Predictive Maintenance using Clustering Methods for the use-case of Bolted Connections in the Automotive Industry

E.L. Bekker[1], J. Grobler[1] & S. Matope[1*]

ARTICLE INFO

Article details
Submitted by authors 12 Sep 2024
Accepted for publication 4 Apr 2024
Available online 30 May 2025

Contact details
* Corresponding author
smatope@sun.ac.za

Author affiliations
1 Department of Industrial Engineering, Stellenbosch University, Stellenbosch, South Africa

ORCID® identifiers
S. Matope
https://orcid.org/0000-0001-5598-1027

J. Grobler
https://orcid.org/0000-0002-1868-0759

E.L. Bekker
https://orcid.org/0009-0004-1284-3248

DOI
http://dx.doi.org//10.7166/36-1-3106

ABSTRACT

The paper explores the feasibility of using clustering for the predictive maintenance of nutrunners in a modern high-volume manufacturing environment. During the period of the study, the failure rate of one nutrunner was significantly higher than that of the others. The clustering algorithms that were evaluated were agglomerative hierarchical clustering (AHC), density-based spatial clustering of applications with noise (DBSCAN), and a self-organising feature map (SOFM). The performance metrics used to compare and evaluate the clusters were the silhouette coefficient score (SC) and the variation rate criterion (VRC). It was found that it is feasible to use clustering to improve the maintenance strategy of nutrunners in the automotive industry.

OPSOMMING

Die studie ondersoek die uitvoerbaarheid van die gebruik van groepering vir die voorspellende instandhouding van moerindraaiers in 'n moderne hoëvolume vervaardigingsomgewing. Gedurende die tydperk van die studie was die mislukkingsyfer van een moerindraaieraansienlik hoër as dié van die ander. Die groeperingsalgoritmes wat geëvalueer is, was agglomeratiewe hiërargiese groepering, digtheid-gebaseerde ruimtelike groepering van toepassings met geraas, en 'n selforganiserende kenmerkkaart. Die prestasiemaatstawwe wat gebruik is om die groepe te vergelyk en te evalueer, was die silhoeëtkoëffisiënttelling en die variasietempo-kriterium. Daar is gevind dat dit haalbaar is om groepering te gebruik om die instandhoudingstrategie van moerindraaiers in die motorbedryf te verbeter.

1. INTRODUCTION

The automotive industry is one of the world’s largest industries by revenue. According to the International Organization of Motor Vehicle Manufacturers, the industry saw an increase in production from 54 million cars in 1997 to a peak of 97 million in 2017 [1]. Two major concerns for the automotive industry are user safety and regulatory compliance. In extreme cases, non-compliance with international and local standards can lead to a product recall and/or costly litigation. To avoid such extreme cases, the ISO 26262 standard is used because it is considered best practice for achieving automotive functional safety [2].

Non-permanent bolting connection quality is one of the critical considerations for compliance with international standards. These connections are generally classified as critical when there would be an immediate danger to the driver of the car if the connection were to fail. The standards require that these connections be performed with electronically controlled (EC) nutrunners that provide reliable and accurate process control of the bolting process and simplify the manufacturing. EC nutrunners are expensive to implement, but are a critical part of the technology systems for modern automotive manufacturing. Nutrunners come equipped with torque transducers that measure and record the applied torque during the bolting process.

All technology systems and items of machinery require maintenance because of operational wear and unexpected breakdowns. Recent progress in industrial practice has resulted in the evolution of maintenance management from ‘a necessary evil’ to a competitive advantage, consequently producing greater profit, better product quality, and increased technical availability of machines. The onset of Industry 4.0 has enabled improved predictive maintenance strategies, as machines now have the ability to track and record operating conditions by using external and internal monitoring and sensors [3]. Access to the process parameters of machinery has made it possible to understand and predict better the condition of the system, allowing improved predictive maintenance.

This study aimed to determine whether the process data that has already been captured and stored could be used to understand better the condition of each nutrunner and then be used to perform predictive maintenance on the system. A predictive maintenance approach leads to better scheduling of maintenance work, improved process quality, and increased uptime. The success of the implementation is measured through the ability to understand better the technical availability of the nutrunner system, its product quality, and the use of maintenance resources through data analytics on the process parameters.

2. LITERATURE REVIEW

The literature review covers the relationship between Industry 4.0 and predictive maintenance in Section 2.1; Section 2.2 provides a brief description of the three algorithms used for clustering; and Section 2.3 describes the performance metrics used.

2.1. Industry 4.0 and predictive maintenance

The fourth industrial revolution started in the 1990s, and is ongoing [4]. It is characterised by the rise of the internet, and has paved the way for predictive maintenance, the most advanced strategy to date. The common definition of predictive maintenance is the ability regularly to monitor actual mechanical conditions, process parameters, and operating efficiencies to manage better the interval between repairs and to reduce the number and cost of unscheduled outages created by machine-train failures [5].

New opportunities to use optimally all the data gathered from machines opened up with the rise of Industry 4.0 and the implementation of ‘smart machines’. Smart factories or machines are highly digitised systems that have the ability to collect and share data continuously throughout the process. They combine modern machinery with Internet of Things (IoT) systems to monitor and measure process parameters better as the machine performs the needed actions [6].

The biggest obstacle in an effective predictive maintenance application is finding the appropriate machine learning (ML) algorithm for the dataset [7]. Bayesian networks have been suggested for diagnosing and predicting faults in larger datasets where little to no information is given on the variables [8]. Recurrent neural networks (RNNs) have been found to be particularly useful for time-series data, but they lack in long-term remaining-useful-life RUL predictions where long short-term memory (LSTM) NNs are preferred [9]. However, an LSTM is sensitive to dataset changes, and so is not ideal for real-life production data [10]. Last, transformer-based approaches have recently been implemented to forecast time-series data, as they seem to outperform other ML models in this regard [11].

