

Fatigue Damage Monitoring for Mining Vehicles using Data Driven Models

Erik Jakobsson¹, Robert Pettersson², Erik Frisk³, and Mattias Krysander⁴

^{1,2} *Epiroc Rock Drills AB, Örebro, 70225, Sweden*
erik.jakobsson@epiroc.com
robert.pettersson@epiroc.com

^{1,3,4} *Linköping University, Linköping, 581 83, Sweden*
erik.frisk@liu.se
mattias.krysander@liu.se

ABSTRACT

The life and condition of a mine truck frame are related to how the machine is used. Damage from stress cycles is accumulated over time, and measurements throughout the life of the machine are needed to monitor the condition. This results in high demands on the durability of sensors, especially in a harsh mining application. To make a monitoring system cheap and robust, sensors already available on the vehicles are preferred rather than additional strain gauges. The main question in this work is whether the existing on-board sensors can give the required information to estimate stress signals and calculate accumulated damage of the frame. Model complexity requirements and sensors selection are also considered. A final question is whether the accumulated damage can be used for prognostics and to increase reliability. The investigation is performed using a large data set from two vehicles operating in real mine applications. Coherence analysis, ARX-models, and rain flow counting are techniques used. The results show that a low number of available on-board sensors like load cells, damper cylinder positions, and angle transducers can give enough information to recreate some of the stress signals measured. The models are also used to show significant differences in usage by different operators, and its effect on the accumulated damage.

1. INTRODUCTION

A common approach for data driven prognostics and health management is to train models based on historic failure data. The method relies heavily on the quality of the data that can be obtained. This work targets how to generate prediction models for a mine truck in a real application.

Erik Jakobsson et al. This is an open-access article distributed under the terms of the Creative Commons Attribution 3.0 United States License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

There are two main approaches to obtain failure data. The first is to induce all imaginable types of faults in controlled tests. This approach can be costly, but also dangerous, for example in the case of nuclear power plants. Further, there are no guarantees the investigated faults are the ones that will actually occur. For frequent faults, like clogged filters or worn out brakes on a vehicle, the method could work well. But for certain types of components, such as the frame structure of a vehicle, it is both hard to detect actual damage through sensor signals, and time consuming to reach the worn out state.

A second method is to let a vehicle run in normal operation, and to collect both sensor data and record any faults. To test different algorithms, such an approach could be quite effective, since a large number of units could be monitored at a fairly low cost. But when the target is to generate a predictive maintenance model for a newly developed machine, the method has some major drawbacks. Collecting enough data for creating a failure model will take a long time, require many units, and are still likely to never see many faults that can occur. Another issue is that the data sets generated will most certainly be imbalanced due to the low number of faults and therefore hard to learn from as discussed in (Japkowicz & Stephen, 2002).

Slowly developing errors, such as cracks in frames, poses a specifically complicated case for collecting data since it takes too long to generate any failure data at all. Before enough data could be collected, the vehicle is likely to be outdated and preceded by the next model. The result is that neither of the two common methods to collect data are suitable in this case. Physics based modeling, PBM, is one approach when failure data is insufficient, as discussed in (Zio, 2016).

For the type of mine truck studied here, damage accumulation can be such a physics based approach. Instead of detecting actual faults, the complete time history of the vehicle can be

used to estimate how much of its operational life that has been consumed and compare to known material data. In this way, known theory on material behavior can be combined with on-line monitoring on a vehicle. This also enables validation of the method using only a small number of vehicles and relatively short time periods, assuming the relation between stress cycles and failure probability is correct. The rate of life consumption is dependent on how the vehicle is used, and on the external loads affecting it (Gurgenci & Guan, 2001). Such an approach is commonly used in fatigue analysis, where strain gauges can be used to measure stress in a structure and accumulated damage be calculated. However for mobile machinery, the introduction of a large number of strain gauges on each machine is of little interest given the increased complexity and high cost to maintain such a system (Molent, Barter, & Foster, 2012; Molent & Aktepe, 2000). Accelerometers instead of strain gauges is one option, but as seen in (Koistinen & Juuso, 2015) also these robust sensors are prone to failure when used on mining machinery.

This paper suggests an alternative to equip machines with high cost sensors, where different types of sensors already available on-board the machines are used to recreate a stress history signal. This is commonly known as operational load monitoring in the aerospace industry (Staszewski, Boller, & Tomlinson, 2004; Abelkis & Potter, 1979). In (Pais & Kim, 2015), an approach where finite element modeling is combined with an analytical model to recreate a stress signal is developed. This requires deep knowledge on how loads affect the structure. To reduce the need for such detailed system knowledge, the contribution of this work is how to use a data set under real operating conditions to find a linear relationship between standard on-board sensors and temporary strain gauge sensors on a vehicle, and how to combine the result with fatigue prediction to reach a measure of accumulated damage. The work also covers how to select signals for such a linear model, and what level of complexity is required. For future products, this relationship can be found using high fidelity simulation models, thus enabling the creation of a condition monitoring system already during the design phase of the vehicle. Some ideas such as the use of coherence, auto regressive models and damage accumulation are shared with previous work (Jakobsson, Frisk, Pettersson, & Krysander, 2017). Major differences from (Jakobsson et al., 2017) include results from different operating modes, segmentation, model complexity considerations and long term damage and prognostics. The data-set used for this work is also vastly different from previous work since other types of sensors and sampling frequencies are used.

