Computers & Industrial Engineering

Predictive Maintenance of Baggage Handling Conveyors using IoT

--Manuscript Draft--

Manuscript Number:	CAIE-D-21-03835R3
Article Type:	Research Paper
Keywords:	condition monitoring; Conveyors; Baggage Handling; Predictive maintenance; Machine Learning
Corresponding Author:	Vishal Gupta Indian Institute of Technology Kharagpur Kharagpur, INDIA
First Author:	Vishal Gupta
Order of Authors:	Vishal Gupta
	Rony Mitra
	Frank Koenig
	Maneesh Kumar
	Manoj Kumar Tiwari
Abstract:	This article discusses issues related to the maintenance of airports' baggage handling systems and assesses the feasibility of using predictive maintenance instead of periodic maintenance. The unique issues related to baggage handling systems are discussed - namely random noise captured by the IoT sensors due to the movement of the luggage and complex interconnected components that constitute the conveyors. The paper presents a scalable and economical maintenance 4.0 solution for the system using data from sensors installed on a live system in absence of historical data. Differentiating between anomaly detection and outlier detection the paper presents an algorithm useful to remove noisy data from the datasets. Using integrated machine learning approaches, it tries to detect and diagnose incumbent defects in the early stage to avoid breakdowns. The paper proposes an automated machine learning pipeline by processing unstructured industrial data. The performance of various machine learning algorithms on the collected data is compared. Finally, the paper discusses avenues for future research.

To **Date: 19-12-2022**

The Editor-in-Chief

Computers & Industrial Engineering

Please find an enclosed copy of the manuscript titled "**Predictive Maintenance of Baggage Handling Conveyors using IoT**" by Vishal Kumar Gupta, Rony Mitra, Frank Koenig, Maneesh Kumar, and Manoj Kumar Tiwari for publication in the regular issue of Computers & Industrial Engineering.

In this study, we discussed the issues related to the maintenance of baggage handling conveyors and assessed the feasibility of using predictive maintenance. Baggage handling conveyors are found in most large airports. Airports like Heathrow can have thousands of conveyors where maintenance is mainly done in conventional ways. Maintenance of such high-volume assets - where most of the parts are hidden - is a daunting task that calls for predictive maintenance. There is a lack of literature dealing with predictive maintenance of conveyors - particularly for baggage handling systems. Therefore, we humbly present some unique challenges and solutions to them in this paper.

We have followed all ethics and publishing standard while preparing this manuscript. We also guarantee that the research work presented in this manuscript is original and not submitted in another journal simultaneously. We would be honoured to get our manuscript processed in your journal.

Sincerely,

Authors

Title Page (with Author names, affiliations, Acknowledgement)

Title: Predictive Maintenance of Baggage Handling Conveyors using IoT

Vishal Gupta[1], Rony Mitra[1], Frank Koenig[3], Maneesh Kumar[4], Manoj Kumar Tiwari[1,2]

[1] Indian Institute of Technology Kharagpur, West Bengal, India

[2] National Institute of Industrial Engineering, Mumbai, India

[3] Business School, University of Buckingham, Buckingham, UK

[4] Business School, Cardiff University, Cardiff, UK

*Corresponding Author Information

Vishal Kumar Gupta

Department of Industrial and Systems Engineering, Indian Institute of Technology, Kharagpur, West Bengal, 721302, India. Email: vkg03@hotmail.com

Email-ids of other authors

Vishal Kumar Gupta	E-mail: vkg03@hotmail.com
Rony Mitra	E-mail: mitrarony92@gmail.com
Frank Koenig	E-mail: frank.koenig@siemens.com
Maneesh Kumar	E-mail: kumarm8@cardiff.ac.uk
Manoj Kumar Tiwari	E-mail: mkt09@hotmail.com

Highlights

- Presented a scalable and economical maintenance 4.0 solution for conveyors used in baggage handling.
- Introduced an algorithm to remove non-white noisy data from the datasets.
- Detected and diagnosed incumbent defects in the early stage to avoid breakdowns by using machine learning approaches.
- Proposed an automated machine learning pipeline by processing unstructured industrial data.

Response Letter

We sincerely thank the editors and referee for their time and effort in processing and reviewing this manuscript. In the review comments, the referee has asked for certain clarifications. Our response to those comments can be found below.

Reviewer #2: The reviewer suggests the authors clarify the following comments before further steps. But other than this comment, the reviewer has no further comments.

Comment - Following up on comment two from the last review, did the authors test on other S-Lifts (for example, training data from S-Lift 1 and test on S-Lift 2)? or is all the data collected from the same S-Lift? If all the data are from the same S-Lift, then it might not be able to apply to other S-Lifts. And the reviewer suggests the author clarify and add to limitations.

Response – We acknowledge the comments of the learned referee.

The reviewer is correct to point out that data from one type of S-Lift may not be applied to another. In our case, the S-Lifts installed in the baggage handling systems are identical. They were sourced from the same manufacturer and have the same specifications/speed. Therefore, the model developed on one set of S-Lifts can be applied to other S-Lifts in this baggage handling system. We want to clarify that in total eight identical conveyors were selected for experimental studies, however, we did not separate those conveyors into training and test conveyors. Training and testing were done after the data was collected and already labeled (ignoring the conveyor-id from which they were collected).

As suggested, we have added limitations stating the non-usability of data from one type of S-Lifts to other kinds of S-Lifts. We have also clarified that training and test data were not collected from separate conveyors; instead, training and testing were done ignoring the conveyor-ids from which they were collected.

Finally, we extend our sincere gratitude for the valuable time that you spent reviewing our paper. We are thankful for all the constructive comments provided to us during the review process. We hope we were able to address the comment.

Credit Author Statement

Vishal Kumar Gupta: Methodology, Data curation, Formal analysis, Investigation, Visualization, Roles/Writing - original draft, Writing - review & editing;

Rony Mitra: Software, Validation, Writing - review & editing;

Frank Koenig: Funding acquisition, Project administration, Writing - review & editing;

Maneesh Kumar: Conceptualization, Supervision, Writing - review & editing;

Manoj Kumar Tiwari: Resources , Supervision.

Predictive Maintenance of Baggage Handling Conveyors using IoT

Abstract

This article discusses issues related to the maintenance of airports' baggage handling systems and assesses the feasibility of using predictive maintenance instead of periodic maintenance. The unique issues related to baggage handling systems are discussed - namely random noise captured by the IoT sensors due to the movement of the luggage and complex interconnected components that constitute the conveyors. The paper presents a scalable and economical maintenance 4.0 solution for such a system using data from sensors installed (on a live system in absence of historical data). Differentiating between anomaly detection and outlier detection the paper presents an algorithm that can be used to remove idle and noisy data from the datasets. Using integrated machine learning approaches, it tries to detect and diagnose incumbent defects in the early stage to avoid breakdowns. The paper proposes an automated machine-learning pipeline by processing unstructured industrial data. The performance of various machine learning algorithms on the collected data is compared. Finally, the paper discusses avenues for future research.

Keywords— Condition monitoring, Conveyors, Baggage handling, Predictive maintenance, Machine learning

1 Introduction

Conveyors are remarkably useful and cost-effective devices that are used to carry items from one location to another. They transport goods cumbersome to handle manually, making them an integral part of material handling and numerous factory operations. They are hassle-free and enable faster operations making businesses more productive and profitable. Hidden from view, almost all international airports operate a dedicated system of conveyors, called baggage handling systems. The purpose of having such a system is to transfer checked-in luggage from one point to the other within the airport. A baggage handling system reduces operational costs in airports while also improving aircraft turn-around times. It improves the working environment for laborers and reduces the volume of damaged and misplaced baggage. Unsurprisingly, the global airport baggage handling system's market was valued at $8,504.2 million in 2017 and was projected to reach $14,509.9 million by 2025, registering a compounded annual growth of 6.7% from 2018 to 2025 (Sawant and Kakade, 2018).

In 2018, airlines carried 4.65 billion bags, and the number of passengers traveling via airways was growing by 7% year-on-year. There was significant pressure on airports to expand their baggage handling systems to accommodate more flights ("Air transport industry insights, The Baggage Report" 2018). Shorter transfer times between flights and reloading windows as airports try to accommodate more airlines meant even the slightest disruption in baggage handling caused snowballing effects. Unplanned downtimes caused luggage to miss flights (H. Peng and Zhu, 2017) or led to flight delays (Koenig, P. A. Found, Kumar, and Rich, 2020). Luggage that misses the flight requires additional logistics at the departure and arrival airports, escalating the costs of baggage transfer. According to the Montreal Convention 1999, a passenger is entitled to $1,600 for expenses incurred because of lost or delayed luggage (Meltzer, 2017). Any delay in the transportation of baggage leads to direct financial losses to the airlines which they typically back-charge to the airport. Therefore, effective baggage handling is a priority for airports around the world.

1.1 Problem Description

A baggage handling system can have several kilometers of belt and tray types of conveyors (see Figure 1). For the maintenance of these conveyors, most airports follow time-based maintenance policies such as monthly inspections and yearly overhauls. During yearly overhauls, most of the moving parts such as bearings and belts are replaced regardless of their condition. Furthermore, these baggage handling systems employ hundreds of technicians to keep the system healthy and operational. The major part of their job is to inspect conveyors as frequently as - every day - for signs of deterioration and abnormal noise. However, even with such a stringent maintenance regime, the maintenance strategy degenerates into run-to-failure which leads to unexpected breakdowns and severe disruption to passengers, airlines, and operations (Koenig, P. A. Found, and Kumar, 2019). Large airports like Heathrow can have anywhere from ten thousand to thirty thousand conveyors in total, periodic maintenance of such high-volume assets - where a majority of the parts are hidden and difficult to assess visually - is indeed a daunting task.

(a) belt conveyor (b) tray conveyor

Figure 1: Types of conveyors

It is known that a breakdown in a conveyor can take anywhere from a few hours to a couple of days to fix. The casing provided for the safety and isolation of the moving parts delays inspection and maintenance activities. Due to these safety measures and the complex structures, conveyors need partial dismantling even to replace a small component. The manual surveys are limited to visual inspections and are heavily dependent on the interpretation of technicians, their motivation, and their experience. To complicate things further, often assets are not easily accessible. The inspection windows in most cases are limited and vary depending on how busy the airports are. Due to these limitations, several major hub airports have suffered fatal failures leading to excessive downtime, disruption, and costly refurbishment. Furthermore, due to a sharp decline in the number of passengers during and post Covid-19, airports can no longer sustain a large workforce for periodic inspections. Limited literature in the field of predictive maintenance for time-critical assets suggests the use of condition-based maintenance solutions to predict the failure of assets in advance and take corrective actions before failure occurs in the system (Rijsenbrij and Ottjes, 2007; Koenig, P. Found, and Kumar, 2019; Koenig, P. A. Found, Kumar, and Rich, 2020).

A typical conveyor in a baggage handling system can have multiple assemblies, namely drive motor assembly, drive shaft assembly, tension shaft assembly, and idler shaft assembly. These assemblies have multiple components of their own, for example, a drive motor assembly consists of a drive motor, a gearbox, a hollow shaft and keyway, and a mounting bracket. Subsequently, we can have sub-components like bearings, gears, pulleys, idlers, belts, etc. What differentiates conveyors in baggage handling systems from other industrial assets is the speed of the operation and frequent start-stops. These conveyors operate only when they detect the presence of luggage. The conveyor stops soon after the luggage is transferred, and therefore the conveyors stand idle for most of the time. Not all luggage bags are the same and they come in various shapes, weights, and builts. The normal operating speeds can range from 1 to 2.5 meters per sec imparting enough inertia to the luggage that erratic noise is generated when they move and stop frequently.

These issues are not trivial when combined with Industry 4.0 construct that the IoT sensors were to be used for data acquisition - because wired sensors are simply not economical or scalable. We found that it is difficult to trigger IoT sensors externally to record clean data - i.e. when the conveyor is operational but not loaded. Another problem with IoT sensors is that they simply can not transfer high-volume data in real-time, therefore, the available sensors are designed to record vibration for 1-2 seconds and then transfer the recorded data over to the data server over the next 1-2 minutes. Due to the non-continuous usage of these assets, sensors are bound to record vibrations when the conveyors are idle (i.e. most of the time). Occasionally when sensors capture non-idle data, the captured data are noisy because of the movement of luggage (see Figure 2). In such a scenario, frequent false positives would make it difficult to carry out appropriate maintenance or inspection planning.