Furthermore, purely technology-based maintenance strategies can easily fail without a clear business objective, as the deployment of IoT sensors can result in a poor return on investment or produce insufficient statistically significant findings. IoT sensors can collect, connect, and send data through network connections or through cloud-based solutions [12]. IoT implementation also needs to go hand-in-hand with a wider shift in a business’ processes and workflow. Buy-in from all crucial stakeholders is important to ensure that the correct strategy and process is followed [6].

2.2. Clustering algorithms

Three clustering algorithms are relevant to this study, namely agglomerative hierarchical clustering (AHC), density-based spatial clustering of applications with noise (DBSCAN), and a self-organising feature map (SOFM) with ward clustering.

2.2.1. *Agglomerative hierarchical clustering*

AHC [13] follows a bottom-up hierarchical approach in which each observation starts in its own cluster, and then the algorithm successively agglomerates pairs of clusters. The process is repeated until all clusters have been merged into a single cluster that represents all the data.

2.2.2. *Density-based spatial clustering of applications with noise*

DBSCAN [14] is a density-based clustering technique that groups together points that are within a user-defined distance (€) from one another and labels points outside these bounds as outliers or noise. The algorithm allocates points as *core points* (of point **P**) if there are at least the minimum number of points *(minPts)*, including point **P** itself, within the allocated neighbouring distance from the core. The distance is commonly calculated as the Euclidean distance, and points outside these bounds are labelled as outliers or noise.

2.2.3. *Self-organising feature map with ward clustering*

SOFM is an unsupervised artificial neural network that is trained with competitive learning techniques [15]. SOFM is commonly used to process high-dimensional data into a low-dimensional data set while retaining the topological structure of the data. The output of the SOFM algorithm is typically a two-dimensional representation of the input data set [16].

2.3. Performance metrics

Two performance metrics, namely the silhouette coefficient (SC) and the variance ratio criterion (VRC), are discussed in this section.

2.3.1. *Silhouette coefficient*

SC is a metric that can be used when the ground-truth of the data is not known or is unavailable [16]. SC is a combination metric of the intra-cluster cohesion and the inter-cluster separation of clusters in a data set, and ranges from -1 to 1. A value close to the -1 limit indicates poor cluster cohesion and separation; on the other hand, a score closer to the upper bound of 1 indicates a well-balanced clustering between intra-cluster cohesion and inter-cluster separation [18,19].

2.3.2. *Variance ratio criterion*

VRC is a ratio between the sum of within-cluster dispersion (SS_B) and between-clusters dispersion (SS_W). A high value of SS_B indicates distinct and well-formed clusters. A small value of SS_W shows well-formed clusters. The VRC value is a positive integer when the highest value indicates optimal clustering, as the between-cluster dispersion is significantly larger than the within-cluster dispersion [20].

3. BOLTED CONNECTIONS

The objective of a bolted connection is to create a non-permanent clamping force between two or more surfaces. A bolted joint presents various difficulties, as the joint changes in response to service and environment [17]. Section 3.1 describes the bolting process, and in Section 3.2 a failure analysis of one stage of the process, the rundown stage, is provided.

3.1. The bolting process

A bolting process occurs in four main stages: rundown, alignment, clamping, and post yield. Figure 1 shows a typical two-phase bolting curve of the input torque versus the angle of turn of the bolt. During the rundown stage, the bolt has not yet made contact with the mating surface, and the required torque is caused by the friction in the threads. Point A on Figure 1 is defined as the maximum torque reading during rundown and, if a sudden spike in this value is measured, it may indicate misalignment.

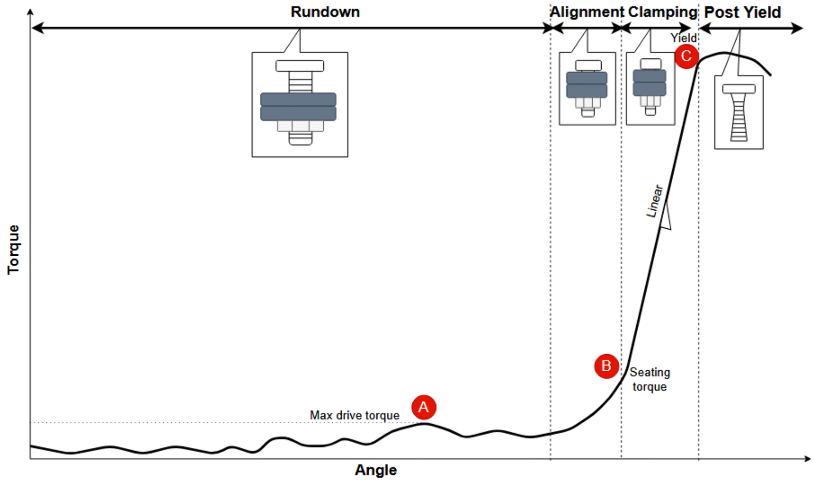

Figure 1: Simple two-phase bolting curve showing the torque applied vs the angle of turn

The alignment stage begins when mating between the two surfaces has occurred. The curved radius of this region is non-linear, and is influenced mainly by the stiffness of the joint. It is the start of preload build-up in the connection. Point B on Figure 1 indicates the seating-torque, and is the point where the curve becomes linear. The clamping stage is linear, and should remain consistent over identical bolting connections. The gradient of the curve is given by Equation 1:

$$Slope = \frac{(K_B K_J)}{(K_B + K_J)} \frac{P}{360} \tag{1}$$

and represents a linear relationship between the pitch of the thread (*P*), the friction coefficients of the bolt (K_B), and the joint members (K_J). The curve remains linear until Point C, where one of the mating surfaces starts to yield and failure in the connection occurs. Post yield, the bolt deforms elastically, which is commonly called 'necking', until ultimately failure occurs and it breaks. The area beneath the curve of Figure 1 can also be used to calculate the input energy.