Underground mining trucks share most mechanical characteristics with other more commonly known earth moving machinery, such as dumpers and dump trucks used in surface applications. Some differences worth mentioning are their more rigid and compact design, and the absence of damping at the rear axle. This makes them suitable to be modeled as

linear systems as proposed. However, the main difference to surface vehicles lies not within the mechanics, but rather in the harsh environment in which the vehicles are operated. Mines are a damp, dusty, and often corrosive environment. This makes sensor installation a difficult task. In many mines the maintenance of machines is also kept to a minimum, and only systems critical for the immediate operation are kept in working order. This is catastrophic for more long term measurements, where benefits for the customer only becomes useful over time, and in particular for the damage accumulation calculation, where a continuous measurement over the full life span of the machine is required. This makes the proposed technique of using available sensors, critical for short term operation, to estimate long term damage especially suited for mining machinery.

The use of existing on-board sensors is an enabler for implementing the models developed here on existing vehicles, and is one step towards the target of fleet wide prognostics (Zio, 2016). By comparing vehicles from different sites, applications and operating conditions the full potential of prognostics and health management can be unlocked.

2. PROBLEM FORMULATION

A mine truck contains numerous sensors intended to supply the on-board control system with information. Load sensors for production target follow up, inclination sensors for guaranteed stability, pressure sensors for controlling the hydraulic system are examples. The main question of this work is: Can these sensors be used to estimate damage on the mine truck frame for prognostic purposes, even though they were not at all designed for such a task? This main research question is divided into the following sub questions:

- Q1. From a large number of non-dedicated sensors, how can one select the most appropriate sensors to use for damage estimation purposes?
- Q2. Is it possible to use a model-free approach, such as statistics on a single on-board sensor, to estimate accumulated damage at a point on the frame?
- Q3. If a model-free single signal approach is not sufficient, is a linear model with multiple inputs such as an autoregressive model, enough to recreate a stress signal usable for damage accumulation calculations? And if so, is this possible for all different damage generating tasks a vehicle performs?
- Q4. How can the developed models be used for prognostic purposes, and in the long run help predicting failure, ease maintenance planning, and improve reliability?

To answer these questions, a large data set with vehicles from two different mine sites during real operation is used. Different model-free and model-based approaches are evaluated, and accumulated damage calculations are verified against the real

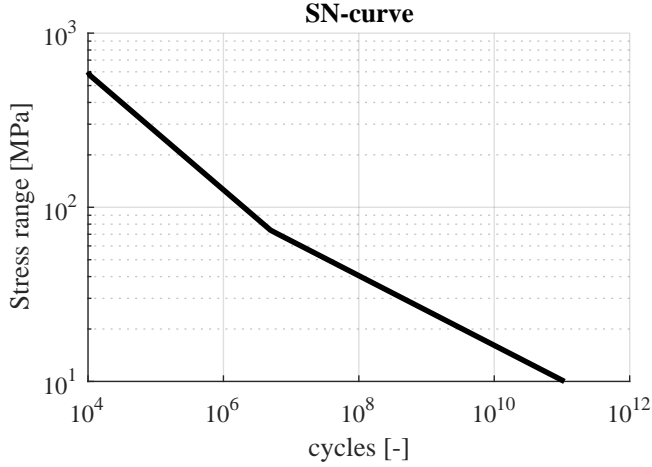


Figure 1. Relationship between stress range and the number of cycles until fatigue failure occurs for a given probability, here 2.3%. For a given stress range level, the structure can withstand a number of cycles as given by the curve.

data.

Once a model is found, the intention is to enable a condition monitoring system where the only modification required on the machine is a software upgrade. The condition monitoring capabilities will preferably cover any point modeled in the structure, but hopefully also generalize to a measure on how roughly the machine is used.

3. BACKGROUND

To give the reader some more background on the techniques used, a short summary follows. Three techniques are presented: Accumulated damage, Coherence, and ARX models.

3.1. Rain flow count and accumulated damage

Cycle counting and damage accumulation is commonly used to evaluate the damage caused by a certain stress-time signal. A typical metallic material can withstand a certain number of load cycles with a given stress range before failure. The number of cycles is dependent on the cycle amplitude, and the relation is often presented in the stress to number of cycles-diagram or SN-diagram for short. More information on metal fatigue and SN-diagrams can be found in (Stephens, Fatemi, Stephens, & Fuchs, 2000). For this work, the SN-diagram (Bygggavelningen, Göransson, & Åkerlund, 1999) in Fig. 1 is used.

A recognized way to handle spectral time-stress signals containing variable amplitude loading is by using the rain flow counting method (ASTME 1049-85 (Reapproved 1997), 1999) and then to apply the Palmgren-Miner rule for damage accumulation (Palmgren, 1924). Rain flow counting is used to define load cycles of varying amplitude from a stress-time-signal. The cycles are sorted in bins according to stress amplitude,

and the Palmgren-Miner rule stated as

$$D = \sum_{i=1}^k \frac{n_i}{N_i} \quad (1)$$

is used to evaluate and sum up the accumulated damage for each bin. In (1), n_i is the number of cycles at the stress amplitude indexed by i from the rain flow count, and N_i is the number of constant amplitude cycles until fatigue failure at the same stress range given by the SN-diagram. D is the accumulated damage.

For future prognostics and real-time implementations, work is available on real-time rain flow counting (Musallam & Johnson, 2012) but was not further investigated in this work since complete time series data was available.