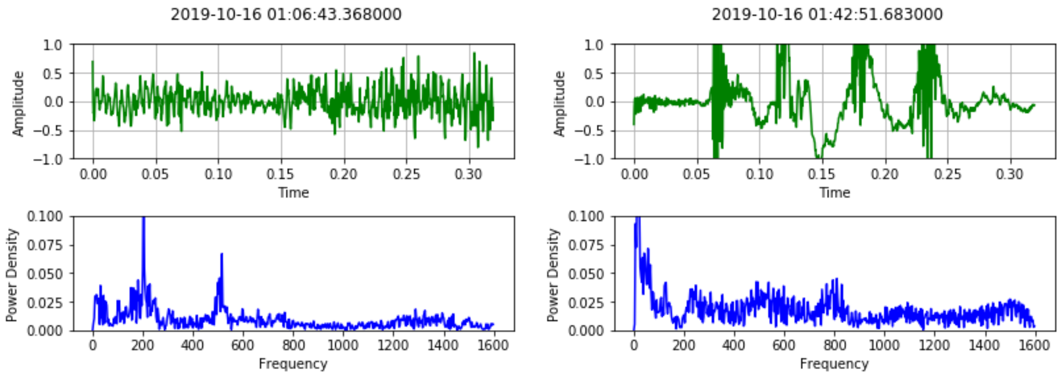

(a) No-load vibration

(b) With-load vibration

Figure 2: Potential load and no-load vibration recorded by a sensor

1.2 Research Gap

Alsyouf, Humaid, and Kamali (2014) documented that it is difficult to apply predictive maintenance (PdM) in baggage handling systems because of the complexity and diversity of the equipment and machinery installed. Therefore, most baggage handling systems still use the time-based maintenance approach. Very few research articles report the successful implementation of predictive maintenance on conveyors (see section 2.3). Few that dealt with conveyors did not consider variable load conditions and erratic noise as an issue that, in our understanding, can not be ignored in baggage handling systems. Furthermore, the low quality of data captured by IoT sensors makes data cleansing indispensable for advanced data analytics. As highlighted by Y. Liu et al. (2020), one of the limitations of IoT sensors is the quality of data that they capture. Cleansing of noisy data is one of the critical steps for accurate predictions, however, few studies distinguish between noise and anomalies while investigating the effects on defect detection results.

Other in-general issues related to predictive maintenance are (a) how to deal with lack of run-to-failure data and (b) how to deal with heterogeneous unstructured data to train machine learning models (Dalzochio et al., 2020). Lack of run-to-failure data: When sensors are newly installed, the data that gets collected does not represent all the possible defects that would eventually be observed if the sensors were installed and individual components were allowed to run till failure. Therefore, it is difficult to train predictive models in the absence of data that does not represent time series degradation and all failure patterns.

Oddly, previously published articles on predictive maintenance typically start at a point when a sufficient amount of data is already available for training and testing purposes. From extant literature, it can be seen that research on PdM is primarily based on near-ideal datasets obtained from accelerated degradation testbeds where one does not face real-life restrictions like (a) pieces of machinery takes years to show signs of degradation while (b) producing lots of poor quality data in the process (Lei et al., 2018).

In predictive maintenance, it is advisable to keep the machine learning models up to date. This requires keeping track of all the activities that take place on the asset under observation and labeling the sensor data accordingly. Not surprisingly, most companies keep records of all such activities in form of maintenance and spare parts usage logs. But, these logs are mostly unstructured and can be in the form of excel sheets stored in decentralized repositories; they often contain manually entered erroneous texts about defects and corrective actions, therefore can not be used directly for data labeling. It is not practical to manually read these reports and label all sensor data so that the predictive models could be updated and trained time-to-time. To the best of our knowledge, papers on predictive maintenance rarely talk about the need to process these unstructured heterogeneous data for model retraining.

1.3 Contribution

To transition from time-based periodic maintenance to condition-based predictive maintenance, one must develop the necessary hardware and software solutions to gather and process the signal and visualize diagnosis reports related to the condition of these assets. These solutions must also be suitable for brownfield projects where the conditions of conveyors are far from perfect. In the absence of historical data, the condition monitoring solution should enable a gradual shift from run-till-failure and time-based to condition-based and ultimately prediction-based with quick tangible benefits like cost savings in service efforts, reduced spare part consumption, and reduced downtimes. Considering these points, we aimed at designing a scalable and economical maintenance 4.0 prototype for baggage handling systems. We understand one-size-fits-all approach can not be applied to maintenance 4.0 and that the solutions are often customized according to the requirement of the industry. Therefore, we believe this article presents a unique perspective from the airport baggage handling industry.

We documented that the movement of luggage generates noisy data and hinders the implementation of predictive maintenance solutions for baggage handling systems. In this paper, we present how we overcame this problem. To solve the issue of lack of training data, we adopted integrated machine learning approaches that enable anomaly detection using multivariate analysis to catch defects in their early stage and thereby avoid breakdowns while machine learning (ML) trains to diagnose defects as more data are available. Furthermore, to facilitate seamless training of supervised learning models - we presented how heterogeneous unstructured data like inspection logs, work orders, spare parts logs, and maintenance logs can be processed. With the help of text processing and ontological reasoning, we present how the labeling and training of ML can be automated. Finally, we compared the efficacy of four ML algorithms, namely logistic regression, multi-layer perceptron, random forest, and support vector machine on real data. The results indicate that random forest outperforms other techniques.

The rest of the paper is structured as described next. In section 2, a literature review on condition monitoring and predictive maintenance is presented. Section 3 shows the solution framework adopted from PdM literature. In that section, we discuss how the framework was implemented for the case presented in this paper. Finally, sections 4 and 5 discuss the results and conclude the paper, respectively.

2 Literature Review

According to European Standards EN13306, maintenance strategies are grouped into (a) corrective maintenance and (b) preventive maintenance (see figure 3). Corrective maintenance takes place after the fault has occurred, and it can be either immediate or deferred. Maintenance activities are deferred if the assets under consideration are not critical. However, if the asset is critical to the throughput of the plant, immediate corrective actions are taken. Immediate corrective maintenance (often known as Run-to-failure) is a risky proposition because it leads to substantial downtimes and intervention costs. To eliminate such breakdowns and costly downtimes, preventive maintenance was introduced. According to the standard, preventive maintenance is either carried out in a predetermined time-based manner or based on the sensed condition of the asset.

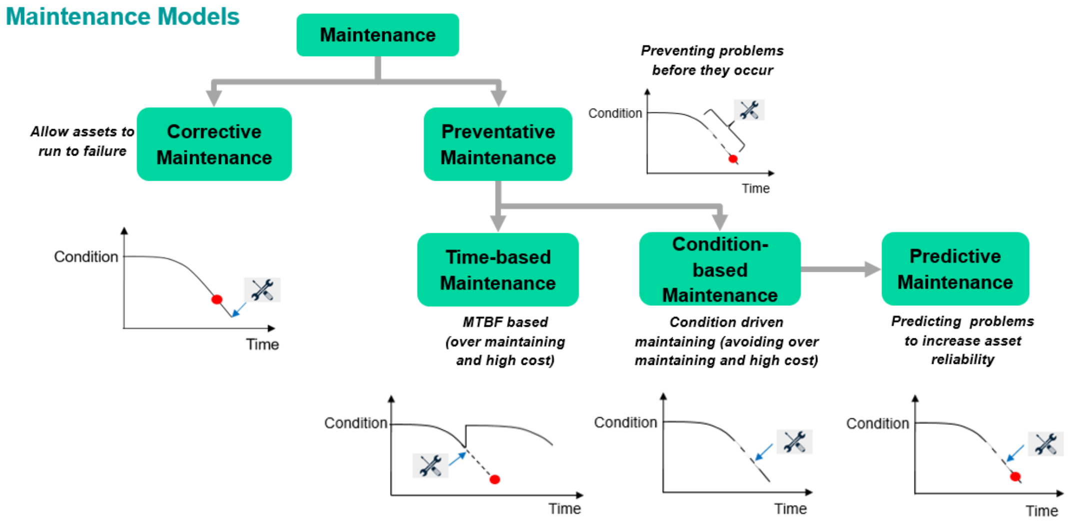

Figure 3: Maintenance Strategies - based on EN 13304 (Source: Schmidt and L. Wang (2018))

In time-based preventive maintenance, inspection and maintenance activities are scheduled in advance. However, such periodic maintenance lead to unnecessary inspection, corrective action, or parts replacement even if parts are not necessarily worn out. Irrespective of that, time-based maintenance does not guarantee zero unplanned downtime (Y. Peng, Dong, and Zuo, 2010). On the contrary, condition-based maintenance (CBM) aims at determining the equipment's working condition. CBM is based on the idea that 99% of asset failures are preceded by peculiar signs or indications that such a failure is about to happen (Ahmad and Kamaruddin, 2012). Theoretically, CBM relies on the detection of degradation patterns in machine components by frequent and sometimes continuous monitoring of health parameters like vibration, temperature, pressure, acoustic emissions, and so on. CBM systems must appraise the condition of equipment in real-time and provide servicing and replacement suggestions only when there is any evidence of abnormal behavior; it thereby aims to reduce operations and maintenance costs.

CBM systems are commonly based on (a) statistical or machine learning, or (b) physical modeling approaches (Jardine, D. Lin, and Banjevic, 2006). The idea behind the physical modeling approach is to utilize specific and explicit mathematical models to identify signals corresponding to defects of one's interest. This approach is mostly confined to rotating equipment like bearings, gearboxes, motors, and turbines as it requires mathematical modeling and needs to be validated in laboratories. However, a shop floor consists of a large variety of assets for which design specifications are not readily available; explicit physical modeling is significantly difficult for the assets operational in modern industries (Jardine, D. Lin, and Banjevic, 2006). Contrary to physical modeling, statistical or machine learning approaches do not require domain knowledge given enough data is available.

2.1 Predictive Maintenance

When condition-based maintenance uses statistical or machine learning tools like advanced data analytics, regression analysis, trend analysis, pattern recognition, multivariate correlation, etc., to anticipate failures in advance and to augment the decision-making process, it is called predictive maintenance (Schmidt and L. Wang, 2018; Sezer et al., 2018). Traditionally, predictive maintenance using statistical or machine learning techniques has not been easy due to the lack of low-cost sensor and information technologies essential for data acquisition (Y. Peng, Dong, and Zuo, 2010). Fortunately, low-cost Internet of Things (IoT) and modern wireless sensors have made real-time data acquisition affordable in recent years (Sezer et al., 2018; Strauss et al., 2018).

Other statistical approaches used in predictive maintenance include approaches like hypothesis testing, statistical process control, distance measures like Euclidean distance or Mahalanobis distance, feature vector correlation coefficient, hidden-Markov model, etc. (Jardine, D. Lin, and Banjevic, 2006). On the other hand, machine-learning-based predictive maintenance approaches use trend analysis and pattern recognition. Such maintenance systems rely on the clustering or classification of signals captured by the sensors. It is argued that in statistical approaches many of the features get overlooked while defining the diagnostics rules. Furthermore, defining rules using statistical inferences for each fault scenario is a time-consuming process and requires specific technical knowledge. Carvalho et al. (2019) state that out of statistical and machine learning approaches, the latter outperforms the former. The advantage that ML has over a statistical approach is that it can handle multivariate data easily and can identify hidden relationships between attributes without human interference; thereby can identify impending defects much in advance (Wuest et al., 2016).

In this regard, R. Liu et al. (2018) presented a review of machine learning and AI techniques for fault diagnosis for rotating machinery. The authors conclude that Support Vector Machine (SVM), Artificial Neural Networks (ANN), and deep learning methods tend to perform better when dealing with multi-dimensions and continuous features; while k-Nearest Neighbor (k-NN) and Naive Bayes tend to perform better when dealing with discrete features. Carvalho et al. (2019) carried out a systematic literature review on machine learning methods applied to PdM, highlighting the performance and limitations of state-of-the-art ML techniques. Authors documented that Random Forest was the most commonly used classification technique for fault diagnosis - followed by Neural Network Based approaches. Other supervised machine learning techniques such as k-NN, Bayesian classifier, and SVM are also frequently used. Recently, deep learning techniques like Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs) are also getting researchers' attention (Zhao et al., 2019).

Although fault prognosis has proven advantages over fault diagnosis, there are constraints to the application of fault prognosis in industrial environments. According to Carvalho et al. (2019), the factory must implement run-to-failure or time-based maintenance strategies to collect a sufficient amount of data to be able to label those data for machine learning. Some of such constraints as listed by Lei et al. (2018) include, (a) machinery may take months and sometimes years to show any sign of degradation while generating a huge amount of data; (b) it is difficult to label the data, as equipment mostly are not allowed to run till failure because that may lead to catastrophic breakdowns and losses; and (c) interference from outside environment degrades the quality of the data captured.

Another relatively less emphasized idea is the use of anomaly detection techniques for predictive maintenance. Anomaly detection, novelty detection, outlier detection, and deviation detection are synonymously used in a wide variety of applications mainly to detect rare patterns or abnormal data. For example - it is used - to detect abnormal traffic patterns in a network to identify hacking attempts, detect fraudulent transactions on a credit card, and detect the presence of a malignant tumor in MRI scans. Applications in domains such as medical and public health, fraud detection, intrusion detection, industrial maintenance, image processing, signal processing, robot behavior, and astronomy are well known. In condition monitoring - anomaly detection techniques like - distance-based approaches like Mahalanobis distance as in Y. Liu et al. (2020), density-based approaches like the Local Outlier Factor (LOF) as in Safaei et al. (2020), ensemble-based techniques like Isolation Forest as in Strauss et al. (2018), and one-class clustering using SVM as in Erfani et al. (2016) has already given appealing results.

