d_K$ , even higher values should be experimented with in the future, though they will lead to higher model complexity and training/inference time. The attention head count  $h$  was also experimented with using the same range of values as for  $d_Z$  and  $d_K$ , however, none performed better than the  $h = 8$  configuration. The results are presented in Table 6, the corresponding precision-recall plot is shown in Figure 7. 91% of the sequences marked as normal were actually normal and 67% of the total number of anomalous sequences in the test set were detected. One example of

Table 6: Precision  $P$ , recall  $R$ ,  $F_1$  score, theoretical best  $F_1$  score  $F_{1,best}$  and area under the precision-recall curve  $A_{PRC}$  result for the best  $d_Z$ ,  $d_K$  and  $h$  values.

$d_Z$	$d_K$	$h$	$P$	$R$	$F_1$	$F_{1,best}$	$A_{PRC}$
64	64	8	0.91	0.67	0.77	0.79	0.74




Figure 7: Precision-recall curve for the final MA-VAE.

the anomalous cycles and the respective reconstructions is plotted in Figure 8.

## 5.7 Benchmarking

Of course, MA-VAE is not the first model proposed for time-series anomaly detection. To underline its anomaly detection performance, it is compared with a series of other models based on variational autoencoders. The chosen subset of models is based on the work discussed in Section 3 which either linked source code or contained enough information for implementation. The models are implemented using hyperparameters specified in their respective publications. To even the playing field, the models are trained on the 512h subset with early stopping, which is parametrised equally across all models. The anomaly detection process specified in Algorithm 1 is also applied to all models, along with the threshold estimation method. The results can be seen in Table 7.

Table 7: Precision  $P$ , recall  $R$ ,  $F_1$  score, theoretical best  $F_1$  score  $F_{1,best}$  and area under the precision-recall curve  $A_{PRC}$  results for competing models and MA-VAE (Ours). The best values for each metric are given in **bold**.

Model	$P$	$R$	$F_1$	$F_{1,best}$	$A_{PRC}$
VS-VAE	<b>1.00</b>	0.33	0.50	0.56	0.51
W-VAE	<b>1.00</b>	0.30	0.46	0.46	0.41
OmniA	0.96	0.37	0.53	0.58	0.53
SISVAE	<b>1.00</b>	0.30	0.46	0.50	0.51
MA-VAE	0.91	<b>0.67</b>	<b>0.77</b>	<b>0.79</b>	<b>0.74</b>




Figure 8: Wheel diameter anomaly plotted in black and the output distribution in red, as well as anomaly score plotted in blue and the threshold as a straight line in orange.

As is evident, MA-VAE outperforms all other models in every metric, except for precision. As stated in Section 4 a high precision figure is important in this type of powertrain testing, however, the reduced precision is still considered tolerable. Also, it comes at the benefit of a much higher recall figure, which is reflected in the superior  $F_1$  figure. Furthermore, the  $F_{1,best}$  figure, which is obtained at  $P = 0.98$  and  $R = 0.67$ , suggests that MA-VAE has the potential to achieve even higher precision without sacrificing recall if the threshold were optimised. The higher  $A_{PRC}$  also shows that MA-VAE has a higher range of thresholds at which it performs well.

## 6 CONCLUSION AND OUTLOOK

In this paper, a multi-head attention variational autoencoder (MA-VAE) for anomaly detection in automotive testing is proposed. It not only features an attention configuration that avoids the bypass phenomenon but also introduces a novel method of remapping windows to whole sequences. A number of experiments are conducted to demonstrate the anomaly detection performance of the model, as well as to underline the benefits of key aspects introduced with the model.

From the results obtained, MA-VAE clearly benefits from the MA mechanism, indicating the avoidance of the bypass phenomenon. Moreover, the proposed approach only requires a small training/validation subset size but fails to obtain a suitable threshold, as with increasing subset size only the calibrated anomaly detection performance increases. Training with different seeds also is shown to have little impact on the anomaly detection metrics, provided the threshold is chosen suitably, further underlining the previous point. Moreover, mean-type reverse windowing fails to significantly outperform its first-type and last-type counterparts, while introducing additional lag if it is applied to online anomaly detection. Lastly, the hyperparameter optimisation revealed that the MA-VAE variant with the largest latent dimension and attention key dimension resulted in the best anomaly detection performance. It is only 9% of the time wrong when an anomaly is flagged and manages to discover 67% of the anomalies present in the test data set. Also, it outperforms all other competing models it is compared with.