3.2. The coherence function

In order to minimize the complexity of a future measurement system, it is important to have as few sensors as possible. To reduce the number of input signals, signals can be discarded based on what input signals contained the most information to describe a given output signal. The ordinary coherence function (Newland, 2012) describes how much of a systems output signal that can be explained using a linear relationship from the input signal, for each frequency covered in the measurement. If the ordinary coherence function approaches one, the relationship is purely linear and noise free. A low coherence indicates either a lack of relation between input and output, or that the relation is highly non-linear or contaminated by noise.

Let x_t and y_t be input and output signal respectively. Then the ordinary coherence C_{yx} is defined as

$$C_{yx}(f) = \frac{|G_{yx}(f)|^2}{G_{xx}(f)G_{yy}(f)}, \quad (2)$$

where $G_{yx}(f)$ is the cross spectral density, and $G_{xx}(f)$, $G_{yy}(f)$ the auto-spectral density for the input and output signal respectively.

3.3. System identification and ARX models

System identification can be used to create models from multiple input signals, to a single output signal. The single input single output ARX (AutoRegressive model with eXogenous inputs) model is given by

$$y_t + a_1 y_{t-1} + \dots + a_{n_a} y_{t-n_a} = b_1 u_{t-n_k} + \dots + b_{n_b} u_{t-n_k-n_b+1} + e_t \quad (3)$$

where y_t is the output at time t and u_t is the input at time t . Parameters $a_1, \dots, a_{n_a}, b_1, \dots, b_{n_b}$ are adjusted to create the best fit of the model, and n_k is an optional time delay. The term e_t represents the noise.

In matrix notation, expanded to multiple inputs, the expression can be written as a linear regression

$$y(t) = \varphi_t^T \theta + e_t \quad (4)$$

where

$$\begin{cases} \varphi(t) = (-y_{t-1}, \dots, -y_{t-n_a}, u_t^1, \dots, u_{t-n_b}^1, u_t^{n_b}, \dots, u_{t-n_b}^{n_b})^T \\ \theta = (a_1, \dots, a_{n_a}, b_1^1, \dots, b_{n_b}^1, b_1^{n_u}, \dots, b_{n_b}^{n_u}). \end{cases}$$

The multiple input signals are notated with superscripts from 1 to n_u where n_u is the number of input signals. The same notation is used for the corresponding set of parameters b^1, \dots, b^{n_u}

The parameter vector θ in (4) can be estimated using a least squares approach. This minimizes the residual between the measured and estimated stress signals, and since the least squares problem is convex, a global optimum is found. For more details, see (Ljung, 1999).

4. DATA AND SYSTEM DESCRIPTION

This section describes the data available and also briefly how the particular type of machine is used.

4.1. System description

The Epiroc MT65 mine truck is a heavy-duty machine designed for usage in underground mining. Equipped with a 567 kW engine, it is rated for 65 metric tons payload. A typical usage cycle includes the following:

- Loading, when a wheel loader drops rock material into the dump box of the truck. Typically 3 scoops are required to fill the box, resulting in over 20 metric tons per scoop.
- Hauling, when the truck moves the material. A common scenario is driving up a steep incline, possibly for hours, until the machine reaches the surface of the mine. Roads vary from paved roads to very rough gravel roads.
- Unloading, when the truck lifts the dump box and the load falls off.
- Driving empty, when the truck drives back to be loaded once more.

There is also a risk other use cases contribute with a non-negligible amount to the accumulated damage. Examples of such cases are compaction of the load using a large wheel loader or hitting a rock wall while driving. This type of events are unknown, but possibly present in the data sets available.

4.2. Available sensors

The machines used in the study contain both standard sensors and a number of high sampling rate sensors not commonly available on the Epiroc Mine trucks. Some sensors include

multiple directions, which is indicated with x,y,z in the sensor name. Some key sensor positions are shown in Fig. 2.

Standard sensors on the machine are:

- Vehicle speed, (e1)
- Load sensor at front point, (lc1)
- Load sensor at left pivot point, (lc2)
- Load sensor at right pivot point, (lc3)
- Inclination sensor in vehicle pitch direction, (incx)
- Inclination sensor in vehicle roll direction, (incy)
- Length of left shock absorber, (h1). The equilibrium location can be manually controlled.
- Length of right shock absorber, (h2). The equilibrium location can be manually controlled.
- Steering angle sensor, (ang1)
- Steering pressure, (p5). This is hydraulically connected to the dumping cylinders while dumping takes place.

Additional sensors, available during these trials only are:

- Two 3-axis accelerometers, (a1 and a2)
- One 2-axis accelerometer, (a3)
- Two pressure sensors for steering cylinders, (p1 and p2)
- Two pressure sensors for damping cylinders, (p3 and p4)
- One 3-axis gyroscope, (g1, only on a few experiments)

To be able to calculate reference damage, the data set also contains 25 strain gauges attached to the frame and dump box of the machine, s1-s25. The locations of these sensors were chosen as critical locations with the purpose to verify the development Finite Element Analysis models. The selection was not part of this work. Not all sensors were available at all times. In particular sensors s11, s12, and s14-s16 on the dumpbox were removed due to a change of dumpbox at some time during the trials. From the available sensors, s6 was selected for further investigation and model generation for two main reasons: First, it was one of few strain sensors that didn't suffer from various kinds of faults during the measurements. Second, the location of s6 makes it interesting since strain at this location is affected both when driving and when loading/unloading.

For a shorter measurement period it is possible to use strain gauges, even if they are considered too fragile to be useful during the full life of a machine.

4.3. Processing of data

Data was collected using two separate logging systems. One system handled logging of strain gauge and accelerometer data, at a sampling rate of 500Hz. The other system logged all messages available on the vehicle's communication bus. Different messages were sent at different rates, but most signals of interest were available at approximately 25Hz.