Sikorska, Hodkiewicz, and Ma (2011) suggests: "... organizations must implement a staged approach; which involves maturing their existing diagnostic programs before progressing to advanced prognostics levels". We propose - because of the lack of historical (labeled) data - anomaly detection can be considered the first step toward predictive maintenance before progressing towards diagnostics and then prognostics.

2.2 PdM in Industries

This section reviews practical applications of condition-based or predictive maintenance. We have ignored the research articles if the data was generated in ideal lab-like environments or dealt with any specific mechanical or electrical component. Readers are requested to refer Carvalho et al. (2019) for PdM on individual components like pumps, fans, motors, turbines, gearboxes, and so on.

Implementation of predictive maintenance in an industrial setting was first documented by H. Li et al. (2014). The paper discusses how in the railroad industry intermediate maintenance was performed using data from strategically located sensors to avoid derailments. The authors highlighted that the pre-existing system used a rule-based approach for alarm generation but those alarms were not reliable and considerably reduced the network velocity. Authors thereby used decision trees and support vector machines to automatically learn rules from historical data to predict which rail cars are likely to have problems. Syafrudin et al. (2018) proposed a hybrid fault prediction system that used density-based spatial clustering for outlier detection (to clean the data) and random forest to classify events as normal or abnormal (to predict defects) in an automotive manufacturing assembly line in Korea. For data acquisition and data processing, the authors used Apache Kafka as a message queue, Apache Storm as a data processing engine, and MongoDB for data storage. The authors claimed that the solution was scalable considering its low network delays, low computational power, and low memory usage.

Fernandes et al. (2020) used ARIMA for fault detection in a metallurgical industry where without previous fault information the algorithm learns in an unsupervised manner. The authors used real sensor data (like spindle load, noise, vibration, and coolant level) obtained from CNC machines for anomaly detection and maintenance planning. In Ruiz-Sarmiento et al. (2020) carried out defect diagnosis on coiler drums within Steckel mills used in hot rolling processes in a steel sheets factory. Steckel mills work under severe mechanical and thermal stress and involved expensive replacements. The authors used Bayesian filters to estimate and predict gradual degradation patterns in the machinery. The proposed model fused experts' knowledge and replacement logs along with real data from the sensors to iteratively update the model. Bekar, Nyqvist, and Skoogh (2020) addressed the issue of data preprocessing using principal component analysis (PCA) and proposed the use of K-means clustering to gain insights into the behavior of machines in different working conditions. The authors carried out a case study to demonstrate the applicability of the approach to the real data obtained from two bottleneck machines installed in an engine-component production line in a Swedish company. Since in an unsupervised learning approach like K-Means clustering, the number of clusters is not well defined and depends on the user's intuition, the authors used the elbow method to determine the optimal number of clusters. Naskos et al. (2020) used acoustic sensor data and time-series discretization techniques for defect diagnosis in the cold forming press as manual inspections were not possible. Authors state that, unlike traditional classification problems, in PdM events are very rare and features are sparse. Therefore, to reduce the burden of labeling the sensor measurements with failure events, the authors used a sequence of artificial events to map time-series data. To this end, the authors used a data structure technique called Matrix Profile for motif and discord discovery. Danishvar, Angadi, and Mousavi (2020) stated that all [machine] learning models require large datasets (about the state of the machinery and potential breakdown) for training purposes but such data rarely exist in real-life situations. Therefore, the authors used 'event-clustering' and 'event-sequence prediction' techniques for what they called 'genomic of machine breakdown'. The paper used a regression-based event tracker to estimate the mean time to failure in an industrial case of a continuous compression molding machine that produces plastic bottle caps for the beverage industry.

More recently, Aivaliotis et al. (2021) stated that researchers have mainly used data-driven techniques for predictive maintenance of individual components used in industrial machinery, however, the intricate design of industrial machines and lack of historical data for data-driven predictive maintenance calls for the adoption of advanced concepts like simulation and digital twins. The Authors emphasized that simulation can help in estimating the future behavior of the industrial asset and thereafter presented a case study from the white-goods industry. The paper discussed the steps for physical modeling like model creation, selection of degradation curve, degradation data extraction, and finally remaining useful life estimation. Betti et al. (2021) presented a case study about predictive maintenance of components like penstocks, turbines, generators, and high voltage transformers in two hydro-power plants in Italy. Authors compared a 'self-organizing map' neural network with Hotelling t^2 index as a multivariate process control tool and found that the former outperformed the latter on accuracy and sensitivity. Researchers took advantage of the fact that in the given case most of the measured data corresponded to nominal behavior and very few data represented anomalous patterns. In Ayvaz and Alpay (2021), the authors presented an ML-based predictive maintenance solution for a consumer goods manufacturing plant located in Turkey. Researchers used the MQTT protocol so that the data generated by the sensors are collected in a database located in a private cloud. Since in predictive maintenance majority of the data that gets collected represents healthy operating conditions, therefore the issue of class distribution imbalance is common - resulting in ML classifiers' bias toward larger classes. To deal with the issue, the authors used data sampling strategies such as random under-sampling and over-sampling, syntactic over-sampling, and bagging and boosting. Authors conclude that models of Random Forest and XGBoost outperformed other techniques like gradient boosting, multilayer perceptron regressor, and support vector regression.

After reviewing recently published papers, we can say that the issue holding back wider applications of predictive maintenance is the lack of training data that is a prerequisite for supervised learning. To overcome this problem physical modeling and digital twin can be used to simulate the physical properties of machinery as we see in some papers, but the approach is extremely complex, very technical, time-consuming, and costly to carry out. In data-driven approaches, it is relatively easy to identify a defect if it occurs frequently, however, if a defect has never manifested in the past, there is no way we can diagnose that defect correctly in the future. This is a major limitation of supervised learning and maybe this is why most of the papers that we reviewed relied either on anomaly detection or clustering for predictive maintenance. Nevertheless, supervised learning allows for better planning for the future by assisting in maintenance schedule optimization and inventory optimization, and is therefore desirable. As we go from fault detection to identification to remaining useful life prediction and potentially towards prescriptive maintenance - the value to the business increases (Sikorska, Hodkiewicz, and Ma, 2011).

2.3 PdM for Conveyors

Conveyors are automated material handling systems and mostly consist of two or more pulleys with a closed-loop belt. The loop then moves weights from one point to another on the belt. Literature acknowledges that the digitization of conveyors can increase productivity and reduce downtime. For example, W. Li et al. (2013) developed an analytical tool based on failure data of a conveyor system at a distribution warehouse to determine the effectiveness of predictive maintenance with production planning strategies. They combined wavelet packet decomposition and SVM to implement a fault diagnosis system and designed an online monitoring system for the belt conveyor. Stefaniak, Wodecki, and Zimroz (2018) developed a decision support system for the maintenance of an underground copper ore mine by integrating data from different sources present in the network of conveyors. They carried out multivariate analyses and demonstrated how analytical models, reasoning criteria, and reporting tools can allow for constant monitoring of the network and the maintenance staff. Kiangala and Zenghui Wang (2018) introduced an experimental predictive maintenance framework for conveyor motors to detect impairments while minimizing incorrect fault diagnosis in a bottling plant. Later in Kiangala and Zenghui Wang (2020), authors also developed a machine learning model that classifies whether the abnormalities observed are production-threatening or not. The classification model was built using a combination of time-series imaging and a convolutional neural network.

Since belts are integral to conveyors, we found several papers dealing with the condition monitoring of belts. For example, Q. He and Zongqiang Wang (2021) studied conveyor belts and their relation to engine failures using phase-sensitive optical time-domain reflectometers. They proposed a 4th-order cumulant algorithm that can detect and analyze Gaussian-like vibration in belts to capture data with a better signal-to-noise ratio. Andrejiova, Grincova, and Marasova (2021) carried out experimental research on a belt conveyor system, where materials are dropped to the conveyor from a height. They found a correlation between the erosion in conveyor belts and the type of falling material and the height of the fall. They also classified damage by applying ML algorithms such as the Naive Bayes classifier, decision trees, and logistic regression. To diagnose and monitor the longitudinal tear in conveyor belts, Qu et al. (2021) proposed an adaptive deep convolutional network. The aim was to detect damages in the belt by extracting features from images captured at 20 frames per second. The authors found that the proposed technique was well suited to detect scratches and tearing of conveyor belt surfaces, and performed better than the traditional Support Vector Machine.

Our review of the literature indicates that there is little work that focuses on the condition monitoring of conveyors. Researchers mainly focused on detecting damages in belts in the conveyor systems. However, conveyors have numerous other components like motors, gearboxes, coupling, rollers, idlers, etc., that require condition monitoring. Though these components are individually studied in several papers, few take a holistic approach to assembly-level condition monitoring. Importantly, we were not able to find any papers that dealt with condition monitoring of conveyors that carry variable loads.

3 Solution Framework

ISO 13374-1:2003 and Open system architecture for condition-based maintenance (OSA-CBM) are the two commonly referred standards that have been adopted in many condition-based maintenance systems (Guillén et al., 2016). OSA-CBM consists of six generic functional blocks namely (1) Data Acquisition, (2) Data Manipulation, (3) State Detection, (4) Health Assessment, (5) Prognosis, and (6) Advisory Generation. In this study, we chose OSA-CBM as the base framework that accommodates state detection and machine learning as shown in Figure 4.

In step 1 in fig 4, the IoT sensors collect and send raw data to data servers over a wireless network wherein the data is stored for both long and short-term usage. Sending raw data directly to a cloud platform is problematic because an enormous volume of data would be captured from thousands of conveyors every day. J. Lin et al. (2017) states that because of bandwidth costs and network limitations, there is a need to incorporate edge analytics in predictive maintenance frameworks. The literature recommends edge computing because it enables local processing of the data instead of sending data directly to the cloud. The idea is to bring the intelligence and analytics as close to the source as possible and reduce the volume of data so that the

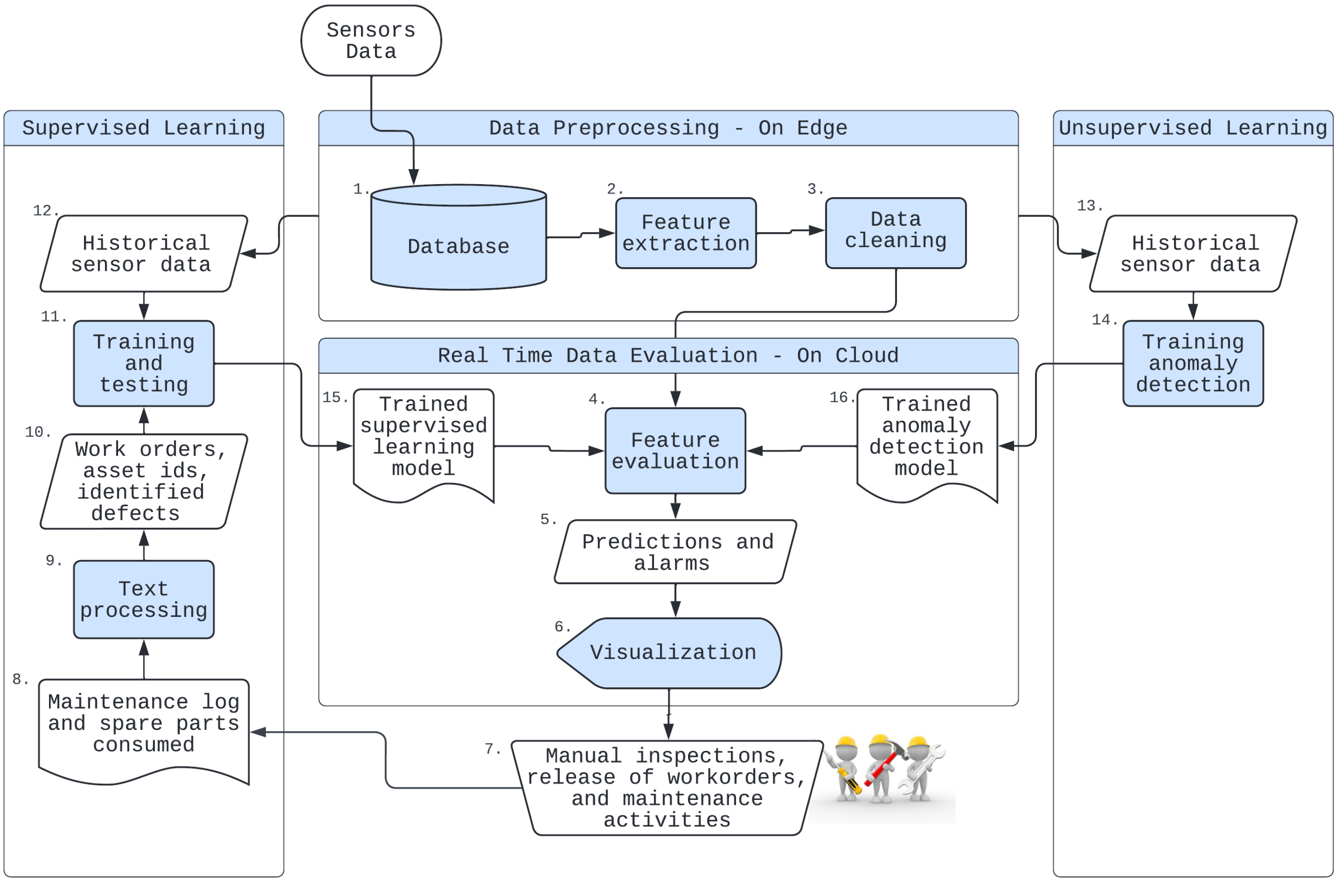

Figure 4: Maintenance 4.0 PdM Framework (adopted from OSA-CBM)

bandwidth, cloud processing, and cloud storage are reduced (Sittón-Candanedo et al., 2019). Therefore, steps 2 and 3 indicate that data collected in the data server are feature-engineered and cleaned before sending them to the cloud over the internet. Also, a copy of the cleaned data could be saved in the database for referencing and machine-learning purposes.