3.2. Failure analysis for rundown

Rundown is defined as the time from the start of the bolting process to the point where mating of the two surfaces has started. During this phase, it is expected that there will be no torque build-up, as there is no clamping force between the surfaces before mating. There should also be very limited variation in torque readings. A sudden spike during rundown could indicate misalignment of the manipulator, the nutrunner, or the bolt. In the case of the misalignment of the manipulator, all nutrunners in the system show similar spikes, as the entire system is misaligned. If only one nutrunner shows sudden spikes, that indicates that the misalignment is limited to a singular nutrunner.

Multiple spikes and irregularities during rundown could indicate a worn socket, as the process deviates from the expected graph. Furthermore, achieving the target torque takes longer, as the socket slipping does not result in the same amount of true angle turn.

Another prominent failure is when there is no torque build-up for the entire cycle time. No torque build-up indicates that the nutrunner is not turning the bolt, as there is no sign of mating. This situation triggers the system to stop, and the cycle is labelled 'NOK', as the bolt did not reach the target function.

Table 1 summarises the most frequent failures identified during rundown with the added effect on the graph in Figure 2

Table 1: Failure mode and effects analysis (FMEA) for rundown stage

Possible cause	Affected parts	Recovery	Graph	Feature	Figure 2
Socket worn	Socket	Replace socket	Torque vs time	Jagged rundown; many spikes in rundown	Green
Misalignment of manipulator	Manipulator/robot	Correct alignment according to jig	Torque vs time	No torque build-up; sudden spike in rundown; only failures of bolts associated with that manipulator	Yellow
Output drive square worn	Output drive	Replace output drive	Torque vs time	Deviations in phase; only one nutrunner affected	Purple
Output drive square worn	Output drive	Replace output drive	Torque vs time	No torque build-up during rundown	Purple

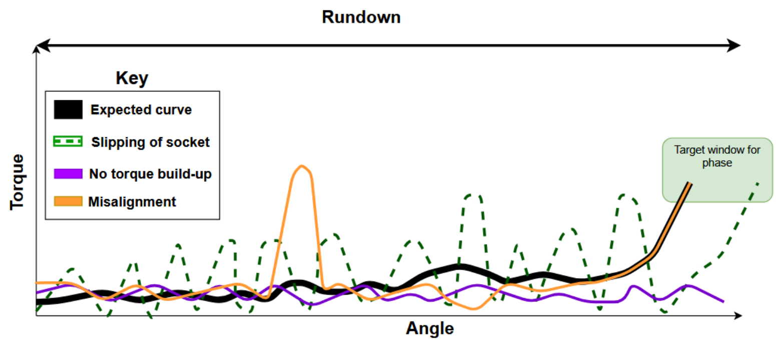

Figure 2: Example illustration of failure features for rundown region

4. METHODOLOGY AND EXPERIMENTAL SETUP

This section describes the methodology and experimental setup used for the study. Section 4.1 outlines the data collection setup, while Section 4.2 explains the operator process during the wheel bolting process, and Section 4.3 illustrates the engineered features used to generate the final dataset.

4.1. Data collection setup

Data collection is crucial in this setup, as regulations require that the data for every car be stored for up to 30 years. It is required to save only the 'NOK' (if applicable) and the subsequent 'OK' result for each bolt and the final measurements. The process data is received directly from each controller and saved in a separate data lake. The data is sent only once there are data for all five bolts on each wheel. The supervisory control and data acquisition (SCADA) overview is shown in Figure 3.

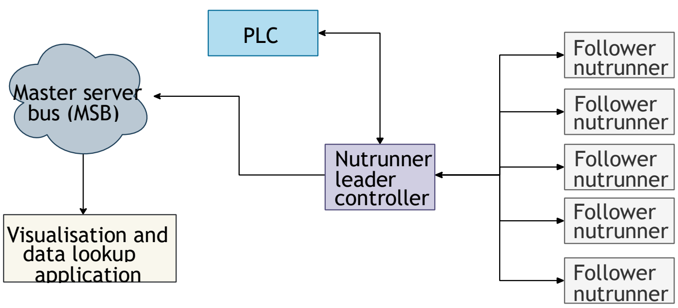

Figure 3: SCADA overview of the bolting system

The controller sends the information to the master server bus (MSB), where it is stored until it is queried by the visualisation software currently used by the plant technicians. The data sets of torque, angle, and time used in this case study were retrieved from the same software. The 30 years of results that the regulations required are sent to a separate database. They were not tampered with, downloaded, or used in any way for this case study.

4.2. Process setup

The process setup prescribed by the manufacturing guidelines could not be altered for this paper, but it would be important to understand it fully. The bolting process is specified by the design engineers, and cannot be changed without investigation and a proper procedure of process change.