Figure 2. A selection of sensor locations on the mine truck frame. Remaining sensors s18-s25 were located on the dump box, not included in this image. Important standard sensors are highlighted in green, and the investigated output signal in red.

Data was post-processed to generate a homogeneous data set suitable for system identification and other forms of machine learning. The following procedure was used:

- Resample both data-sets to 50Hz using an anti-aliasing low pass filter and zero padding.
- Use correlation analysis on a pressure, that was measured using both logging systems, to find time shifts between the signals.
- Merge 50Hz data into a single array.
- Validate signal quality and reject faulty measurements (mainly failing strain gauges).

To ensure that the down sampling of the signals did not cause major loss of information, accumulated damage for down sampled stress signals are compared to accumulated damage for the high frequency signals. Fig. 3 shows how much difference it makes on the accumulated damage to down sample the strain signals to 50Hz. As seen, some sensor positions capture most of the accumulated damage at 50Hz, while other positions do not. The reason is that damage in some positions are driven by high frequency components, and other by low frequency components.

Accurate synchronization of the signals from the different measurement systems showed to be of high importance for the modeling technique used. Due to differences between the two measurement systems' clocks, the signals in the two systems slowly drifted apart and this can cause problems when several measurements are used to generate a model. Since the modeling technique used relies on a linear relation between input and output signal, any increase/decrease in time delay between the signals cause inaccuracies. To solve the issue, the data is synchronized using shorter time intervals.

Table 1 shows a summary of the available data set. Not much is known about the actual usage of the machine, since no detailed log-book was kept during the trials. From investigating the full data set one can draw some conclusions about the main differences of the sites:

- The truck at site 1 haul smaller loads, with an average

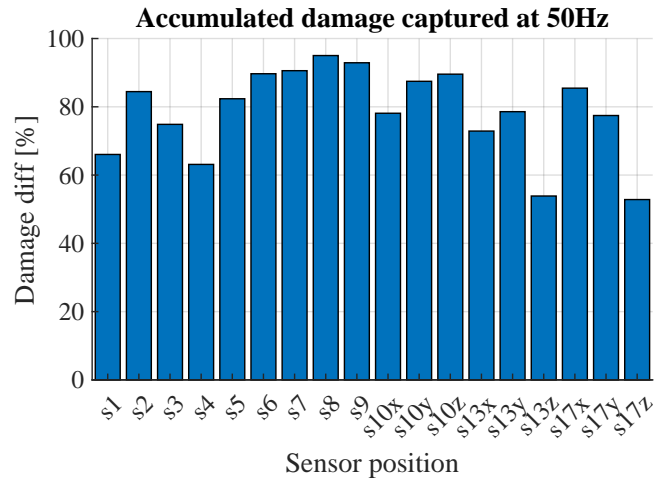


Figure 3. Accumulated damage captured by a downsampled strain signal compared to the original signal. For signal s6, around 90% is captured, meaning that most damage originates from low frequency components at this position.

Table 1. The data set available

	Site 1	Site 2
Number of vehicles	1	1
Number of load cycles	775	190
Calendar time	4mo, 17d	2mo, 14d
Total time	442h	238h
Hauling time	130h	116h
Driving empty time	102h	83h

around 38 tonnes. The truck at site 2 often uses the full payload capability with an average around 53 tonnes.

- The truck at site 1 drives shorter distances, with a typical duration of 0.5 hours complete cycle. The truck at site 2 drives longer distances, with a typical duration of 2.5 hours.

5. SIGNAL SELECTION AND AVAILABLE INFORMATION

The data for this work contains a large number of sensors and an important task is to investigate what input signals are most informative about the output signals. This section describes how signal selection can be done with the purpose to create a linear model, and thus answers question Q1 from the problem formulation. The selection is also important for a future implementation of such a system, where the number of sensors should be kept to a minimum for cost and complexity reasons. Linear models are of interest, which lets us use the coherence function to give valuable information on what input signals contain information on the output signals.

Fig. 4 shows the average coherence in the range 0 to 12.5Hz between different signals for a half-hour long driving segment.




Figure 4. Average coherence in the range 0-12.5Hz for input-output signal combinations. Black indicates high coherence and white low coherence. Sensors above the red line are standard sensors on the vehicles today, and are the sensors considered for further investigation.

The frequency range 0 to 12.5Hz corresponds to the Nyquist frequency of the signals with the lowest sampling rate, and is thus the range where the signals are valid without aliasing. The figure gives an overview of which input signal that can be used to estimate an output signal with a linear model. A black field indicates high coherence between the two signals, and a white field indicates no coherence. No field is completely black, since this would mean that there is no noise present in the measurement. It would also mean that no other input signal contains any additional information, since the output is completely explained by the single input signal and a linear model through the definition of coherence.

More details on the frequency dependency of coherence are seen in Fig. 5 that shows how some signals have high coherence in specific frequency regions. Signal lc3 is a good choice of the standard sensors for use in a linear model, showing high coherence from 0-5 Hz. This is a reasonable result, since lc3 is a load cell located near the s6 strain sensor. Signal h1 is a length sensor for the vehicle suspension, showing some coherence in the region 0-6Hz, but much lower than the load cell. Signals p5 and ang1 are a pump pressure and the waist steering angle respectively. Their low coherence shows they do not need to be part of the model development as they contain little information on strain at s6 while driving.