The next step (marked 4) in fig 4 relates to state detection and health assessment using trained ML models (marked 15 and 16). At this stage, the edge-processed real-time data is fed into the trained ML models that should detect anomalies and diagnose defects. The anomaly detection model should detect abnormalities in data to safeguard against unknown defects. Note, during the initial phase - just after the installation of the sensors, there won't be sufficient labeled data for supervised learning. At this stage, one may solely depend on anomaly detection that relates to unsupervised learning (see steps 13-14 & 16) and wait for defects to happen. We discuss this in detail in section 3.2.1.

Since for consistent performance and accurate predictions, it is essential to periodically update and retrain ML modules (Carvalho et al., 2019), therefore the framework in fig 4 includes periodic training of the supervised and unsupervised machine learning through steps shown on the left and right side of the framework. Steps 5 and 6 relate visualization of the cleaned data, alarms, predictions, and advisory generation through interactive dashboards created using applications like PowerBI or Tableau. Step 7 is about manual inspections (initiated after alarms are raised), the release of work orders (if any problem is identified after inspection), and maintenance activities. Following maintenance activities, maintenance and spare parts logs are updated (mostly by field supervisors).

An early and effective diagnosis of a defect can provide immense value by aiding in manpower scheduling and inventory optimization. Therefore, steps marked 8 to 12 relate to supervised learning that aims to diagnose potential defects and limit the need for manual inspections. Since it is not practical to manually read and label defects for supervised learning, in step 9, we incorporate text processing that should extract labels from the maintenance and spare parts logs automatically. The model trained in step 11 is then deployed in the cloud application that can evaluate the data from the sensor almost in a real-time.

The next subsections discuss how the aforementioned framework was implemented for the use case - i.e. for PdM in the baggage handling system. In subsection 3.1, steps carried out in the edge are discussed; this includes (a) data acquisition and aggregation in the edge, (b) feature extraction from the captured signals, and (c) data cleaning to filter out idle and noisy data. In subsection 3.2, steps carried out in the cloud are discussed; that includes (a) anomaly detection to catch any abnormalities, and (b) supervised learning for defect diagnosis.

3.1 Edge Processing

In this section, steps in the framework (Figure 4) that can be carried out near the source of the data (aka edge) are discussed. We have discussed (a) data acquisition, (b) data cleaning, and (c) feature extraction in greater detail.

3.1.1 Data Acquisition

In general, this step focuses on the selection and placement of sensors on the asset to be monitored. The optimal sensor placement is needed for better fault detection (Cao, Niu, and Z. He, 2012). The location that is least susceptible to noise must always be preferred. Another reason for sensor placement optimization could stem from constraints like a limited budget for the procurement of the sensors. For condition monitoring, there are a variety of sensors available - but vibration and acoustic sensors are commonly used. Out of the two, accelerometers are frequently used in bearings, gearboxes, motors, and turbines. Other types of sensors that may be considered include temperature sensors and vision-based sensors, however, adding more sensors could escalate the cost with diminishing return on investment.

Since components in baggage handling conveyors consist of primarily moving parts, we expect those to produce mainly vibrational and acoustic signatures. Based on our preliminary study, we concluded that (a) acoustic sensors are susceptible to noise from surrounding conveyors and may not capture weaker signals from sub-components, and (b) vibration needs structural contact and is relatively localized. A vibration sensor optimally placed in the middle of the assembly was found to be sufficient because of the relative compactness of the assemblies. Although we do not expect extremely high precision for all the components with just one sensor per assembly, because of budgetary constraints and relatively low-cost sub-components involved, placing more sensors per assembly was not economical. That said, in case one needs to add more sensors, features extracted from multiple sensors could easily be concatenated side by side without affecting the overall framework.

When one looks for accelerometers specifically, there are two types available in the market, (a) traditional wired accelerometers and (b) wireless IoT accelerometers. Both have their advantages and disadvantages. The traditional wired sensor such as 'Siemens-Siplus cms2000' can capture frequencies up to 15kHz at a range of up to 50 m/sec^2. However, installation of these sensors requires lengthy wiring where each sensor needs its wire directly connected to a high-end computer via PLC. These sensors were originally designed for condition monitoring of heavy and costly assets like wind turbines but are not suitable if the cost of the asset to be monitored is not very high. On the other side, the IoT sensors can capture frequencies up to 1 to 3kHz at a measuring range of 16 m/sec^2. The major difference between wired and wireless sensors is that the former transfers the data continuously, while the latter takes 1-2 minutes to transfer a second of captured data to the server - so may not send recordings continuously. It is stated that IoT sensors are most suitable for brown-field projects Strauss et al. (2018).

The selection of wired or wireless sensors should depend mainly on the use case and some preliminary study must be conducted to assess their suitability. Considering (a) the low cost of the components like bearings, belts, gearbox, motor, idlers, and couplings, (b) defects take substantial time to develop, and (c) the operating speed of the components ranges from 48 to 1200 revolutions per minute, the use of costly and highly sensitive accelerometer was avoided. IoT accelerometers finally used in the study captured 1000 samples in 0.625 seconds at every 15 minutes intervals. This gave visibility into frequencies from 0 to 800 Hz.

3.1.2 Feature Extraction

Broadly speaking, feature extraction involves the transformation of the raw vibration data to low-dimension feature vectors for machine learning. Feature extraction techniques are effective in reducing the volume of data, and therefore to minimize the volume of the data to be transferred to the cloud it is recommended that feature extraction is done on edge. In vibration-based condition monitoring, feature extraction can be done in the time domain, frequency domain, and frequency-time domain as in Caesarendra and Tjahjowidodo (2017). Time-domain feature extraction evaluates time-series parameters like root-mean-square (RMS), kurtosis, mean, variance, crest factor, and skewness, among others. Frequency domain feature extraction involves Fast Fourier Transformation of the raw accelerometer data. Time-frequency feature extraction can include Short-Time Fourier Transformation, Wavelet Transformation, and Wavelet Decomposition. Recently, there has been an increase in interest in automated feature extraction using deep learning techniques like CNN as in Khan and Yairi (2018) and Auto-encoders as in Thirukovalluru et al. (2016). Features may not necessarily be "handcrafted" and researchers do not have to manually extract them using either statistical measures or signal transformations.

In this use case, multiple "time domain" and "frequency domain" features are extracted from the vibrational data. The time domain features include peak-peak, RMS, kurtosis, mean, variance, crest factor, and skewness. For frequency domain features, first, time domain data is converted into the frequency domain using fast-Fourier transformation. The resulting power spectrum is split into bins of width 100Hz each, as in Thirukovalluru et al. (2016). For every bin, the means of the power spectral values are calculated (in g^2/Hz). Instead of calculating the means, one may choose the maximums as the features from the frequency bins (but this can make the system extremely sensitive to small changes). Although numerous other features can be extracted from the same data, the aforementioned 7 time-domain features and 8 frequency-domain features were finalized for the use case presented in this paper.

We advocate traditional approaches over sophisticated autoencoder or CNN for dimensionality reduction or feature extraction because of the availability of limited data at the initial phase of the project. Machinery ideally generates distinct signals depending upon the impending defects and the data from a newly installed sensor system can not represent the entire spectrum of defects that would occur in the future. Training autoencoders for feature extraction or dimensionality reduction on such data will not be valid for long and such models would need frequent training. Nevertheless, when there is a significant amount of data, deep learning models can be deployed on the edge for feature extraction and dimensionality reduction.

3.1.3 Data Cleaning

In many use cases, including the case in this paper, the major concern could be noise that gets captured along with the data. The major sources of erratic noise/vibration in conveyors in baggage handling systems are the movement of luggage and manual activities nearby. Therefore, to reduce such noise from getting captured, it is recommended to place a sensor on a rigid structure near the critical assembly that needs to be monitored. Specifically, while using IoT sensors - that record vibration intermittently and have limited computational power, data must be cleaned smartly to avoid erroneous alarms.

Conventionally, data cleansing is done using outlier detection techniques. While using those techniques it is assumed that the appearance of noisy data in the datasets is rare and that the quality of data can be assessed by comparing it with the majority of the observations in the same dataset. However, Y. Liu et al. (2020) states that the concept of outlier is not well-defined as it can include both noise - "which is meaningless and should be removed or corrected" and anomalies - "which are items of interest indicating an unhealthy condition".

In baggage handling systems, we observe that data captured by the sensors are predominantly noisy. When heavy luggage pass over conveyors, interactions between the moving and static parts create random vibrations. In figure 2 we present how the vibrations as recorded by the accelerometers vary because of the interaction with the luggage. One can see that in figure 2b the frequencies in the range 0-100Hz are abnormally high when compared to figure 2a - even though both the signals are only 40 minutes apart. This variation is likely due to mechanical disturbance from the luggage and can not be because of any defect. Since the weight of luggage varies widely, the vibrations that they induce are not consistent nor predictable.

We estimate approximately 70-80% of the captured data can be either idle or noisy and not useful. Because conveyors operate mainly when they have to carry luggage, there are fewer occasions when clean data are captured (like the one shown in Figure 2a). Therefore, any application of conventional outlier detection techniques to identify and remove noisy data is futile. Such noisy data if used for predictive maintenance can lead to false alarms and unnecessary inspections thereby reducing the reliability of the PdM system. The issue is aggravated because multiple IoT sensors just can not send high-volume data continuously to the database. To overcome these limitations, we propose a simple algorithm 1 with justification based on insights from histogram plotted from historical RMS data.

In figure 5, one can see that a clear separation between idle and non-idle data can be found - in this case - RMS from 0.2 to 0.4. This gap can allow for filtering out idle data. However, see that there is no clear separation between noisy and clean data (from RMS 0.4 to 1.2). Noisy data from the movement of luggage are not significantly different nor consistent as compared to clean data.

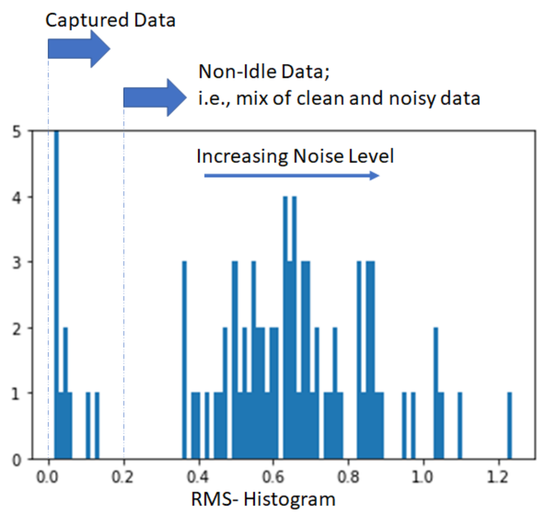

Figure 5: Histogram of root mean squared values of vibration captured during a period

A clear separation between idle and non-idle data allows the use of thresholding or clustering methods to clean useless-idle readings. Note clustering is suitable if there exist persistent background noises that vary with location (because sensors can pick surrounding noise even if the assets on which they are placed are idle). Additionally, to overcome the issue of identifying clean data from noisy data, we can rely on the understanding that operations with luggage create more absolute vibration as compared to operations without luggage. Ideally, the first quartile of the sample would be the cleanest while the last quartile would be the noisiest. Therefore, to clean the data captured in baggage handling systems, we came up with an algorithm that is based on these understandings (see Algorithm 1).

Algorithm 1: Dynamic data cleansing using RMS value

Input: RMS values of the signals captured by the sensors for time-period T
Output: Clean data for the time period T where $t \in T$; `// where` t `represents a day`
Step-1: Apply clustering to identify two groups using RMS as an input feature;
Step-2: Label the data belonging to the cluster with lower mean RMS as group 'Idle-Data';
Step-3: Label the data belonging to the cluster with a higher mean RMS as group 'NonIdle-Data';
for $t = 1 : T$ **do**
  Step-4: Fetch data belonging to group 'NonIdle-Data' from time-period t ;
  Step-5: Sort the selected data by their RMS values in ascending order;
  Step-6: Store the first $n\%$ of the sorted data in a database as 'Clean-Data' ;
end
Step-7: Use database 'Clean-Data' for anomaly detection and diagnosis ;

In step 6 of Algorithm 1 - $n\%$ is a parameter that needs to be tuned depending on the probability that sensors would capture clean data; n in the range of 20 to 30 can be a good starting point.

