

## Data Science Case Study

# Predictive Maintenance Modeling for Industrial Machinery

*Mosaic was contracted by multi-national manufacturer of construction and mining equipment to develop a proof-of-concept predictive maintenance model to predict equipment failure before it happens to minimize downtime and optimize maintenance schedules.*



### Industry

Manufacturing



### Use Case

Predictive Maintenance,  
IoT



### Techniques

Deep Learning (LSTM) &  
Classification



### Outcome

Proved that with the right labeled data, PdM was possible using Deep Learning.

# INTRO

Every machine degrades over time. This degradation is a natural part of equipment's lifecycle – wear and tear through normal use cases any machine to operate less efficiently, and by performing regular maintenance, we can minimize the severity of the degradation and keep our machines working at maximum efficiency. In principle, there are two fundamental maintenance issues, especially for manufacturing companies where machine up-time is mission critical: planned and unplanned. Unplanned maintenance can have a significant financial impact on organizations due to the upkeep itself, factory-line disruption, and having a fixed asset unutilized for some time. Regularly scheduled planned maintenance can reduce the frequency of unplanned maintenance and ensure that everything runs smoothly. However, even with a regular maintenance schedule, sometimes machines fail before we expect them to. With new IoT technologies, there is an abundance of sensor equipment constantly monitoring the performance of machinery – and it is with these sensors that a powerful AI tool, predictive maintenance, can begin to shine.

Any company which operates or sells heavy machinery like mining equipment can benefit greatly by having some way to predict when these machines will fail. Predictive maintenance at its core is a process for these companies to not only minimize downtime but to improve their efficiency and productivity. Deep learning can be a powerful technique for predictive maintenance due to its innate ability to identify complex patterns in scenarios involving large, complex datasets containing multiple types of data. Ultimately, applying deep learning to predictive maintenance can help improve machine or service degradation.

Mosaic has a long history of successfully designing and deploying custom sensor-based solutions powered by deep learning. Our models continuously monitor IoT sensor data of operational behavior to ensure smooth operation and prevent unexpected disruption. The ability to better foresee the future and to anticipate and plan for future events and disruptions helps companies smoothly continue their operations in the face of disruptive events such as maintenance failures.

# THE PROBLEM

Mosaic was contracted by one of the world's largest machine manufacturers to develop a proof-of-concept predictive maintenance model to predict engine failure in large mining equipment. The customer needed to understand if measurements of engine pressure and temperature at various points in the engine and exhaust system could distinguish between different types of leaks and cooling system failures – and accurately predict how much time was left before the leak or part would need to be repaired.

Many large industrial engines have two semi-independent sides to the engine. While these sides may have temperature and pressure fluctuations throughout the day, each side should remain in sync with the other. When the engine is not operating as it should before or during a failure, however, the temperature and pressure difference between these two sides of the engine begins to fluctuate and drift, a sign in the data that something is going wrong.



Mosaic was challenged with using this signal and other sensor data to understand if the data science team could predict if an engine was going to fail so that maintenance could be scheduled and denote the kind of maintenance that should be scheduled. Mosaic designed and deployed custom deep learning models to test this concept.

Once deployed, the models would continuously monitor engine behavior to ensure smooth operation and predict engine failure in advance to prevent unexpected disruption.

## DETECTING THE SIGNAL

When building and training the predictive maintenance deep learning model, Mosaic used in part pressure and temperature differentials between specific components of the engine which may indicate an imminent engine failure, be it a catastrophic failure or a slow insidious leak. To help inform the severity of the imminent failure – and what kind of maintenance would be required – the data science team trained a model that could use these signals to distinguish between different failure modes, namely different leak types and part failures. Some examples of these failures are shown in Figure 1.




Figure 1 Explained: The above image displays some examples of the apparent signal that arises during a failure event. In the left figure, we display the changes of a pressure differential typical of two different kinds of failures. In the right figure, we show the behavior of a third kind of failure, and its effect on a measured temperature differential.

Figure 1 displays the three types of failure that the data science team were asked to distinguish and predict in advance. Noteworthy in the above figures is that each failure, to the human observer, is not subtle in the chosen examples. There is an apparent change in the data when the engine is failing. In these examples, with appropriate subject matter expertise labeling these events, one might expect an LSTM, a deep neural net architecture designed to process sequential data, or similar models to perform well in predicting a failure and distinguishing between the failure type classes.

## TRAINING AND TUNING THE DEEP LEARNING ALGORITHMS

Integral to developing a proof-of-concept model is understanding if this approach has merit – can a data scientist use a model to predict only if any failure is imminent? The confusion matrix in Figure 2 shows the performance of an LSTM at predicting on a per-hour basis if the sensors in the engine are recording data that indicate the machine needs maintenance.

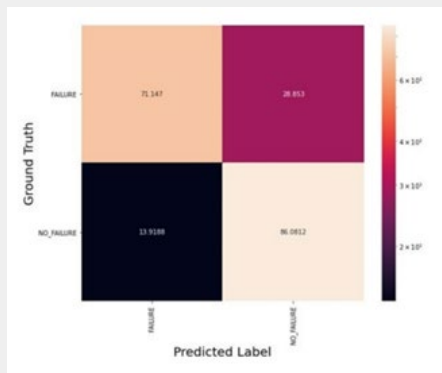


Figure 2 explained: This figure is a confusion matrix for the LSTMs performance on an hour-to-hour basis if an engine failure is imminent. Here one sees that we correctly predict no imminent failure 86% of the time, with a false-negative rate of 13%. Similarly, when a failure is imminent, the model correctly predicts this 71% of the time with a false negative rate of about 29%.