The specific selection of signals for different tasks of the vehicle are seen in Fig. 6. The signals with the highest coherence are chosen for each case, and the selection for further work is marked by "x". Driving tasks, empty and hauling, depend on the same sensors as expected. The loading task is mainly dependent on the load cells located close to the box being filled with material. The unloading task shows low coherence for sensor lc1, and high coherence for p5. This is intuitive since force is shifted from the load cell lc1, to the hydraulic




Figure 5. Coherence between a selection of input signals, and the s6 output signal. Signal lc3 shows a large linear relationship as is probably useful to create a linear model. p5 and ang1 can be deselected for further work since they show hardly no coherence at all for the s6 signal.




Figure 6. Level of coherence (shaded). Symbol "x" marks the input signals selected (horizontal axis) for model generation for signal s6. The choice is based on the level of coherence for each signal during the performed task (vertical axis).

cylinders with pressure p5 during unloading.

The coherence investigation shows there are linear relations available between a number of available on-board sensors in the data set. This suggests that available on-board sensors do contain useful information on damage generating processes. Whether this is sufficient to create useful models is investigated in the following sections.

6. MODEL FREE APPROACHES

For a condition monitoring system for a mine truck, the simplest possible algorithm is preferred. To correlate individual statistic measures from the input signals to accumulated damage (see Section 3.1) at the output signals represents such algorithms. A number of investigations are presented below to answer question Q2 in the problem formulation. The target to find if a model free approach is sufficient for the task of estimating accumulated damage. Only driving with the vehicle is included in these examples.

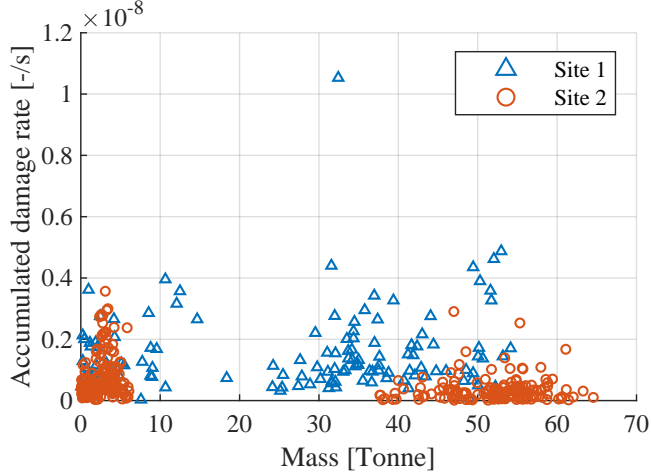


Figure 7. Damage accumulation rate for a large number of hauling segments. The figure shows no clear relationship between mass carried by the vehicle and the accumulated damage suffered.

Accumulated damage for signal s is given by (1) as

$$D = \sum_{i=1}^k \frac{n_i}{N_i} \quad (5)$$

where the number of cycles n_i for each stress range i during time segment τ is given by

$$\{n_1, \dots, n_k\} = \text{rfc}(s(t)), t \in \tau \quad (6)$$

and damage accumulation rate for signal s is defined as

$$D_R = \frac{1}{|\tau|} \sum_{i=1}^k \frac{n_i}{N_i} \quad (7)$$

where $s(t)$ is the stress signal for time t , τ a fixed size time segment, D the Accumulated damage, D_R the Damage accumulation rate, $\text{rfc}()$ the rain flow counting algorithm, N_i the number of cycles until failure of amplitude range i , as given by Fig. 1, and k is the number of cycle ranges.

6.1. Mass carried

The mass carried by the vehicle is believed to influence the amount of damage generated. To verify this the damage accumulation rate, i.e., the damage accumulation for the hauling segment divided by its length was plotted against the load carried during the hauling segment. This was done for a large number of hauling segments from both available sites. Fig. 7 shows no clear trend between the load carried and the damage accumulation at location s6.

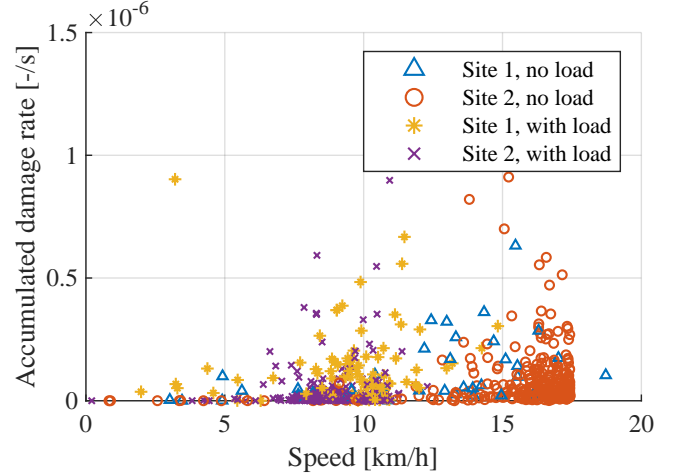


Figure 8. Damage accumulation rate for a large number of hauling segments. The figure shows no clear relationship between vehicle speed and the accumulated damage suffered, needed for the unloaded nor the loaded case of driving.

6.2. Travelling speed

Looking at the average speed traveled for a full hauling segment is not a good way to evaluate the influence of speed on accumulated damage, since the segments typically contain a range of speeds. A better option is to divide the segments into shorter sections with less varying speed, and compare the damage generation for such sections. Fig. 8 shows accumulated damage vs. vehicle speed for a large number of two-minute driving segments. The lack of trend between speed and accumulated damage shows that no simple relation exists. One reason for this could be that the driving speed is related to the road condition through the operator. If the road is rough, the operator is likely to reduce speed, to maintain a certain level of perceived vibration. This behavior would mask the influence of speed on accumulated damage.