3.2 Cloud Analytics

In this section, we have discussed anomaly detection, data labeling, and supervised learning in more detail. The focus is on steps carried out on the cloud as shown in the framework in Figure 4.

3.2.1 Anomaly Detection

Novelty detection or anomaly detection can only be applied when there is enough historical data to compare with - such that the null hypothesis can be tested and the significance of the difference can be assessed. It is simple to apply anomaly detection if it is certain that the asset is in good working condition to start with. In those assets, any deviation from the initial working condition indicates deterioration. However, in brownfield projects, one can not guarantee that the assets are in their best operating condition even if they have been recently overhauled (Do et al., 2015). To solve such problems, researchers have recommended using a cohort of similar equipment and generating a meta-model that reflects the collective learning (Dalzochio et al., 2020). Therefore, to enable learning from a group of similar equipment, it was decided to place sensors identically on multiple similar conveyors. These datasets were then combined to train anomaly detection models.

There are several anomaly detection methods to choose from, for example, distance-based approaches like Mahalanobis distance, density-based approaches like LOF, ensemble-based techniques like isolation forest, and one-class clustering using SVM. In this study, anomaly detection was carried out using the Mahalanobis distance and LOF. The advantage that LOF has over Mahalanobis distance is that the former learns to identify clusters in the training set and would outperform the latter for datasets containing two or more natural clusters. Mahalanobis distance, in comparison, is computationally less challenging.

Mahalanobis distance is a statistical distance metric that calculates the similarity between data points in multivariate data sets. The similarity computation takes into account any correlation between variables. Additionally, one would require a threshold to determine the upper limit of the distance to identify unhealthy working conditions. The threshold could be defined at a value that corresponds to a low probability in the probability distribution that best fits the Mahalanobis distances computed during the training phase. Or simply the threshold can be defined as the 90% of the largest Mahalanobis distances (Entezami, Shariatmadar, and Mariani, 2020).

LOF (Breunig et al., 2000) is loosely related to density-based clustering and computes the local density for a data point vis-a-vis its neighbors. The density is local in the sense that the number of neighbors to consider for the computation of the density is limited and is user-defined. The locality is defined using k-nearest neighbors. The score simply indicates how isolated the object is as compared to objects in its neighborhood. A substantially lower density than neighbors indicates novelty and could be a cause of concern in condition monitoring.

3.2.2 Supervised Learning

Anomaly detection as shown in the previous section has two limitations: (a) components with higher anomaly score does not mean the asset is at the end of its life and does not necessarily warrant part replacements or corrective actions, and (b) it does not help in defect diagnosis or spare part inventory optimization or maintenance scheduling. To move further in the predictive maintenance paradigm one must identify faulty parts much in advance so that the inventory and maintenance activities can be optimized. Therefore, supervised learning techniques like logistic regression, support vector machine, and random forest classifiers are widely used.

Since supervised learning requires labeled data, one may rely on maintenance and spare-parts consumption logs, which are readily available in operations and maintenance departments. However, these logs are not well structured and often stored in Excel sheets. Please see table 1 for examples of such entries. To retrieve class information from semi-structured data as shown, one must go through these entries manually. That would be an easy task if only a few months of data for a few conveyors has to be analyzed. However, such an approach is not scalable for hundreds of conveyors.

Table 1: Samples from a maintenance log

Finish Date	Work Orders	Work Type	Asset Id	Description
2019-09-25 21:17:00	17948132	CM	xxx.16.1-abc	CM:Maint:PM: shaft tensioning worn out
2020-01-09 13:30:00	18781373	CM	xxx.24.1-abc	CM:Maint:PM: pallet slat wear tear
2020-02-28 09:41:00	19194785	BD	xxx.36.1-abc	BD:OPS: torque arm mounting bracket bolts worn
2020-03-04 20:51:00	19231735	CM	xxx.36.1-bcd	CM:Maint:PM: drive shaft bearing warn out due to again

To scale up the solution for numerous conveyors, and for periodic training of the supervised learning models, the sensor-generated data must be labeled automatically. To extract labels, we propose that the techniques from the domain of text processing can be applied. Figure 6 shows steps that can be used for the same. The idea is drawn from the literature, e.g. see Medina-Oliva et al. (2014) and Nuñez and Borsato (2018) where authors have proposed using ontology to capitalize knowledge and ease diagnostic activities for intelligent health monitoring. Ontology can help in building a common understanding among people and software agents by the standardization of terms and concepts. It makes domain assumptions explicit; for example, if the maintenance log states "engine has an internal electrical degradation, the ontology can induce that it is an electrical engine" (Medina-Oliva et al., 2014). The approach also draws inspiration from ontology learning - an area dedicated to automated ontology construction and knowledge extraction (Khadir, Aliane, and Guessoum, 2021). In this paper, however, we simply aimed at extracting knowledge (aka defect labels) from the text documents - automatically or semi-automatically.

As shown in figure 6, the process starts by extracting relevant data from maintenance and spare parts logs. Since maintenance logs do contain the date and time of the activity and the identity of the conveyor attended, these data are retrieved directly from respective columns. The manually entered texts about the defect and the actions taken are extracted from the "Description" column. The text thus obtained goes through tokenization, stop word removal, spelling correction, and morphological analysis like stemming. Tokenization is a process by which the text is broken down into smaller units, usually words. During stop word removal, common and undesirable words like articles and pronouns are eliminated from further processing. The spelling correction is important because the chances of erroneous entries are high given the maintenance personnel may not be native English speakers. Finally, morphological analysis like stemming or lemmatization is done to bring the words to their common root words.

For ontological reasoning, taxonomic hierarchies can be created for conveyors, assemblies, components, and defects classes. The classes are associated with the help of properties (Medina-Oliva et al., 2014). For example, conveyor-type - hasAssembly - assembly-type can be used to define what kind of assemblies a conveyor type has. Such structural information help in matching the defects with the sensor data and limit the number of labels for diagnosis. As an example, entries like "bearing in the tensioning unit warn" and "tension bearing is worn out" are essentially the same and must be labeled so for machine learning purposes. Through tokenization, stop word removal, spelling correction, and morphological analysis - we can reduce the texts to "bearing tension worn" and "tension bearing worn" respectively, which then can be easily arranged into assembly-type - hasComponent - component-type - hasDefect - defect-type format reducing ambiguity and improving diagnostic accuracy.

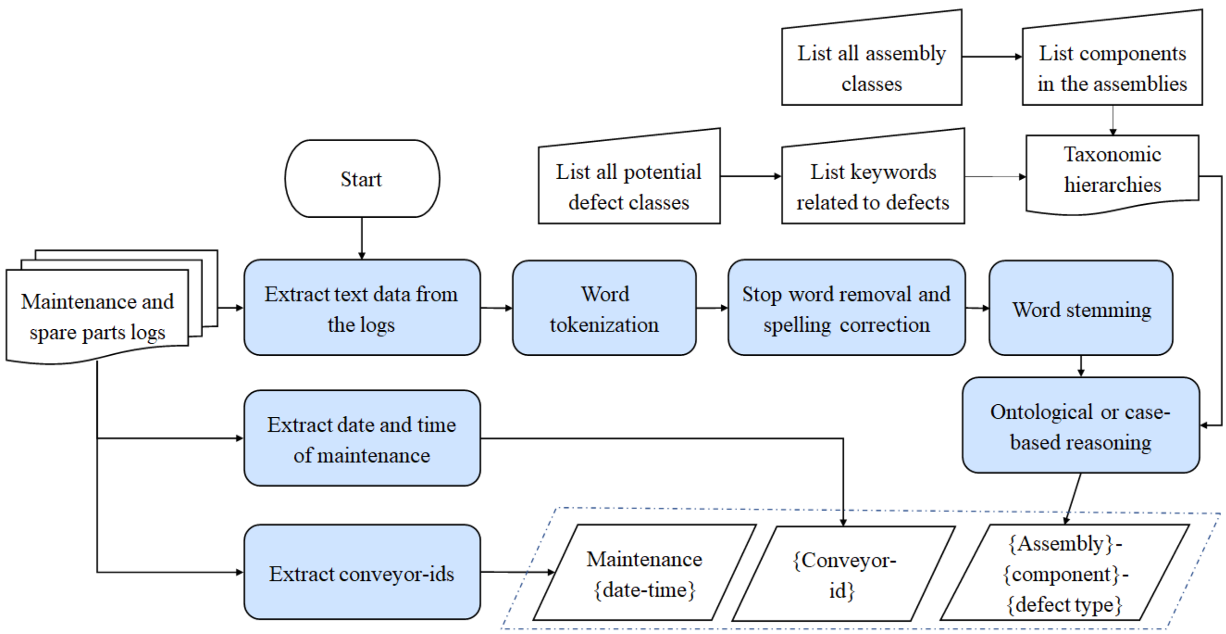

Figure 6: Processing of the text in the maintenance log for automated labeling of data

4 Results and Discussion

In extension to the solution framework, in this section, condition monitoring for a type of conveyor known as S-Lifts is presented. S-Lifts are tray-carrying conveyors that lift loads up or down within baggage handling systems (see figure 7). These conveyors were chosen because they are critical assets and bottlenecks in the system. A simple fault may take days to correct resulting in serious losses. The approximate placements of the sensors are marked in ×. The sensors were programmed to capture and send the data to receivers every half an hour. In total eight identical conveyors were selected for experimental studies (making a cohort). Most of the conveyors were operational for at least 10 years and were not in the best shape. A total of 72 sensors were installed (nine sensor per conveyor). All the data thus captured were stored in edge databases. We had a development computer for data analysis and coding. The computer had access to the edge databases via a virtual network. The data for the experiments were extracted from the databases directly and processed as discussed in sections 3.1.2 and 3.1.3. Feature extraction and data cleaning algorithms were developed and deployed to the edge as shown. These algorithms processed recently captured raw data and sent clean features to the cloud. Anomaly detection and diagnosis models developed were containerized and deployed to the cloud. The containers on the cloud processed cleaned features from the edge and generated reports which were then presented in the dashboard.

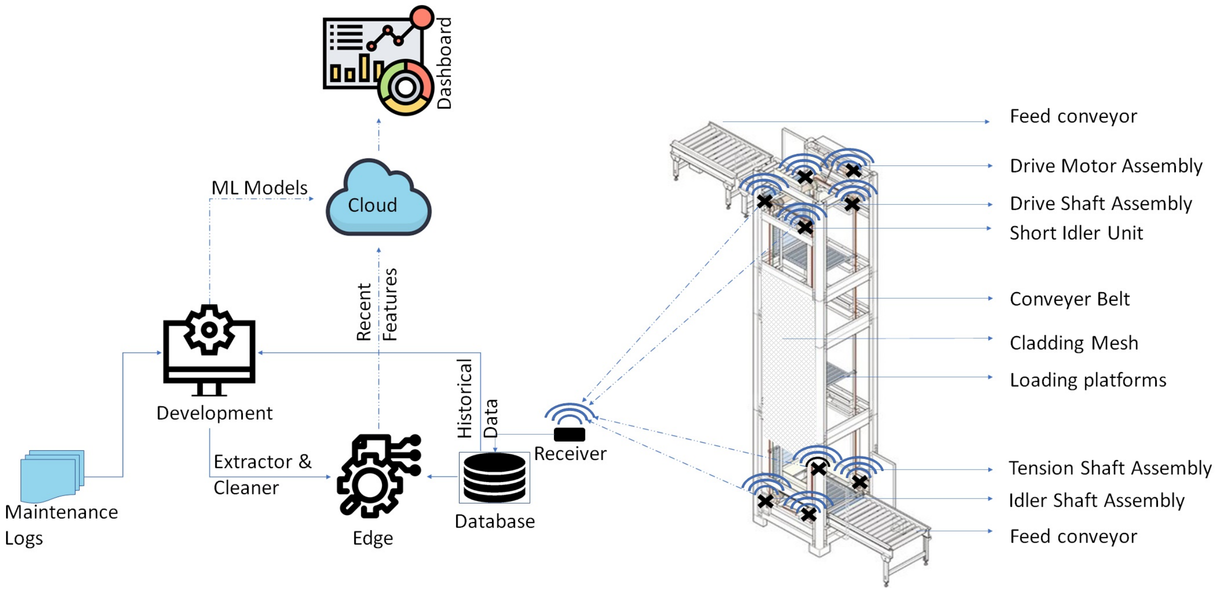

Figure 7: Continuous vertical conveyors (S-Lifts) and experimental setup

The results presented in this section are based on two years of collective data. Following the framework, anomaly detection models and supervised learning models were trained. However, during the initial phase, there were no labels for supervised learning, and alarms were raised primarily using anomaly scores. Table 2 shows a snippet of results from the LOF and Mahalanobis distance analyses. A higher average Mahalanobis distance or higher percent of anomalies as computed using the LOF in the recent data as compared to historical data from the same assembly indicates a problem in the assembly. As shown, assets and corresponding assemblies with a higher percentage of anomalous also had higher Mahalanobis distances. This indicates both these techniques agree and either of them can be used to raise alarms in S-Lifts (as discussed in our framework 4).