The process can be divided into four main steps, as shown in Figure 4. It begins in Block A, where the operators pre-tighten the front wheels of the car with a battery-operated tool set that shuts off when any one of the five bolts reaches 6 Nm. The operators then move the manipulators with the nutrunners to the front wheels and give the input signal to start the bolting process. A programmable logic controller takes over and instructs the controller to start the bolting process. While the bolting process is occurring on the front of the car, the operators bring the wheels to the rear of the car and pre-tighten the bolts with the same 6 Nm battery-operated tool. Once the manipulators have released and finished the bolting process on the front wheels, the operators move the manipulator to the rear wheels and the cycle reaches Block C. While the manipulator tightens the rear wheels (Block D), the operators pretighten the front wheels of the new car that has entered the station, and the cycle is repeated.

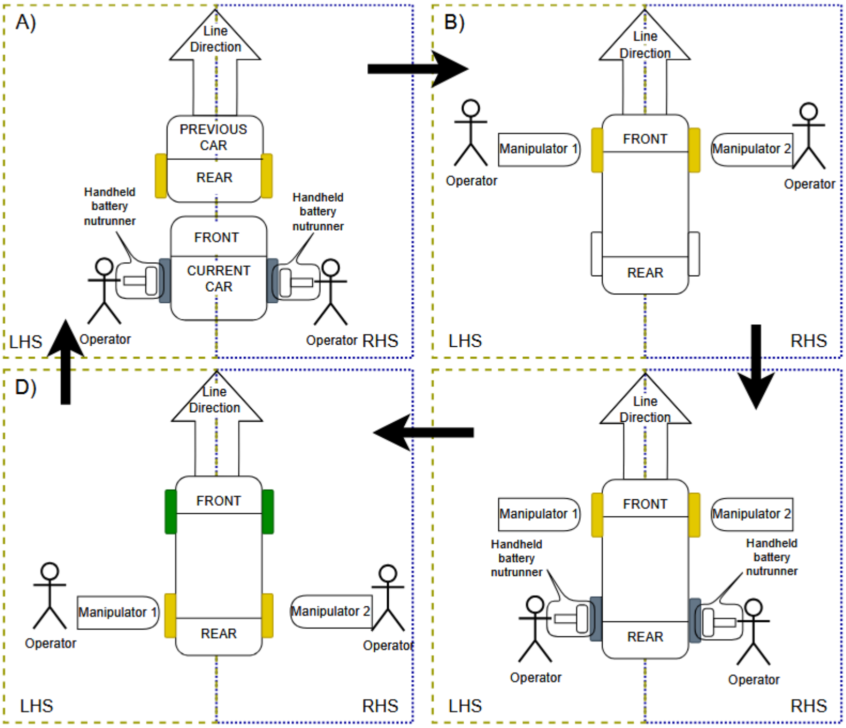

Figure 4: Process setup in factory

The station consists of two manipulators on either side of the car. Each manipulator has five nutrunners that bolt the front and rear wheels of each car. Thus 20 bolts per car per cycle are tighterned by 10 different nutrunners.

The data was collected from February 2020 to May 2021 from the data lake as a compressed comma-separated-values file. It contained all the information for torque, time, and angle measurements for all 16 stages. Figure 5 shows an example of a graph plotted for the torque and time data.

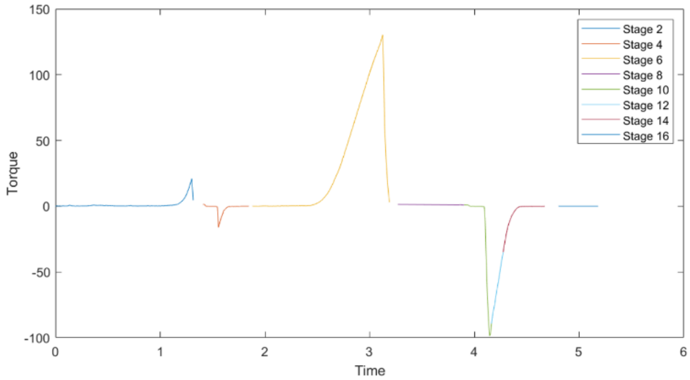

Figure 5: Example plot of the data for all 16 states for torque vs time

The system records each bolt that reaches the target function for that stage with an 'OK' result; and if the target function is not reached, it is marked with a 'NOK'. Once all five bolts have results, regardless of the status, the results are sent to the MSB and stored in the data lake. Each stage has a separate data entry, and contains a machine-evaluated 'OK' or 'NOK' for that specific phase. If a certain bolt fails, for example, at stage 6, that stage is the final entry for that specific car.

As a starting point, the chosen data set was for the LHS machine, nutrunners 1 and 4, and only stage 2. The starting point was based primarily on the 'NOK' rates, which indicated that nutrunner 4 had the worst 'OK' rate (97%), whereas nutrunner 1 had the best (99%); the majority of the failures (98%) occurred in Stage 2. From an analysis of the failures, it could be seen that the area of interest for the failure mode of socket wearing was mainly in the rundown stage. For this particular study, the failure mode to be investigated was the socket failure and rundown happening in stage 2.

Nutrunner 4 was selected because of the socket that was changed on 15 May 2021. Nutrunner 1 had the best 'OK' rate of the five nutrunners. The starting data set was thus stage 2 of nutrunner 1 and nutrunner 4 for the LHS machine. The most informative data for this particular case study was found in the rundown region and particularly in stage 2. The raw data needed to be converted into values that aligned with the FMEA findings described in section 3.

4.3. Feature engineering

The first set of features that were engineered is a summary of the statistics for each quarter of stage 2 for every car, as shown in Figure 6. These features were chosen to flag any unexpected spikes and deviations during the rundown that could have indicated a worn socket because the connection between the nutrunner and the bolt was not secure. The time to the maximum torque was also recorded and, in the case of a result that did not reach 20 Nm, the maximum time was recorded.