6.3. Standard deviation of load sensors

The load sensors gives important information about the stress signals, as shown in the coherence analysis. Since oscillations drives accumulated damage, the standard deviation of these signals could be closely related, if the assumption of similar frequency content is done. Fig. 9 shows the damage accumulation rate vs standard deviation of the load cell signal for a large number of hauling segments for two different mine sites. Some trend can be seen that higher standard deviation corresponds to a larger damage accumulation rate. The large difference between the two different sites show that some other parameter is also effecting the relation, and only using the standard deviation will not be sufficient to predict the damage accumulation accurately.




Figure 9. Damage accumulation rate vs. standard deviation of a load sensor signal for a large number of hauling segments. The figure shows some relation between standard deviation and the accumulated damage suffered.

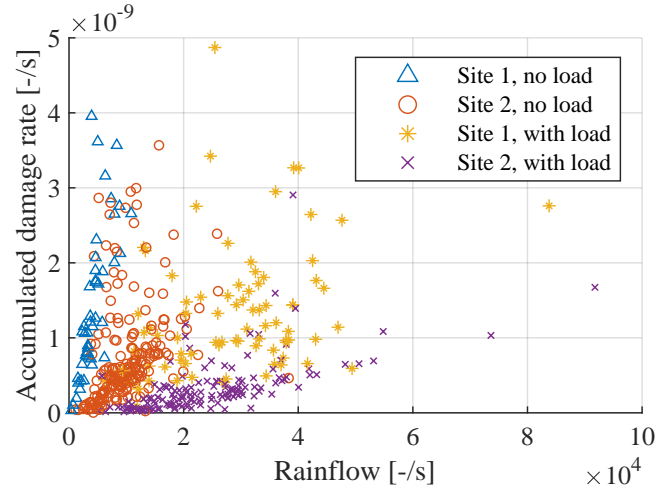


Figure 10. Damage accumulation rate vs. rain flow count on a load sensor signal directly. A slight linear trend can be seen, but the noise is large. Variation between the sites also exist.

Table 2. Correlation results against accumulated damage rate for the different measured quantities.

Measured quantity	Combined	Unloaded	Loaded
Mass	-0.14	-	-
Speed	0.09	-0.06	0.35
Std dev. load	0.09	0.02	0.59
Rainflow load	0.33	0.16	0.61

6.4. Cycle count on load cell signal

Using a rain flow cycle counting algorithm directly on an input signal is a way to capture similar load changes as is done in the strain signal. Fig. 10 show accumulated damage plotted against a rain flow count on one of the load cells. The linear trends in the figure indicate some relationship between the rain flow count on the load sensor and the associated damage. As in the standard deviation plot, the spread is large, and variations exist between the two test sites.

6.5. Correlation

Table 2 shows correlations in summarized form. Some moderate correlation is seen, in particular when separating out the driving cases with load. None of the investigated statistics are deemed sufficiently good for estimating accumulated damage. To use standard deviation or cycle counting on one of the load cell signals would however give a better result than to simply assume constant damage generation over time. Since large variations in damage are seen that are not explained by the single statistics, model-based approaches with multiple inputs are investigated.

7. DIVIDING THE DRIVING CYCLE

Since statistics on a single signal did not explain the variations in accumulated damage, models with multiple inputs are investigated next. First, a single model for the entire cycle is presented and then an improved approach using segments based on what operation the vehicle performs is shown.

7.1. A single model for entire driving cycle

To use a single model for the full driving cycle would be the least complicated approach of using a linear model. There would be no need to keep track on what the machine is doing, or when to switch between different models. Fig. 11 shows that such an approach is not possible using a single linear model from all available input signals. During the first section, the machine is driving without load and stress is underestimated. During loading the model is unable to recreate the amplitude of the steps as mass increases. Hauling is fairly well represented but stresses from unloading are missing completely. Looking at the vehicle structure, these results are intuitive. While driving, the vehicle is largely linear as long as the load does not move around considerably. The main source of excitation is the road profile, entering the vehicle through the wheels. When unloading the vehicle, the load paths change considerably. All force acting on the front load cell is shifted over to the dump cylinders. As the load is lifted, load also falls off, generating even more changes to the previous mechanical system which had constant mass. The driving force changes from wheel input, to oscillations from load cylinder end stops, and falling rocks. To have a single linear model capturing such variations does not work well and a multi-model approach is needed.




Figure 11. Results from having a single model for the entire operating cycle of the machine. The driving portion is estimated fairly well. The loading portion fails to capture the magnitude of stress as a result of load increase. Worst result is seen for unloading, where the model completely misses the very important oscillation peaks.

7.2. A need for segmentation

To find better models, the data needs to be partitioned into different segments where the vehicle behavior is as close to linear as possible within the same segment. The main changes of the vehicle creating non-linearities are change of mass in the box, and change of shape of the vehicle when raising the box to unload. This gives four main segments: driving empty, loading, hauling, and unloading. A simple and sufficient method developed for dividing the measurement data into these segments is described next.

First the break points for different segments are located, based on two conditions. If the vehicle speed passes a threshold above the noise level of the speed signal, a new segment is created. For this case 0.1km/h is chosen. Loading and unloading typically takes place while the machine is stationary, and this condition creates new segments for each load change. If the sensor for checking if the box is closed changes state, a new segment is created since this indicates that the machine has changed shape by opening the box. The mass signal is not a good candidate for setting the break points, due to its high noise level while driving.