To raise alarms, LOF was finally used for its ability to accommodate multiple clusters in the data - considering the three-axis of freedom in vibration signifies three clusters by default. For every assembly, we had a separate anomaly detection model. The models for every assembly were trained every week using a sample of 1000 clean historical data from the cohort. Samples of the 20 most recent clean readings were retrieved every day from every assembly - to check for anomalies in the assembly. Alarms were raised when 75% or more readings were detected as anomalies. Once alarms were raised, the assemblies were inspected and findings were recorded by technicians. In case any component was found to be in poor working condition, corrective actions were taken. These corrective actions included cleaning, lubrication/greasing, fastener adjustments, re-alignments, and sometimes parts replacements. The parameters of the algorithms were tuned when there were too many or too few alarms. That is, we went back and forth on data cleaning and anomaly detection steps till alarms were mostly accurate.

Table 2: Sample output from anomaly detection using one class clustering and Mahalanobis distance

Asset Id	Assembly	LOF - Anomalies	Avg. Mahala Dist.
xxx.24.1-abc	Drive Assembly	78%	56
xxx.11.1-acb	Drive Assembly	46%	54
xxx.36.1-aed	Drive Assembly	59%	46
xxx.01.2-bcd	Drive Assembly	32%	19
xxx.01.1-acd	Drive Assembly	12%	16

To label the data automatically and to train supervised learning models periodically, maintenance and spare parts logs were pulled from the operations and maintenance database. A taxonomy was created in consultation with the maintenance team. Following figure 6, the texts in the logs were processed to extract labels. Conveyor-id and date-time fields were also retrieved. The collective information was then used to identify the assemblies and the components that were repaired. The anomaly scores from the previous step were used in case there were ambiguities in identifying the sensor which would have captured the defect. To collect the sensor data, the edge database was queried by date-time as mentioned in the maintenance log. The data for the previous five days from the date and time of maintenance were extracted. Finally, everything was structured as shown in table 3.

Apart from diagnosing defects, we wanted to diagnose if there is no defect. For that, a new label was introduced named 'No Issue Detected'. Since more than 99% of the data collected by the sensors represented 'no defect' condition, data for predictive maintenance are by default highly imbalanced. Using imbalanced data for training and testing a supervised learning technique would give misleading high prediction accuracy. To overcome the limitation of imbalance in data - because of predominantly no defect condition - and to collect data for the label 'No Issue Detected', we retrieved a sample of historical data that were having low anomaly scores.

Table 3: Labeled data for supervised learning from drive motor assembly

S.No	RMS	Kurt	...	Mean-700Hz	Mean-800Hz	Defect Class
1	0.30	0.50	...	0.07	0.01	Drive - Bracket - Damag
2	0.31	0.55	...	0.08	0.04	Drive - Bracket - Damag
3	0.5	0.03	...	0.11	0.21	Drive - Gearbox - Noise
...						

Supervised learning was carried out for individual assemblies - namely drive motor assembly, drive shaft assembly, tension shaft assembly, and idler shaft assembly, independently. Since defects take years to appear, the data available may not represent all the scenarios possible in the future, we decided to forgo feature selection. Commonly used classifiers namely logistics regression, multi-layer perceptron, support vector machine, and random forest were evaluated. For training and testing purposes, Scikit-learn's Sklearn Python package was used. The machine learning model parameters were kept at their default values and therefore no parameter tuning was done. Table 4 shows the performance of the classification techniques on the test data. It indicates that the random forest classifier outperforms the other classifiers. As shown in figure 8, the cross-validation scores in the learning curve for defects near motor-gearbox assembly indicate that the models were adequately trained.

Table 4: Comparison of classification techniques for fault identification near motor-gearbox assembly

ML Classifier	Precision	Recall	F1-Score
Logistics Regression	0.63	0.57	0.58
Multi-layer Perceptron	0.73	0.64	0.66
Support Vector Machine	0.79	0.76	0.78
Random Forest	0.86	0.86	0.86

Figure 8: Learning Curve for fault identification near motor-gearbox assembly using Random Forest

Ultimately, for defect diagnosis, the random forest classifier was employed because of better overall results. The confusion matrix in figure 9 shows the defects as observed (see 'True Label' - based on maintenance log) and classifications as done by a trained random forest classifier (see 'Predicted Label'). The figure indicates that most of the labels in the test data were identified correctly, except in the case when defects were related to the gearbox. The misclassification for 'Gearbox - Noise' and 'Gearbox - Wear' can be attributed to (a) possible similarity of the vibrational signal coming out of gearboxes or (b) inaccurate entries in the maintenance log and feedback from technicians. While the misclassification due to similarity in the vibrational signal can potentially be solved using feature engineering and deep learning algorithms, it is equally important to acquire quality feedback from the technicians. We observe that entering accurate information in maintenance and inspection logs may not be given high importance. It is important to incentivize and motivate teams on the site to log correct and accurate data. We also recommended updating the format of the maintenance log so that technicians can specify assembly, component, and defect identified separately so that we don't have to rely on text processing heavily.

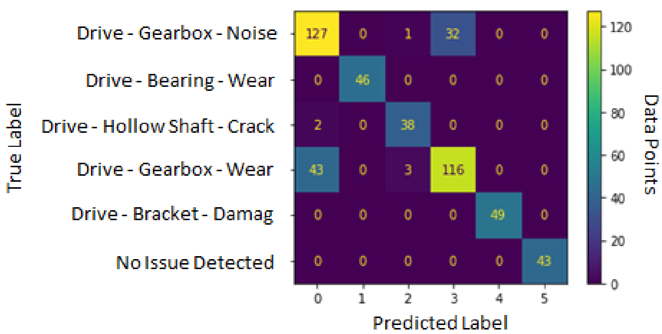

Figure 9: Confusion matrix for defects in motor assemblies using Random Forest

Visualization: Figure 10 shows how we presented the results of anomaly detection and defect diagnosis. We listed all the assemblies and all the conveyors in tabular form. The rows were assorted by conveyors. One could also scroll through the table and sort the table by columns. The table contained conveyor IDs, the name of the assembly, anomaly scores, defect diagnosis results, and priority. When anomaly scores (in %) were higher than 0.75, the priorities were set as one and corrective actions were requested. When anomaly scores were higher than 0.5, the priorities were set as two, and requests for inspection were raised. For anomaly scores less than 0.5, no actions were requested. The rows were highlighted in red and orange for priorities one and two, respectively. Hovering over a row shows the trend in anomaly scores over the last month.

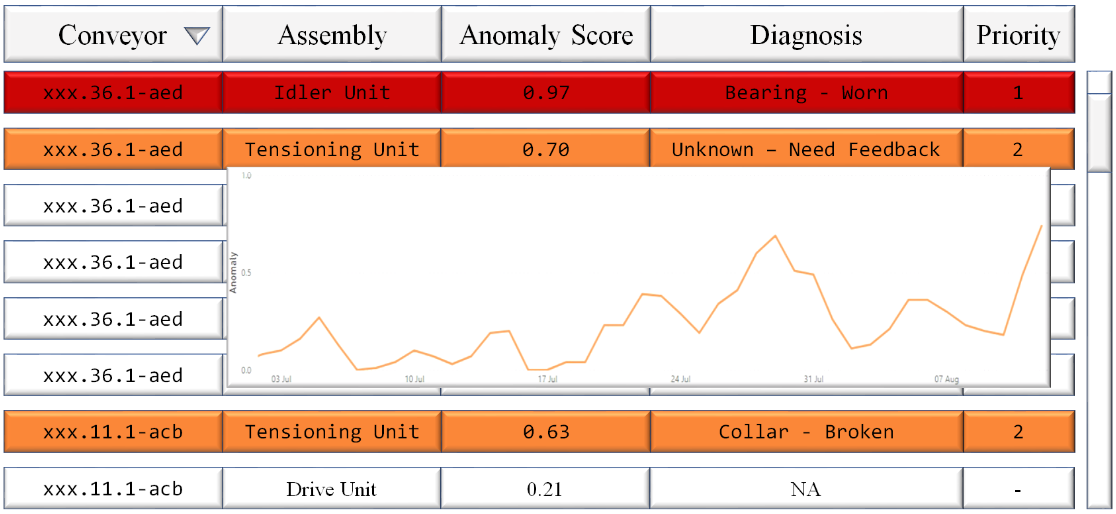

Figure 10: Dashboarding - Anomaly scores and diagnosis reports

Another dashboard (figure 11) was developed in case we wanted to validate the alarms and observe the trend in the raw features like RMS and frequencies. See how features like RMS and mean values in frequency bins changed over time for one of the drive motor assemblies. Initially, the mean power density of 200-300Hz was high. By the end of September 2019, there was corrective maintenance. Although the mean power density of 200-300Hz was subdued, the mean power density of 0-100Hz increased. Another corrective action was done at the start of November 2019. However, RMS and the mean power density of 600-700Hz became abnormally high which was then corrected by end of November 2019. Multiple changes in the signal captured during the period can potentially be attributed to incomplete maintenance (Do et al., 2015) and multiple simultaneous issues related to bearings and belts. Although after November 2019 the frequencies were mostly in control, RMS was still high. It was found that the motor bracket was damaged. By the mid of March, we can see that all the features of the signal got within a range that represented no defect condition.

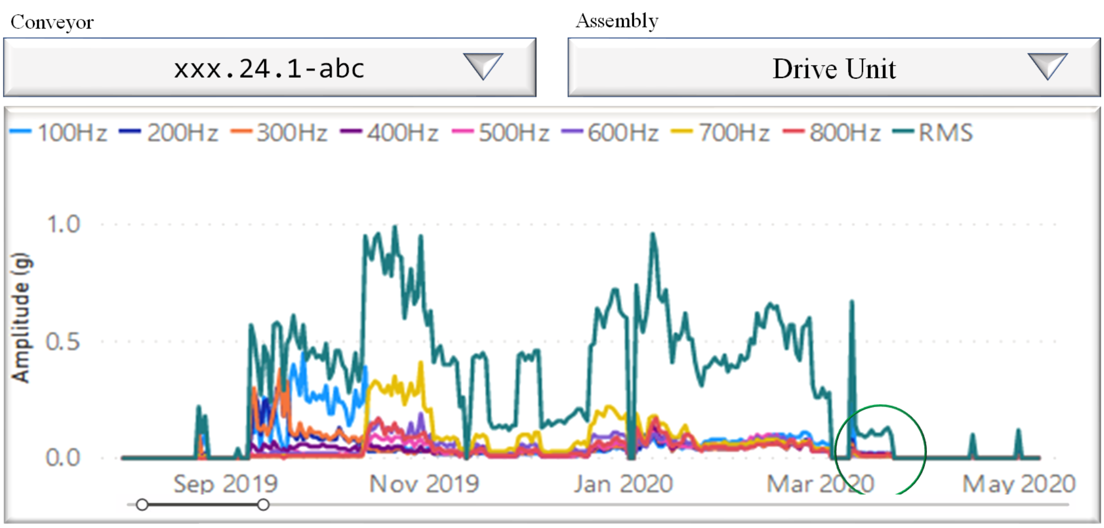

Figure 11: Visualization - Feature vectors for validation

4.1 Limitations

A limitation of the study is that the steps in Algorithm 1 are based mainly on RMS values (as a proxy for absolute vibration). Though the use of RMS was a simple and effective idea, experiments can be conducted to see if other features could also be used to clean the data. Our attempt to use one-class classification on multi-variate data gave inconsistent results (clean data were getting identified as outliers). If it is required to use other features, we suppose semi-supervised outlier detection can also be used to replace steps 5 and 6 in the algorithm.

We focused mainly on diagnosing individual defects because of the limited time that we had to collect labeled data. As in the case discussed in the previous paragraph, there can be any combination of defects in these assemblies. For example, the bearing and the motor bracket (two different components) can simultaneously be damaged. An incomplete diagnosis can result in multiple corrective actions on the same assembly within a small period. It is desirable to identify all defects simultaneously when there are multiple issues in an assembly. Therefore, over time it will be required to develop machine learning models that would identify combinations of defects instead of individual defects. As the number of classes for defect diagnosis will increase, using multi-class classification is recommended.

For training and testing, we ignored the conveyor-id from which the data were collected, however, as an alternative one may also separate those conveyors into training and test conveyors, such that training data is collected from one set of conveyors, and testing data is collected from another set of conveyors. The ML models developed for a type of conveyor can not be applied to other types of conveyors. For the same reason, the models developed for S-Lifts in this study can not be applied to other S-Lifts if they have different operational speeds/specifications or if they are sourced from different manufacturers.

4.2 Future Scope

In the predictive maintenance paradigm - diagnosis and prognosis are complementary tools that support operational decisions like how many spare parts to hold in the inventory or how to schedule human resources (Diez-Olivan et al., 2019). Fault diagnosis simply means detecting and identifying the problem. Prognostics, on the other hand, try to predict how the fault will develop and when the asset would fail. Prognosis is more complex as it takes into account various developmental stages of a fault and as with other predictive models not 100% accurate (Jardine, D. Lin, and Banjevic, 2006). Moreover, moving from diagnosis to prognosis is a difficult task primarily because of the unpredictability of faults in industrial settings.