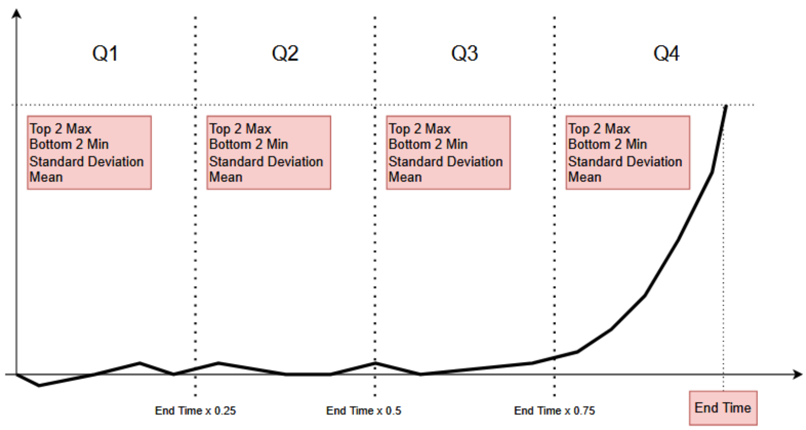

Figure 6: Illustration of the statistical features engineered for the final data set from stage 2

The second set of features represented the power dissipated during the alignment phase by calculating the area under the curve from the maximum torque to 20% of the maximum torque, as shown in Figure 7. This logic allowed the algorithm to evaluate the results for which the expected maximum torque (20 Nm) was not reached, as it was dynamically established with every cycle. All the features were checked for correlations, but none were found. The final dataset thus contained all the engineered features.

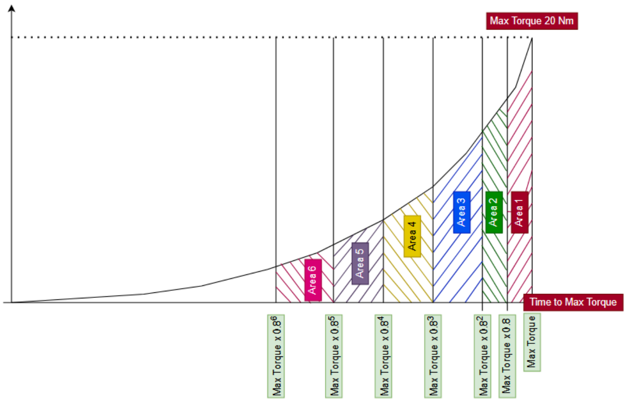

Figure 7: Illustration of the engineered features, based on the various areas under the graph for the final data set for stage 2

5. RESULTS AND CLUSTER ANALYSIS

Section 5.1 explains the method for choosing the clusters that were used for the cluster analysis in Section 5.2. Recommendations for deployment are given in Section 5.3.

5.1. Clustering algorithm evaluation

The final dataset that was developed, as described in Section 4, was clustered using the AHC, DBSCAN, and SOFM algorithms, which underwent a hyperparameter tuning process. The best-performing algorithm was then used for further cluster analysis. The hypothesis test had two possible outcomes: if one algorithm statistically outperformed the other, the winning algorithm would be allocated a 'win' and the losing algorithm would be allocated a 'loss'; while the second possible outcome would occur when neither algorithm statistically outperformed the other, and both would be allocated a *draw*. The resulting hypothesis score was calculated by Equation 2:

$$Total = Wins - Draws - Losses \tag{2}$$

Table 2 shows the results from the hypothesis scores when comparing the algorithms on the basis of the SC metric. The SOFM implementation was the clear winner, with a score of 2. The AHC's best performing parameters, with n = 5 and *ward* linkage, was the worst-performing algorithm of the clustering algorithms, with a hypothesis score of -2 based on the SC.

Table 2: Overall hypothesis test for SC to compare the three best-performing results from each clustering algorithm

Iteration	Wins	Losses	Draws	Total
AHC	0	1	1	-2
DBSCAN	1	0	1	0
SOFM	2	0	0	2

Table 3 shows the results from the hypothesis scores when comparing the algorithms on the basis of the VRC metric. The SOFM implementation was the clear winner, with a score of 2. The AHC's best-performing parameters, with n = 5 and *ward* linkage, was the worst-performing algorithm of the clustering algorithms, with a hypothesis score of -2 based on the VRC. The best-performing DBSCAN algorithm's parameters were tuned to 70 *minPts* and a € of 0.15.

Table 3: Overall hypothesis test for VRC to compare the three best-performing results from each clustering algorithm

Iteration	Wins	Loses	Draws	Total
AHC	0	1	1	-2
DBSCAN	1	0	1	0
SOFM	2	0	0	2

The best-performing algorithm was thus the SOFM implementation, with a gridsize of 3 x 3 and sigma of 10, as it statistically outperformed the other two algorithms. The resulting clusters are analysed further in the next section.

5.2. Cluster analysis

The cluster results for run 8 were analysed further to determine whether the results were valid. Cluster 6 contained 97% of the 'NOK' results. This label was not a feature, and was not used during the algorithm implementation. Therefore, cluster 6 best represented the 'NOK' and 'OK' results, which had abnormally high times to 20 Nm.

Figure 8A shows the average time to 20 Nm per cluster. Cluster 7 was closest to the overall average of 0.1, while Cluster 6 had the highest average. The variation in the clusters was explored further in the breakdown of each cluster.