The second step is to determine which of the following conditions are true between each break point from step one. These conditions ensure that no single segment can be placed in two different categories.

- Idle: Average speed < 0.1 km/h and no mass change > 3 tonne and no box open.
- Hauling: Average speed > 0.1 km/h and mass in box > 10 tonne and no box open and no mass change > 3 tonne.
- Unloading: (Box open or a mass decrease > 3 tonne or




Figure 12. Drive cycle segments for one full working cycle based on speed, mass and box location

next segment) and average speed < 0.1 km/h .

- Loading: Mass increase > 3 tonne during the segment and average speed < 0.1 km/h and no box open.
- Driving empty: Average speed > 0.1 km/h and mass in box < 10 tonne and no box open and no mass change > 3 tonne.
- Mixed: Average speed > 0.1 km/h and (mass change > 3 tonne or box open), i.e., driving while loading or unloading.

The segments of a complete driving cycle are shown for each segment on the measured mass signal in Fig. 12.

7.3. Damage accumulation for different segments

To show the need for modeling the different segments, the accumulated damage (see Section 3.1) for each segment is shown in Fig. 13 and Fig. 14. For signal s2, located on a suspension component in the front of the truck, hauling and driving empty causes most of the damage as seen in Fig. 14. For signal s6, it is the unloading task that generates the main part of the damage as seen in Fig. 13. Hence the unloading model is very important to capture most of the damage on signal s6. The need for different models to capture the full cycle is clear since loading, driving empty, and hauling all have large contributions to the overall damage.

8. LINEAR MODELS FOR DIFFERENT SEGMENTS

This section targets question Q3 from the problem formulation, that is whether linear models are sufficient to recreate the stress history with purpose to calculate accumulated damage. Fig. 15 gives an overview of the process of first generating a model from sensor data, and then to compare both the stress and the accumulated damage to evaluate the model performance.




Figure 13. Accumulated damage by segment for sensor location s6




Figure 14. Accumulated damage by segment for sensor location s2

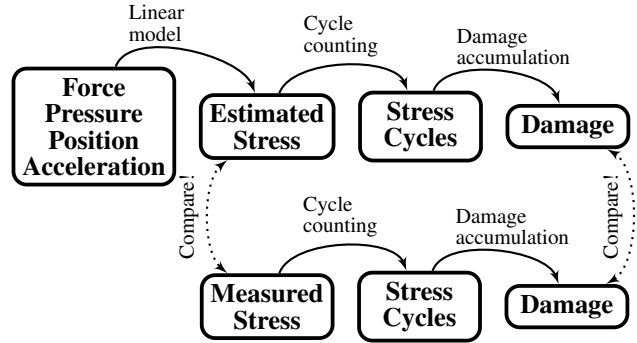


Figure 15. Schematic view of the process for calculating accumulated damage and comparing the results.

8.1. Model generation for the driving segments

The ARX (AutoRegressive model with eXogenous inputs) structure is considered in this work, because more advanced linear models did not show any major improvements in a previous work (Jakobsson et al., 2017).

A model according to (4) with $n_a=8$, $n_b=16$ and $n_k=0$ is shown to be a good compromise of accuracy and complexity, by testing a large number of model orders. The ARX model structure allows such testing with a low computational cost given its convex closed form solution. Fig. 16 shows a comparison between models generated from the standard on-board sensors, and models generated from external sensors such as accelerometers only. It also shows a combined model consisting of all the sensors. Using standard on-board sensors outperforms the use of only accelerometers as seen on the higher visual similarity between on-board sensor simulation and measurement. Fig. 17 shows the accumulated damage for the same case over a longer time period. Even though the accelerometer based model captures fast dynamics, it misses much of the slow moving changes. And since the slow moving changes typically result in large amplitude, it is better to have a model accurate in those regions. For this example a linear model is sufficient to capture over 80% of the measured damage during the 500 second hauling segment.

One issue in (Jakobsson et al., 2017) was how to handle the effect of different amounts of load in the box. In the previous work, not much data with varying load level was present, preventing the investigation of load effects. The current data set contains many more cycles with different load levels and thus enable such investigations. Two alternatives for handling varying load levels are presented below.

The most convenient approach is to let many different mass levels be part of the training data, and identify a model suitable to all. The background to such an approach is that the load cells for measuring the mass in the box is available as dynamic sensor signals. However, using an ARX model did not yield satisfactory results for damage accumulation using




Figure 16. Signals generated from different models compared to measured data. The standard on-board sensors, a combination of load cells, damper position and angular sensors outperform the use of accelerometers only. This is seen on the much better resemblance between these signals to the measured. Best result is obtained using all the sensors.

this approach. An alternative approach to a single model, is to estimate one model for each suitably divided mass range. Fig. 18 shows the mean squared error (MSE) between simulated and verification data, for a large number of driving cycles. Each dot corresponds to one driving segment, evaluated using one model and compared to verification data. A perfect model would result in zero error for each cycle evaluated. Three different models are compared. The “0-80” model is trained using cycles with different load levels between 0-80 tonnes. The “0-20” model and “30-80” models are estimated using data from respective mass range.

Despite the load mass information being available to the “0-80” model through the load sensor signals, the structure is unable to capture the different dynamics involved when driving with/without load and the result is average for both low loads and high loads. A good result for the “0-80” model would show as small error for both low loads and high loads. One possible reason for the poor result could be nonlinear relations between the load level and the rate of damage generation. Other possible reasons could be that the load cells were somewhat unreliable during the measurement, causing load changes without the mass actually changing.