Because remaining useful life prediction techniques are mostly extrapolation based, these techniques can work only if the chosen health indicator has a monotonic increase or decrease in trend (Lei et al., 2018). However, as shown in figure 11, there can be frequent changes in the data captured by the sensors creating a major roadblock toward defect prognosis. These changes can be attributed to imperfect maintenance (Do et al., 2015) and occlusions caused by loose baggage items like straps, handles, buckles, locks, etc. Run-to-failure data is unattainable because the cost of breakdown is too high. Furthermore, these conveyors are not operated continuously or equally because their usage depends on the number of flights landing at the airport and the availability of redundant or bypass conveyor lines. Because of the reasons just mentioned, these machines rarely follow their normal degradation making defect prognosis a challenging task. Therefore, we leave that for the future.

As stated in the limitations section, ML models developed for a type of conveyor can not be applied to other types of conveyors. Generalizing the model for other machines is a difficult task, however, the concepts of ‘transfer learning’ can be applied in the future enabling faster training on similar but different machines (given some amount of labeled data can be collected from those non-identical machines).

5 Conclusion

Airports around the world use conveyors extensively in their baggage handling systems, but thus far predominantly use periodic maintenance policies even though that has proven to be a costly and not a fail-safe approach. Arguably, predictive maintenance is not as common in conveyors as in other industrial assets like turbines, pumps, motors, and so on. Failure in any one of the conveyors can choke hundreds of upstream and downstream conveyors directly affecting the throughput of the baggage handling system. Considering unnecessary corrective actions and part replacement that relates to time-based preventive maintenance, and eliminating unplanned breakdowns while reducing the dependency on manual inspections, the objective of this research was to design a maintenance 4.0 solution applicable to baggage handling systems for airports. Importantly, this paper draws researchers' attention to a relatively unexplored topic where condition monitoring is done in a non-white noisy environment - as in conveyors carrying luggage in airports.

Maintenance 4.0 calls for cloud- and IoT- enabled condition monitoring and AI-based predictive maintenance, however, we found that it was not straightforward to deploy such solutions to baggage handling systems. A major issue in the predictive maintenance of conveyors particularly for airports is that the data captured by the sensors contain noise from the movement of luggage. Since no luggage is the same, the noise generated by their movement is not consistent. Outlier detection techniques gave inconsistent results because the majority of the data that got captured were noisy. Therefore, we developed an algorithm that can help clean such data.

Even though the cost of storing the data in the cloud is declining over the years, if the volume of data in the cloud is high, the cost of storage can lead to a significant increase in the cost of the entire predictive maintenance system. Therefore, to keep the solution scalable, we incorporated decentralized analytics whereby data cleaning and feature extraction were carried out on the edge. Furthermore, we used cohorts of similar equipment and generated meta-models that reflected the collective learning from multiple conveyors. Integrated machine learning approaches allowed anomaly detection and defect diagnosis in parallel enabling early adoption of the solution for real-life applications. Since intelligent fault diagnosis requires automated data labeling and training of Supervised ML models, we presented how text processing with ontological reasoning can extract labels from unstructured data such as inspection reports and maintenance logs. Results obtained from anomaly detection and supervised learning algorithms were presented, and finally, limitations and scope for future research were discussed.

References

Ahmad, Rosmaini and Shahrul Kamaruddin (Aug. 2012). "An overview of time-based and condition-based maintenance in industrial application". In: *Computers & Industrial Engineering* 63.1, pp. 135–149. ISSN: 03608352. DOI: 10.1016/j.cie.2012.02.002. URL: http://dx.doi.org/10.1016/j.cie.2012.02.002%20https://linkinghub.elsevier.com/retrieve/pii/S0360835212000484.

"Air transport industry insights, The Baggage Report" (2018). In: *SITA* 2.of 97. URL: https://www.sita.aero/resources/type/surveys-reports/baggage-report-2018.

Aivaliotis, P. et al. (Oct. 2021). "Degradation curves integration in physics-based models: Towards the predictive maintenance of industrial robots". In: *Robotics and Computer-Integrated Manufacturing* 71, p. 102177. ISSN: 07365845. DOI: 10.1016/j.rcim.2021.102177. URL: https://doi.org/10.1016/j.rcim.2021.102177%20https://linkinghub.elsevier.com/retrieve/pii/S0736584521000600.

Alsyouf, I., F. Humaid, and S. Al Kamali (Dec. 2014). “Mishandled baggage problem: Causes and improvement suggestions”. In: *2014 IEEE International Conference on Industrial Engineering and Engineering Management*. Vol. 2015-Janua. IEEE, pp. 154–158. ISBN: 978-1-4799-6410-9. DOI: 10.1109/IEEM.2014.7058619. URL: http://ieeexplore.ieee.org/document/7058619/.

Andrejiova, Miriam, Anna Grincova, and Daniela Marasova (Sept. 2021). “Identification with machine learning techniques of a classification model for the degree of damage to rubber-textile conveyor belts with the aim to achieve sustainability”. In: *Engineering Failure Analysis* 127.February, p. 105564. ISSN: 13506307. DOI: 10.1016/j.engfailanal.2021.105564. URL: https://doi.org/10.1016/j.engfailanal.2021.105564%20https://linkinghub.elsevier.com/retrieve/pii/S1350630721004258.

Ayvaz, Serkan and Koray Alpay (July 2021). “Predictive maintenance system for production lines in manufacturing: A machine learning approach using IoT data in real-time”. In: *Expert Systems with Applications* 173.November 2020, p. 114598. ISSN: 09574174. DOI: 10.1016/j.eswa.2021.114598. URL: https://doi.org/10.1016/j.eswa.2021.114598%20https://linkinghub.elsevier.com/retrieve/pii/S0957417421000397.

Bekar, Ebru Turanoglu, Per Nyqvist, and Anders Skoogh (May 2020). “An intelligent approach for data pre-processing and analysis in predictive maintenance with an industrial case study”. In: *Advances in Mechanical Engineering* 12.5, p. 168781402091920. ISSN: 1687-8140. DOI: 10.1177/1687814020919207. URL: http://journals.sagepub.com/doi/10.1177/1687814020919207.

Betti, Alessandro et al. (June 2021). “Condition monitoring and predictive maintenance methodologies for hydropower plants equipment”. In: *Renewable Energy* 171, pp. 246–253. ISSN: 09601481. DOI: 10.1016/j.renene.2021.02.102. URL: https://linkinghub.elsevier.com/retrieve/pii/S0960148121002810.

Breunig, Markus M et al. (2000). “LOF: identifying density-based local outliers”. In: *Proceedings of the 2000 ACM SIGMOD international conference on Management of data*, pp. 93–104.

Caesarendra, Wahyu and Tegoeh Tjahjowidodo (Sept. 2017). “A Review of Feature Extraction Methods in Vibration-Based Condition Monitoring and Its Application for Degradation Trend Estimation of Low-Speed Slew Bearing”. In: *Machines* 5.4, p. 21. ISSN: 2075-1702. DOI: 10.3390/machines5040021. URL: http://www.mdpi.com/2075-1702/5/4/21.

Cao, Hongrui, Linkai Niu, and Zhengjia He (June 2012). “Method for Vibration Response Simulation and Sensor Placement Optimization of a Machine Tool Spindle System with a Bearing Defect”. In: *Sensors* 12.7, pp. 8732–8754. ISSN: 1424-8220. DOI: 10.3390/s120708732. URL: http://www.mdpi.com/1424-8220/12/7/8732.

Carvalho, Thyago P. et al. (Nov. 2019). “A systematic literature review of machine learning methods applied to predictive maintenance”. In: *Computers & Industrial Engineering* 137.August, p. 106024. ISSN: 03608352. DOI: 10.1016/j.cie.2019.106024. URL: https://doi.org/10.1016/j.cie.2019.106024%20https://linkinghub.elsevier.com/retrieve/pii/S0360835219304838.

Dalzochio, Jovani et al. (Dec. 2020). “Machine learning and reasoning for predictive maintenance in Industry 4.0: Current status and challenges”. In: *Computers in Industry* 123, p. 103298. ISSN: 01663615. DOI: 10.1016/j.compind.2020.103298. URL: https://doi.org/10.1016/j.compind.2020.103298%20https://linkinghub.elsevier.com/retrieve/pii/S0166361520305327.

Danishvar, Morad, Veerendra C. Angadi, and Alireza Mousavi (Aug. 2020). “A PdM Framework Through the Event-based Genomics of Machine Breakdown”. In: *2020 Asia-Pacific International Symposium on Advanced Reliability and Maintenance Modeling (APARM)*. IEEE, pp. 1–6. ISBN: 978-1-7281-7102-9. DOI: 10.1109/APARM49247.2020.9209530. URL: https://ieeexplore.ieee.org/document/9209530/.

Diez-Olivan, Alberto et al. (Oct. 2019). “Data fusion and machine learning for industrial prognosis: Trends and perspectives towards Industry 4.0”. In: *Information Fusion* 50.September 2018, pp. 92–111. ISSN: 15662535. DOI: 10.1016/j.inffus.2018.10.005. URL: https://doi.org/10.1016/j.inffus.2018.10.005%20https://linkinghub.elsevier.com/retrieve/pii/S1566253518304706.

Do, Phuc et al. (Jan. 2015). “A proactive condition-based maintenance strategy with both perfect and imperfect maintenance actions”. In: *Reliability Engineering & System Safety* 133, pp. 22–32. ISSN: 09518320. DOI: 10.1016/j.ress.2014.08.011. URL: http://dx.doi.org/10.1016/j.ress.2014.08.011%20https://linkinghub.elsevier.com/retrieve/pii/S095183201400204X.

Entezami, Alireza, Hashem Shariatmadar, and Stefano Mariani (Dec. 2020). “Early damage assessment in large-scale structures by innovative statistical pattern recognition methods based on time series modeling and novelty detection”. In: *Advances in Engineering Software* 150.September, p. 102923. ISSN: 09659978. DOI: 10.1016/j.advengsoft.2020.102923. URL: https://doi.org/10.1016/j.advengsoft.2020.102923%20https://linkinghub.elsevier.com/retrieve/pii/S0965997820309698.

Erfani, Sarah M. et al. (Oct. 2016). “High-dimensional and large-scale anomaly detection using a linear one-class SVM with deep learning”. In: *Pattern Recognition* 58, pp. 121–134. ISSN: 00313203. DOI: 10.1016/j.patcog.2016.03.028. URL: http://dx.doi.org/10.1016/j.patcog.2016.03.028%20https://linkinghub.elsevier.com/retrieve/pii/S0031320316300267.

Fernandes, Marta et al. (2020). “Fault Detection Mechanism of a Predictive Maintenance System Based on Autoregressive Integrated Moving Average Models”. In: *Advances in Intelligent Systems and Computing*. Vol. 1003. Springer International Publishing, pp. 171–180. ISBN: 9783030238865. DOI: 10.1007/978-3-030-23887-2_20. URL: http://dx.doi.org/10.1007/978-3-030-23887-2%7B%5C_%7D20%20http://link.springer.com/10.1007/978-3-030-23887-2%7B%5C_%7D20.

Guillén, Antonio J. et al. (Oct. 2016). “A framework for effective management of condition based maintenance programs in the context of industrial development of E-Maintenance strategies”. In: *Computers in Industry* 82, pp. 170–185. ISSN: 01663615. DOI: 10.1016/j.compind.2016.07.003. URL: https://linkinghub.elsevier.com/retrieve/pii/S0166361516301178.

He, Qian and Zongqiang Wang (July 2021). “Detection of Gaussian-like vibration in fiber optical distributed sensing by 4th-order cumulants”. In: *2021 9th International Symposium on Next Generation Electronics (ISNE)*. 1. IEEE, pp. 1–4. ISBN: 978-1-7281-6258-4. DOI: 10.1109/ISNE48910.2021.9493608. URL: https://ieeexplore.ieee.org/document/9493608/.

Jardine, Andrew K.S., Daming Lin, and Dragan Banjevic (Oct. 2006). “A review on machinery diagnostics and prognostics implementing condition-based maintenance”. In: *Mechanical Systems and Signal Processing* 20.7, pp. 1483–1510. ISSN: 08883270. DOI: 10.1016/j.ymssp.2005.09.012. URL: https://linkinghub.elsevier.com/retrieve/pii/S0888327005001512.

Khadir, Ahlem Chérifa, Hassina Aliane, and Ahmed Guessoum (Feb. 2021). "Ontology learning: Grand tour and challenges". In: *Computer Science Review* 39, p. 100339. ISSN: 15740137. DOI: 10.1016/j.cosrev.2020.100339. URL: https://doi.org/10.1016/j.cosrev.2020.100339%20https://linkinghub.elsevier.com/retrieve/pii/S1574013720304391.

Khan, Samir and Takehisa Yairi (July 2018). "A review on the application of deep learning in system health management". In: *Mechanical Systems and Signal Processing* 107, pp. 241–265. ISSN: 08883270. DOI: 10.1016/j.ymssp.2017.11.024. URL: https://doi.org/10.1016/j.ymssp.2017.11.024%20https://linkinghub.elsevier.com/retrieve/pii/S0888327017306064.