Figure 2 shows the confusion matrix for a binary model trained to predict if any of the three failures are upcoming on an hour-to-hour basis. One red flag with results from this model is the false negative rate predicting if a failure is imminent. A high false positive rate on the simplified binary problem would pose future complications considering the multi-class prediction problem. However, at this point in our analysis the Mosaic team were optimistic – the model was performing better than random guessing, which would not be the case if this data didn't offer any predictive insights, and redefining what it means to predict a failure (i.e., daily aggregates vs. hour-to-hour predictions) could reduce false negatives and positives.

Mosaic attempted to understand how the model would perform on the multi-class problem that not only predicts failure but also distinguishes between the different modes of failure. Unfortunately, the multiclass models struggled to distinguish between different failure types. While the data science team were still predicting long-term events reasonably well, such as the absence of failure and part failures (with a recall of approximately 80% and 70% for each case respectively on an hour-to-hour basis), distinguishing the leaks from each other and other kinds of failures was near impossible for this particular model in many cases.

This led Mosaic's team of data scientists to ask two important questions. First, is an hour-to-hour comparison for accuracy incorrect? Perhaps some of these failures happen so slowly and subtly that the team are first misclassifying them, and as the problem progresses, the team eventually get it right. Or perhaps the team start with correct predictions, but for longer lived failures our model eventually gets confused and misclassifies this event.

The second, and more problematic concern, was that the training and testing labels might be inconsistent, meaning our model is struggling to learn due to a sometimes erroneous source of truth. The modeling teased out data and labeling problems that the manufacturer did not know were problems, a critical goal of a proof-of-concept machine learning effort.

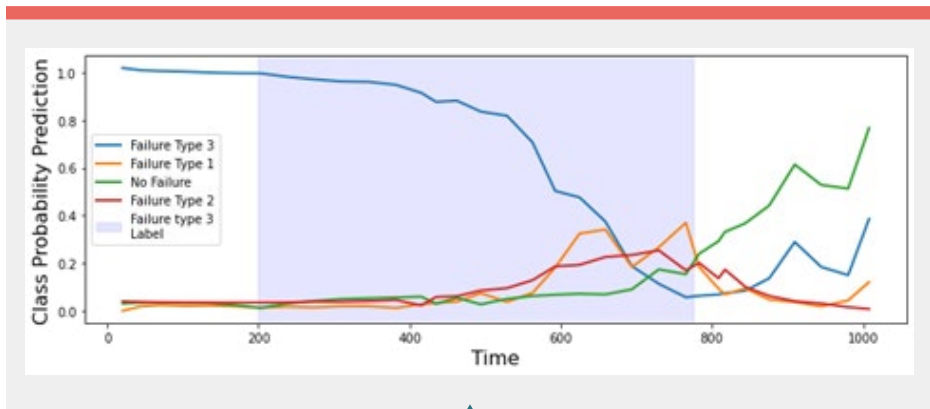


Figure 3: Mosaic shows predictions of the failure class for a specific part.

## LEARNING FROM THE PROOF-OF-CONCEPT

From Figure 3, the data science team sees that within the highlighted region of the labeled failure, the model first predicts this correctly before eventually losing confidence in the classified region. However, this misclassification begins when the failure starts to worsen – the system is changing, and the model loses confidence in what kind of failure we should expect.

The more significant issue is more subtle. The labeled failure begins at approximately 2019-03, but if one looks at the second figure in Figure 1 (the source of our data), one notices that it appears this failure started well before the label says it did – if the data is indistinguishable before and after our label, is it fair to expect the model to get this right? Similarly, in Figure 3, there are several examples of these data exhibiting wild variations with similar characteristics to what is frequently seen during failures – but the model is being told that there is no failure happening in those regions.

Based on these observations concerning potential issues around the accuracy of data labels, the manufacturing customer realized that they need to take a step back and invest the time and effort to create more accurate and complete labels for their data set.



## THE CONCLUSION

This predictive maintenance project for an industry leader in construction and mining equipment manufacturing involved a dynamic and collaborative approach with the customer to understand the trends Mosaic saw in the data, and what it means for a failure to be on the horizon. At the end of the engagement, the primary conclusion of the analysis was that predictive maintenance does indeed seem possible – but new, carefully labeled data would be needed to achieve production-ready results.

This project successfully demonstrated one of the core challenges of machine learning and AI: obtaining accurate and consistent training labels. Models can only be as good as the data used to learn and generalize from. When a model isn't improving, it is often essential to take a step back and ask ourselves if the information we are asking the model to learn from is as representative as it possibly can be.

The customer now understands the steps necessary to build accurate predictive maintenance models for their fleet of industrial machinery. Armed with these insights, the company can prioritize creating accurate and consistent labels for their data and in order to continue developing, testing, and integrating machine learning into solving its predictive maintenance tasks.