In the future, a method of threshold choice involving active learning will be investigated, which can use user feedback to hone in on a better threshold. Also, MA-VAE is set to be tested in the context of online anomaly detection, i.e. during the driving cycle measurement.

## REFERENCES

- Bahdanau, D., Cho, K., and Bengio, Y. (2015). Neural Machine Translation by Jointly Learning to Align and Translate. In *International Conference on Learning Representations (ICLR)*.
- Bahuleyan, H., Mou, L., Vechtomova, O., and Poupard, P. (2018). Variational Attention for Sequence-to-Sequence Models. In *International Conference on Computational Linguistics (COLING)*.
- Bridle, J. S. (1990). Probabilistic Interpretation of Feedforward Classification Network Outputs, with Relationships to Statistical Pattern Recognition. *Neurocomputing*, pages 227–236.
- Chen, T., Liu, X., Xia, B., Wang, W., and Lai, Y. (2020). Unsupervised Anomaly Detection of Industrial Robots Using Sliding-Window Convolutional Variational Autoencoder. *IEEE Access*, 8:47072–47081.
- Chollet, F. (2021). *Deep Learning with Python*. Manning Publications.
- Fu, H., Li, C., Liu, X., Gao, J., Celikyilmaz, A., and Carin, L. (2019). Cyclical Annealing Schedule: A Simple Approach to Mitigating. In *Conference of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*.
- Goodfellow, I., Bengio, Y., and Courville, A. (2016). *Deep Learning*. MIT Press.
- Kingma, D. P. and Welling, M. (2014). Auto-Encoding Variational Bayes. In *International Conference on Learning Representations (ICLR)*.
- Li, L., Yan, J., Wang, H., and Jin, Y. (2021). Anomaly Detection of Time Series With Smoothness-Inducing Sequential Variational Auto-Encoder. *Transactions on Neural Networks and Learning Systems*, 32(3):1177–1191.
- Park, D., Hoshi, Y., and Kemp, C. C. (2018). A Multimodal Anomaly Detector for Robot-Assisted Feeding Using an LSTM-Based Variational Autoencoder. *IEEE Robotics and Automation Letters*, 3(3):1544–1551.
- Pereira, J. and Silveira, M. (2018). Unsupervised Anomaly Detection in Energy Time Series Data Using Variational Recurrent Autoencoders with Attention. In *International Conference on Machine Learning and Applications (ICMLA)*.
- Pereira, J. and Silveira, M. (2019). Unsupervised representation learning and anomaly detection in ECG sequences. *International Journal of Data Mining and Bioinformatics*, 22(4):389.
- Rezende, D. J. and Mohamed, S. (2015). Variational Inference with Normalizing Flows. In *International Conference on Machine Learning (ICML)*.
- Rezende, D. J., Mohamed, S., and Wierstra, D. (2014). Stochastic Backpropagation and Approximate Inference in Deep Generative Models. In *International Conference on Machine Learning (ICML)*.
- Shannon, C. (1949). Communication in the Presence of Noise. *Proceedings of the IRE*, 37(1):10–21.
- Su, Y., Zhao, Y., Niu, C., Liu, R., Sun, W., and Pei, D. (2019). Robust Anomaly Detection for Multivariate Time Series through Stochastic Recurrent Neural Network. In *International Conference on Knowledge Discovery & Data Mining (KDD)*.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., and Polosukhin, I. (2017). Attention Is All You Need. In *Conference on Neural Information Processing Systems (NIPS)*.
- von Schleinitz, J., Graf, M., Trutschnig, W., and Schröder, A. (2021). VASP: An autoencoder-based approach for multivariate anomaly detection and robust time series prediction with application in motorsport. *Engineering Applications of Artificial Intelligence*, 104:104354.