Kiangala, Kahiomba Sonia and Zenghui Wang (Aug. 2018). "Initiating predictive maintenance for a conveyor motor in a bottling plant using industry 4.0 concepts". In: *The International Journal of Advanced Manufacturing Technology* 97.9-12, pp. 3251–3271. ISSN: 0268-3768. DOI: 10.1007/s00170-018-2093-8. URL: http://link.springer.com/10.1007/s00170-018-2093-8.

— (2020). "An Effective Predictive Maintenance Framework for Conveyor Motors Using Dual Time-Series Imaging and Convolutional Neural Network in an Industry 4.0 Environment". In: *IEEE Access* 8, pp. 121033–121049. ISSN: 2169-3536. DOI: 10.1109/ACCESS.2020.3006788. URL: https://ieeexplore.ieee.org/document/9131796/.

Koenig, Frank, Pauline Found, and Maneesh Kumar (Sept. 2019). "Improving maintenance quality in airport baggage handling operations". In: *Total Quality Management & Business Excellence* 30.sup1, S35–S52. ISSN: 1478-3363. DOI: 10.1080/14783363.2019.1665772. URL: https://www.tandfonline.com/doi/full/10.1080/14783363.2019.1665772.

Koenig, Frank, Pauline Anne Found, and Maneesh Kumar (Aug. 2019). "Condition monitoring for airport baggage handling in the era of industry 4.0". In: *Journal of Quality in Maintenance Engineering* 25.3, pp. 435–451. ISSN: 1355-2511. DOI: 10.1108/JQME-03-2018-0014. URL: https://www.emerald.com/insight/content/doi/10.1108/JQME-03-2018-0014/full/html.

Koenig, Frank, Pauline Anne Found, Maneesh Kumar, and Nicholas Rich (July 2020). "Condition-based maintenance for major airport baggage systems". In: *Journal of Manufacturing Technology Management* 32.3, pp. 722–741. ISSN: 1741-038X. DOI: 10.1108/JMTM-04-2019-0144. URL: https://www.emerald.com/insight/content/doi/10.1108/JMTM-04-2019-0144/full/html.

Lei, Yaguo et al. (May 2018). "Machinery health prognostics: A systematic review from data acquisition to RUL prediction". In: *Mechanical Systems and Signal Processing* 104, pp. 799–834. ISSN: 08883270. DOI: 10.1016/j.ymssp.2017.11.016. URL: https://doi.org/10.1016/j.ymssp.2017.11.016%20https://linkinghub.elsevier.com/retrieve/pii/S0888327017305988.

Li, Hongfei et al. (Aug. 2014). "Improving rail network velocity: A machine learning approach to predictive maintenance". In: *Transportation Research Part C: Emerging Technologies* 45, pp. 17–26. ISSN: 0968090X. DOI: 10.1016/j.trc.2014.04.013. URL: http://dx.doi.org/10.1016/j.trc.2014.04.013%20https://linkinghub.elsevier.com/retrieve/pii/S0968090X14001107.

Li, Wei et al. (Jan. 2013). "Design of Online Monitoring and Fault Diagnosis System for Belt Conveyors Based on Wavelet Packet Decomposition and Support Vector Machine". In: *Advances in Mechanical Engineering* 5, p. 797183. ISSN: 1687-8140. DOI: 10.1155/2013/797183. URL: http://journals.sagepub.com/doi/10.1155/2013/797183.

Lin, Jie et al. (Oct. 2017). "A Survey on Internet of Things: Architecture, Enabling Technologies, Security and Privacy, and Applications". In: *IEEE Internet of Things Journal* 4.5, pp. 1125–1142. ISSN: 2327-4662. DOI: 10.1109/JIOT.2017.2683200. URL: https://ieeexplore.ieee.org/document/7879243/.

Liu, Ruonan et al. (Aug. 2018). "Artificial intelligence for fault diagnosis of rotating machinery: A review". In: *Mechanical Systems and Signal Processing* 108, pp. 33–47. ISSN: 08883270. DOI: 10.1016/j.ymssp.2018.02.016. URL: https://doi.org/10.1016/j.ymssp.2018.02.016%20https://linkinghub.elsevier.com/retrieve/pii/S0888327018300748.

Liu, Yuehua et al. (Aug. 2020). "Noise Removal in the Presence of Significant Anomalies for Industrial IoT Sensor Data in Manufacturing". In: *IEEE Internet of Things Journal* 7.8, pp. 7084–7096. ISSN: 2327-4662. DOI: 10.1109/JIOT.2020.2981476. URL: https://ieeexplore.ieee.org/document/9039599/.

Medina-Oliva, Gabriela et al. (Sept. 2014). "Predictive diagnosis based on a fleet-wide ontology approach". In: *Knowledge-Based Systems* 68, pp. 40–57. ISSN: 09507051. DOI: 10.1016/j.knosys.2013.12.020. URL: http://dx.doi.org/10.1016/j.knosys.2013.12.020%20https://linkinghub.elsevier.com/retrieve/pii/S0950705113004000.

Meltzer, Matt (2017). *Delayed Baggage Compensation: What Airlines Really Owe You.* URL: https://www.cntraveler.com/story/delayed-baggage-compensation-what-airlines-really-owe-you (visited on 12/19/2019).

Naskos, Athanasios et al. (2020). "Event-Based Predictive Maintenance on Top of Sensor Data in a Real Industry 4.0 Case Study". In: *Communications in Computer and Information Science*. Vol. 1168 CCIS. Springer International Publishing, pp. 345–356. ISBN: 9783030438869. DOI: 10.1007/978-3-030-43887-6_28. URL: http://dx.doi.org/10.1007/978-3-030-43887-6%7B%5C_%7D28%20http://link.springer.com/10.1007/978-3-030-43887-6%7B%5C_%7D28.

Nuñez, David Lira and Milton Borsato (Oct. 2018). "OntoProg: An ontology-based model for implementing Prognostics Health Management in mechanical machines". In: *Advanced Engineering Informatics* 38.October, pp. 746–759. ISSN: 14740346. DOI: 10.1016/j.aei.2018.10.006. URL: https://doi.org/10.1016/j.aei.2018.10.006%20https://linkinghub.elsevier.com/retrieve/pii/S1474034617306080.

Peng, Hao and Qiushi Zhu (July 2017). "Approximate evaluation of average downtime under an integrated approach of opportunistic maintenance for multi-component systems". In: *Computers & Industrial Engineering* 109, pp. 335–346. ISSN: 03608352. DOI: 10.1016/j.cie.2017.04.043. URL: http://dx.doi.org/10.1016/j.cie.2017.04.043%20https://linkinghub.elsevier.com/retrieve/pii/S0360835217301900.

Peng, Ying, Ming Dong, and Ming Jian Zuo (Sept. 2010). "Current status of machine prognostics in condition-based maintenance: a review". In: *The International Journal of Advanced Manufacturing Technology* 50.1-4, pp. 297–313. ISSN: 0268-3768. DOI: 10.1007/s00170-009-2482-0. URL: http://link.springer.com/10.1007/s00170-009-2482-0.

Qu, Dingran et al. (Mar. 2021). "Research On ADCN Method for Damage Detection of Mining Conveyor Belt". In: *IEEE Sensors Journal* 21.6, pp. 8662–8669. ISSN: 1530-437X. DOI: 10.1109/JSEN.2020.3048057. URL: https://ieeexplore.ieee.org/document/9310211/.

Rijsenbrij, Joan C. and Jaap A. Ottjes (Aug. 2007). “New Developments in Airport Baggage Handling Systems”. In: *Transportation Planning and Technology* 30.4, pp. 417–430. ISSN: 0308-1060. DOI: 10.1080/03081060701461899. URL: https://www.tandfonline.com/doi/full/10.1080/03081060701461899.

Ruiz-Sarmiento, Jose-Raul et al. (Jan. 2020). “A predictive model for the maintenance of industrial machinery in the context of industry 4.0”. In: *Engineering Applications of Artificial Intelligence* 87.January 2019, p. 103289. ISSN: 09521976. DOI: 10.1016/j.engappai.2019.103289. URL: https://doi.org/10.1016/j.engappai.2019.103289%20https://linkinghub.elsevier.com/retrieve/pii/S0952197619302489.

Safaei, Mahmood et al. (Apr. 2020). “Standalone noise and anomaly detection in wireless sensor networks: A novel time-series and adaptive Bayesian-network-based approach”. In: *Software: Practice and Experience* 50.4, pp. 428–446. ISSN: 0038-0644. DOI: 10.1002/spe.2785. URL: https://onlinelibrary.wiley.com/doi/10.1002/spe.2785.

Sawant, Rohit and Prasad Kakade (Nov. 2018). “Airport Baggage Handling System Market Outlook: 2025”. In: *Allied Market Research*. URL: https://www.alliedmarketresearch.com/airport-baggage-handling-system-market.

Schmidt, Bernard and Lihui Wang (Oct. 2018). “Cloud-enhanced predictive maintenance”. In: *The International Journal of Advanced Manufacturing Technology* 99.1-4, pp. 5–13. ISSN: 0268-3768. DOI: 10.1007/s00170-016-8983-8. URL: http://dx.doi.org/10.1007/s00170-016-8983-8%20http://link.springer.com/10.1007/s00170-016-8983-8.

Sezer, Erim et al. (June 2018). “An Industry 4.0-Enabled Low Cost Predictive Maintenance Approach for SMEs”. In: *2018 IEEE International Conference on Engineering, Technology and Innovation (ICE/ITMC)*. June. IEEE, pp. 1–8. ISBN: 978-1-5386-1469-3. DOI: 10.1109/ICE.2018.8436307. URL: https://ieeexplore.ieee.org/document/8436307/.

Sikorska, J.Z., M. Hodkiewicz, and L. Ma (July 2011). “Prognostic modelling options for remaining useful life estimation by industry”. In: *Mechanical Systems and Signal Processing* 25.5, pp. 1803–1836. ISSN: 08883270. DOI: 10.1016/j.ymssp.2010.11.018. URL: https://linkinghub.elsevier.com/retrieve/pii/S0888327010004218.

Sittón-Candanedo, Inés et al. (Oct. 2019). “A review of edge computing reference architectures and a new global edge proposal”. In: *Future Generation Computer Systems* 99.2019, pp. 278–294. ISSN: 0167739X. DOI: 10.1016/j.future.2019.04.016. URL: https://doi.org/10.1016/j.future.2019.04.016%20https://linkinghub.elsevier.com/retrieve/pii/S0167739X1930264X.

Stefaniak, Paweł, Jacek Wodecki, and Radosław Zimroz (2018). “Maintenance Management of Mining Belt Conveyor System Based on Data Fusion and Advanced Analytics”. In: *Applied Condition Monitoring*. Vol. 10. August, pp. 465–476. ISBN: 978-3-319-62041-1. DOI: 10.1007/978-3-319-62042-8_42. URL: http://link.springer.com/10.1007/978-3-319-62042-8%20http://link.springer.com/10.1007/978-3-319-62042-8%7B%5C_%7D42.

Strauss, Patrick et al. (Dec. 2018). “Enabling of Predictive Maintenance in the Brownfield through Low-Cost Sensors, an IIoT-Architecture and Machine Learning”. In: *2018 IEEE International Conference on Big Data (Big Data)*. IEEE, pp. 1474–1483. ISBN: 978-1-5386-5035-6. DOI: 10.1109/BigData.2018.8622076. URL: https://ieeexplore.ieee.org/document/8622076/.

Syafrudin, Muhammad et al. (Sept. 2018). "Performance Analysis of IoT-Based Sensor, Big Data Processing, and Machine Learning Model for Real-Time Monitoring System in Automotive Manufacturing". In: *Sensors* 18.9, p. 2946. ISSN: 1424-8220. DOI: 10.3390/s18092946. URL: http://www.mdpi.com/1424-8220/18/9/2946.

Thirukovalluru, Raghuveer et al. (June 2016). "Generating feature sets for fault diagnosis using denoising stacked auto-encoder". In: *2016 IEEE International Conference on Prognostics and Health Management (ICPHM)*. IEEE, pp. 1–7. ISBN: 978-1-5090-0382-2. DOI: 10.1109/ICPHM.2016.7542865. URL: http://ieeexplore.ieee.org/document/7542865/.

Wuest, Thorsten et al. (Jan. 2016). "Machine learning in manufacturing: advantages, challenges, and applications". In: *Production & Manufacturing Research* 4.1, pp. 23–45. ISSN: 2169-3277. DOI: 10.1080/21693277.2016.1192517. URL: http://dx.doi.org/10.1080/21693277.2016.1192517%20https://www.tandfonline.com/doi/full/10.1080/21693277.2016.1192517.

Zhao, Rui et al. (Jan. 2019). "Deep learning and its applications to machine health monitoring". In: *Mechanical Systems and Signal Processing* 115, pp. 213–237. ISSN: 08883270. DOI: 10.1016/j.ymssp.2018.05.050. URL: https://doi.org/10.1016/j.ymssp.2018.05.050%20https://linkinghub.elsevier.com/retrieve/pii/S0888327018303108.