Using a self-supervised encoder for anticipating failures in industrial equipment

Anonymous Author(s)
Affiliation
Address
email

Abstract

Failures in industrial systems entail safety risks and production losses. However, with the right choice of technique, they can be forecasted. This is possible because certain measurable variables that are present in the system, respond differently to operation conditions when a failure is nearing. The objective of this work is to calculate the expected response that such variables would have in a healthy system. It will then be possible to compare them with the real time measures, in order to diagnose if there is a problem with the system. If there is a divergence between the two of them, a fault is signaled. The type of fault diagnosed depends on the shape of the divergence, following a set of rules proposed by human experts. It has to be considered that, for the same operation conditions, each system has a different response that is induced by disparities in manufacturing, installation procedures and other factors that can't be directly quantified. In addition, even for a single system, its response varies over time due to its *degradation*, which can't usually be measured by using non-invasive techniques.

The expected response is often calculated by means of parametric models. However, three main problems arise with them. First, they are based on suppositions that don't account for all the nuances present in a production environment, which results in a loss of accuracy. Second, a different model must be fitted for every system in order to account for their particularities. Third, the degradation term must be estimated independently and given to them. These last two shortcoming complicate the deployment of such models. In order to overcome these limitations, we propose a neural network model based on an encoder-decoder architecture.

The nature of industrial data makes this approach possible. Industrial systems are usually stopped and started multiple times during their lifetime. It is called *warm up* to the period at the beginning of each of these cycles when the system goes from idle to its intended operating conditions. The rest of the cycle is known as *operation* and it is when we want to detect possible failures.

Our hypothesis is that the information about the degradation and particularities of each system for the *operation* phase of each cycle, can be inferred from the *warm up* phase. Therefore, we force the encoder to reduce the *warm up* time-sequence to a single point, which is fed to the decoder. Then, for every observation of the *operation* phase, the decoder takes as input the operation conditions and the corresponding encoded point for that cycle and outputs the expected response of the system at that time. The network is trained in TensorFlow 2.0 with a custom distributed loop and using tf.data.Dataset for data preprocessing. The chosen objective function is the MSE with respect to the measured response. Adam is used as the optimizer, with a linear learning rate scheduler after the first epochs. As no labelled data is given, we call it a self-supervised encoder.

As a side gain, the model is also interpretable. Every point represents a single system and cycle and can be plotted in a 2D graph. This is useful because it allows us to perceive similarities between different systems and their evolution in the graph along the cycles. This is represented in Figure 1.

Submitted to 33rd Conference on Neural Information Processing Systems (NeurIPS 2019). Do not distribute.

In order to validate our hypothesis, we train the model with data coming from electrolytical cells. In this particular case, the cell's output voltage is the variable that changes its behavior when nearing a failure. Similar processes are common in industry and we are confident that the same procedure should apply to them as well. Results obtained to date prove that our hypothesis is correct. A single trained model works for many different systems and, without needing external degradation data, improves the accuracy of currently used approaches. More detailed results with comparisons to current methods will be presented, as tests are still being carried out at the moment. A comparison between the measured and the predicted voltage can be found in Figure 2.

Figures

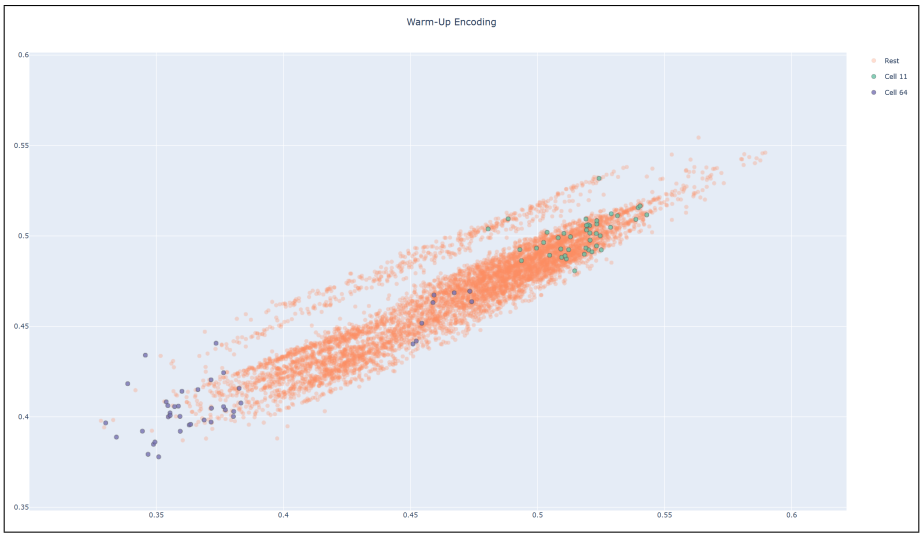

Figure 1: This figure represents all the encoded *warm-up* phases of each system. Two eletrolytical cells are highlighted over the rest. The *green* evolves over time but remains in the same cluster, which means that its degradation over time is stable. However, the *purple* is divided in two clusters. This means that something occured to it at a certain cycle and its degradation changed significantly.

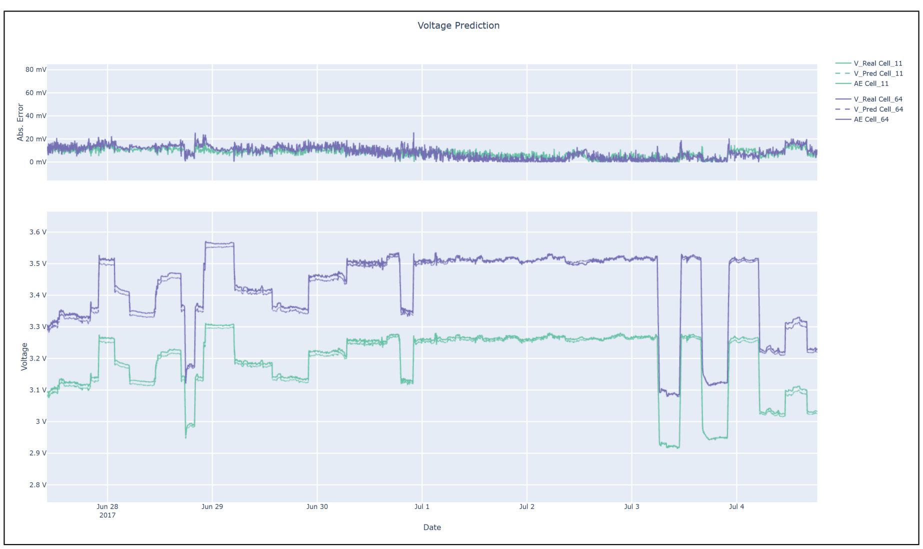

Figure 2: An extract from the *operation* phase of a cycle is represented. The lower graph shows the divergence between the predicted (dashed line) and the measured voltage (solid line) for each of the two cells presented in Figure 1. The upper graph shows the *Absolute Error* between both values. We can see that it stays well under the defined failure threshold of 50mV, which is a satisfactory result.