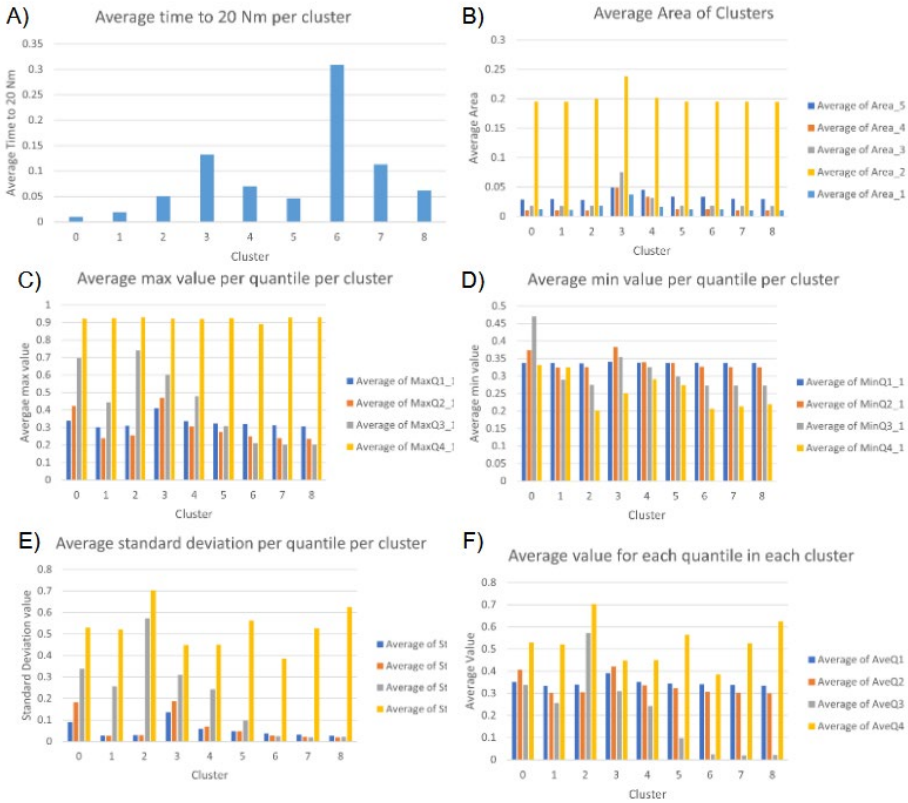

Figure 8: Summary graphs of the average values for features in each cluster

Figure 8B shows the average value of the areas under the curve illustrated in Figure 7. Areas 1, 2, and 3 were expected to be similar as they represented the rundown region. The values for area 4 and area 5 were expected to be similar, as they represented the snug radius of the curve; the hardness of the bolting process remained the same for each bolted wheel. Cluster 3 had the highest average of all the areas. Furthermore, cluster 4 showed a similar trend to that of cluster 3. The areas for all the other clusters were similar, and did not reveal any significant information. Quantiles 1 and 2 represented the rundown regions where mating had not yet taken place and where the torque build-up should not have been high. Quantiles 3 and 4 represented the alignment and clamping regions. Quantile 4 was the end of this bolting stage, with a target torque value of 20 Nm.

Figure 8C shows the average maximum values for each quantile per cluster. The highest average for quantile 1 and 2 occurred in Cluster 3. Quantile 3, which represented the alignment region, had high values in clusters 0, 2, and 3. Last, the average value for quantile 4 was similar throughout all the clusters, which was expected, as the target value for the stage remained the same.

Figure 8D shows the average minimum values for each quantile per cluster. The values were similar across all the clusters except for cluster 0, which had a much higher average value for quantile 3.

Figure 8E shows the average standard deviation for each quantile per cluster. Quantiles 1 and 2 were within the rundown region, and thus the standard deviation was expected to be small. A high standard deviation would indicate slipping of the socket or an unexpected impact on the bolting system. A small standard deviation was expected for quantile 3, as this was the non-linear alignment phase with a gradual increase in torque build-up. The highest standard deviation was expected in quantile 4, as this was the clamping phase and, owing to the high friction and the hardness of this application, a steep gradient was expected. Clusters 5, 6, 7, and 8 aligned with the expected trend. Cluster 3 had the highest average standard deviation for quantiles 1 and 2, which indicated spikes during rundown.

Figure 8F shows the average values of each quantile per cluster. Quantiles 1 and 2 were expected to have a low average. Quantile 3 was expected to have a higher average, with Quantile 4 having the highest average.

The rest of this section gives an in-depth analysis of each cluster, with the example graphs being compared with the expected normal graph. The expected normal graph was randomly selected from a bolting result that reached 20 Nm in the average expected time.

Clusters 0, 1, 5, 7, and **8** indicated, as expected, normal graphs for a pre-tightened connection. **Cluster 2** contained only ‘OK’ machine-labelled entries. It was also the smallest cluster, with only 0.01% of the data falling into it. The cluster had average area values, and the time to 20 Nm also fell within the tolerable range. The most noteworthy observation was the standard deviation values of quantiles 3 and 4. The standard deviation indicated a change in the snug radius, as shown in Figure 9, which indicated a change in the hardness of the mating components. The change in snug radius decreased the bolting time further, as the radius was steeper, resulting in a shorter alignment phase. This cluster indicated a quality defect more than a machine failure, and was useful for quality assurance.

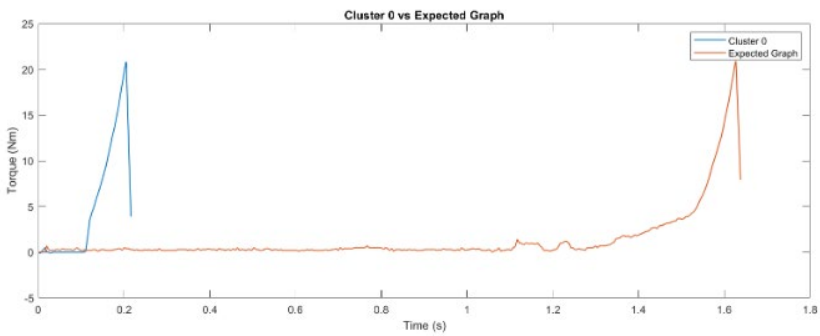

Figure 9: Cluster 2’s example graph compared with an expected normal bolting graph

Cluster 3 had the highest average values for areas 2, 3, 4, and 5. This cluster consisted of ‘OK’ evaluated results. Furthermore, the standard deviation of this cluster indicated spiking in the rundown phase. Figure 10 shows a clear pattern of a slipping socket compared with the expected normal. The spikes in the torque readings indicated the slipping of the socket, as the socket did not make sufficient contact with the bolt. Thus, cluster 3 represented the data of a worn socket.

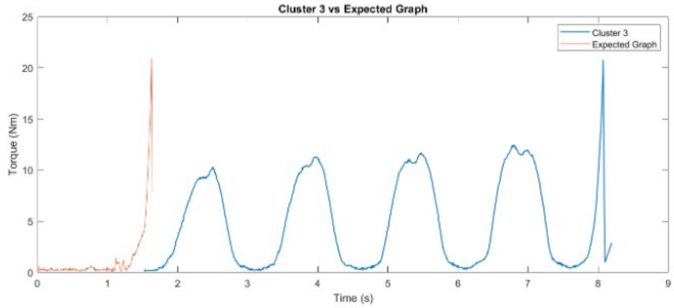

Figure 10: Cluster 3 example graph compared to an expected normal bolting graph. The spikes in the torque readings indicate slipping of the socket as proper contact between the socket and bolt was not achieved.

Cluster 4 contains only "OK" machine-evaluated entries with a slightly higher occurrence in nutrunner 1 than nutrunner 0. Cluster 4's most characteristic difference is the high standard deviation in quantile 3. The high standard deviation indicates a spike in the graph during the alignment region. Figure 11 shows a graph from cluster 4 vs the expected graph. A very clear spike in the graph can be seen which most likely indicates a misalignment in the bolting system. Cluster 4 shows a machine related fault, and should have been flagged for the attention of the maintenance team.

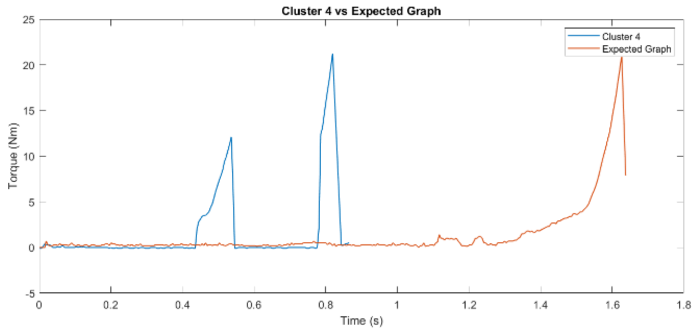

Figure 11: Cluster 4's example graph compared with an expected normal bolting graph

Cluster 6 had the highest number of 'NOK' machine-labelled entries. The machine-labelled status was not provided to the algorithms. This clustering was mainly because the average time to 20 Nm was the highest in this cluster, defaulting to 20 seconds if 20 Nm was not met on the basis of the data preparation algorithm. Figure 12 shows an example data entry of a 'NOK' bolting case in which little to no torque build-up occurred before the machine timed out and terminated the process. Furthermore, cluster 6 contained some data entries in which 20 Nm was reached, but in a significantly longer time than with the other clusters. The longer time could have been a result of operator error or caused by a manipulator alignment error. This cluster was problematic because it showed a process failure; and the majority of the failures had already been flagged and stopped by the current process setup.

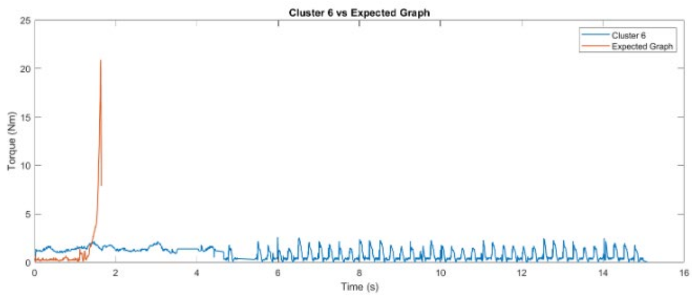

Figure 12: Cluster 6's example graph compared with an expected normal bolting graph

5.3. Recommendation about deployment

Table 4 summarises all nine clusters along with the required maintenance actions per cluster. Clusters 0, 1, 5, 6, 7, and 8 showed normal operating conditions with some expected variations. Another important consideration should be the extensive pre-processing and data cleaning that was needed. The raw data has a number of data quality issues that need to be cleaned and filtered before the clustering algorithms is possible to use. This phase is specific to the wheel bolting workstation and inspiration can be drawn from the results, but each bolting station will require a similar study to establish the data cleaning requirements. Multiple workstations are needed to better understand which data-cleaning techniques could be generalised and which were situation-specific.

Table 4: Summary of cluster classification showing whether the cluster indicated normal operating and, if not, whether maintenance was required

Cluster	Normal operation?	Root cause	Maintenance
0	Yes	N/A	N/A
1	Yes	N/A	N/A
2	No	Change in material hardness	Quality-related
3	No	Socket wearing	Change socket
4	No	Misalignment in system	Check alignment and adjust
5	Yes	N/A	N/A
6	No	‘NOK’ results due to various reasons	Maintenance team to investigate
7	Yes	N/A	N/A
8	Yes	N/A	N/A

5.4. Industry feedback

The findings from the clustering analysis were presented to key role players in the automotive industry. The participants in the feedback workshop included bolting specialists, maintenance managers, and process engineers. Their feedback was positive. Practical considerations for further implementation of the findings were raised, such as processing power requirements, generalisation of the model, and expanding the model to incorporate more failures.

6. CONCLUSION

This study investigated the feasibility of using a clustering algorithm as a baseline for the predictive maintenance of a bolting system. Three clustering algorithms were compared: AHC, DBSCAN, and SOFM. The SOFM statistically significantly outperformed the other two algorithms. The study showed promising results, and managed to identify all expected failures in the rundown region. However, the data cleaning and processing required significant time and domain knowledge, and was workstation-specific. The computational requirement should be taken into account, as the current system does not have enough computational power to run the needed algorithms. Future work would include the development of a generalised data cleaning and processing algorithm to deploy a plant-wide solution for all bolting workstations.

REFERENCES

[1] **International Organization of Motor Vehicle Manufacturers.** 2017. *2017 production statistics*. Retrieved 21 April 2022. Available from https://www.oica.net/category/production-statistics/

[2] **International Organization for Standardization.** 2018. *ISO 26262-10:2018. Road vehicles — Functional safety — Part 10: Guideline on ISO 26262*. ISO.

[3] **Poór, P., Ženíšek, D., & Basl, J.** 2019. Historical overview of maintenance management strategies: Development from breakdown maintenance to predictive maintenance in accordance with four industrial revolutions. In *Proceedings of the International Conference on Industrial Engineering and Operations Management, Pilsen,* Czech Republic, pp. 23–26.

[4] **Hořánek, V., & Basl,J.** 2018. Overview and comparison of tools for assessing the readiness of companies for Industry 4.0. *In: the 32nd International Business Information Management Association Conference,* pp. 1797 – 1807

[5] **Ayvaz, S., & Alpay, K.** 2021. Predictive maintenance system for production lines in manufacturing: A machine learning approach using IOT data in real-time. *Expert Systems with Applications*, 173, 114598.

[6] **Ippolito, M., Nikolakis, N., Cerquitelli, T., O'Mahony, N., Makris, S., & Macii, E.** 2021. Industrial digitisation and maintenance: Present and future. In Cerquitelli, T., Nikolakis, N., O'Mahony, N., Macii, E., Ippolito, M., & Makris, S. (eds), *Predictive maintenance in smart factories,* Springer, pp. 3–18.

[7] **Zonta, T., Da Costa, C. A., Da Rosa Righi, R., De Lima, M. J., Da Trindade, E. S., & Li, G. P.** 2020. Predictive maintenance in the Industry 4.0: A systematic literature review. *Computers & Industrial Engineering,* 150, 106889.

[8] **Carbery, C. M., Woods, R., & Marshall, A. H.** 2018. A Bayesian network based learning system for modelling faults in large-scale manufacturing. In *2018 IEEE International Conference on Industrial Technology (ICIT),* IEEE, pp. 1357-1362.

[9] **Yan, H., Qin, Y., Xiang, S., Wang, Y., & Chen, H.** 2020. Long-term gear life prediction based on ordered neurons LSTM neural networks. *Measurement*, 165, 108205.

[10] **Sayah, M., Guebli, D., Al Masry, Z., & Zerhouni, N.** 2021. Robustness testing framework for RUL prediction deep LSTM networks. *ISA Transactions*, 113, pp. 28–38.

[11] **Wu, N., Green, B., Ben, X., & O'Banion, S.** 2020. Deep transformer models for time series forecasting: The influenza prevalence case. *arXiv preprint arXiv:2001.08317.*

[12] **Li, S., Xu, L. D., & Zhao, S.** 2015. The Internet of Things: A survey. *Information Systems Frontiers*, 17, pp. 243-259.

[13] **Day, W. H., & Edelsbrunner, H.** 1984. Efficient algorithms for agglomerative hierarchical clustering methods. *Journal of Classification*, 1(1), pp. 7-24.

[14] **Ester, M., Kriegel, H. P., Sander, J., & Xu, X.** 1996. Density-based spatial clustering of applications with noise. In *International Conference Knowledge Discovery and Data Mining*, Vol. 240, No. 6, pp 1-8

[15] **Kubat, M.** 2017. *An introduction to machine learning* 2nd ed. Springer.

[16] **Kohonen, T.** 1990. The self-organizing map. *Proceedings of the IEEE*, 78(9), pp. 1464–1480.

[17] **Bickford J. H.** 2007. *Introduction to the design and behavior of bolted joints: Non-gasketed joints* 4th ed. CRC Press.

[18] **Bonaccorso, G.** 2020. *Mastering machine learning algorithms: Expert techniques for implementing popular machine learning algorithms, fine-tuning your models, and understanding how they work* 2nd ed. Packt Publishing Ltd.

[19] **Rousseeuw, P. J.** 1987. Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *Journal of Computational and Applied Mathematics*, 20, pp. 53–65.

[20] **Caliński, T., & Harabasz, J.** 1974. A dendrite method for cluster analysis. *Communications in Statistics: Theory and Methods*, 3(1), pp. 1-27.