The “0-20” model show low mean square error for the low load cases, but very high MSE for the high load cases, and the “30-80” model shows the opposite. This shows that a linear model is able to capture the driving dynamics well, but fail to capture the dependence from different mass levels.

By using a separate model for the low- and high load case, a direct comparison with correlation results from the model-free approach in Table 2 is possible. The ARX model output gives a correlation of 0.26 in the unloaded case, 0.74 in the loaded




Figure 17. The accumulated damage for location s6 is dominated by large cycles. Since the on-board sensors do a better job capturing such changes, it outperforms the use of accelerometers only. A combined model using all sensors, gives a slightly better result.

case and 0.56 for the combined case, all higher than the best results from Table 2.

8.2. Model generation for the loading segment

During loading, the truck is typically stationary. A wheel loader drops rocks and gravel into the box until it is filled, which typically takes 3-4 buckets depending on the size of the wheel loader used. This is seen in the stress signal in Fig. 19 as four distinct plateaus.

The model was trained from a number of loading cycles from site 1. Validation shows that the model capture fast dynamics well, i.e., the oscillations that occurs at each mass-increase. However the model fails to get the static levels correct, and simulations for specific loading cases often drift. One example of such a case suffering from drift is shown in Fig. 19. This has significant effect on the accumulated damage, since large erroneous cycles emerge from the drift.

The drift in model output is caused by the strain gauges drifting during the measurements. The model tries to capture this output signal drift, by adding a pure integrator to the ARX model. The drift is not constant between different loading cycles, and does not correspond to any change in input signal. When the model later is used to simulate an output, the integrator causes drift in the output that do not exist in the measured data. To improve the models for the loading segment, either new measurements without drift issues, or a model structure better capable of handling such errors in the training data are required.

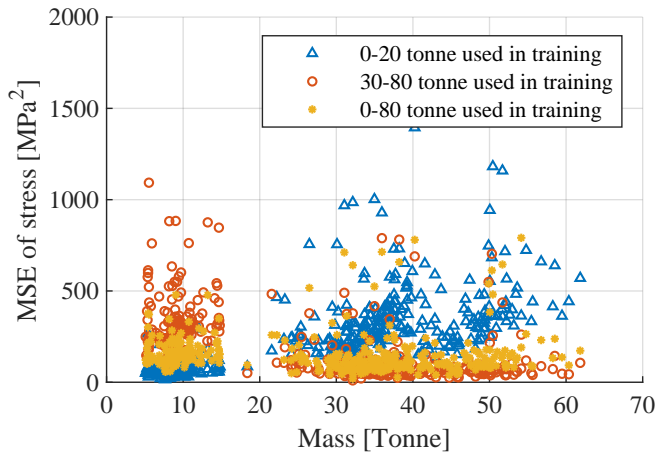


Figure 18. Mean squared error between simulated and measured stress is used to compare three models for a large number of driving cycles. A model trained from a wide range of mass levels results in a compromise for both high and low loads. Training separate models for low/high load reduce the errors in the range where it was trained.

8.3. Model generation for the unloading segment

During unloading, hydraulic pressure is applied to two lift cylinders. This causes stress close to the attachment points of the cylinders and is seen at (A) in Fig. 20. The lift cylinders are 4-step telescopic cylinders, and for each change of steps a small stress oscillation is seen (B,C,D). At (E), the box is first raised to its maximal position and hits the end stop. Rocks and gravel fall out, and the box is closed again. The operator can choose to lift/lower the box arbitrarily, and may also hit the end stop multiple times if needed to get the gravel to fall out. This is seen at (F,G,H). These additional end-stop hits create large stress oscillations. When lowering the box, the operator can either use the main pump, or simply letting the box free fall to its minimum location. Each step in the cylinders closing is seen at (I,J,K), but causes only minor oscillations.

The model captures all major events in the signal, but the accumulated damage is underestimated as seen in Fig. 21. This is due to two effects. One is inability to capture slow moving changes of the signal, causing some of the largest cycles to be distorted. The other is a constant underestimation of peak height for oscillations. This underestimation could be caused by noise in the timing of different signals.

To demonstrate the difference in performance between different unloading cycles, accumulated damage for a large number of cycles is plotted. The cycles are sorted by the simulated damage of the cycle. As seen in Fig. 22 the model consistently captures around half the actual damage, and this is true for both low-damage cycles and high-damage cycles. This enables the model to distinguish a gentle unloading cycle from a harsh unloading cycle. The importance of this distinction is discussed further in Section 10.2. An alternative approach




Figure 19. The model captures dynamic oscillations, but drifts considerably. This cause very bad performance when calculating damage.

would be to assume a constant amount of damage contribution for each unloading cycle. As seen in Fig. 22, choosing the level for such a contribution would be hard given the large difference in accumulated damage between different cycles. Despite the large spread in how much damage the linear model predicts, it outperforms the constant-damage-per-unloading approach since it takes into account the variability between different cycles.

9. LIMITATIONS OF THE MODELS

A number of limitations exists both for the proposed types of models, and the methods to generate them. Most prominent is the general issue of data driven methods: It is very hard to estimate how well the models will generalize to unseen cases, such as the machine hitting the wall while driving. This needs to be kept in mind when using the models.

Another important limitation comes from the assumption that damage is only generated from forces that can be detected in the measured signals. If there is a way to create strain in the structure that do not effect the monitored load cells, inclination sensor and pressure sensors, etc, this strain will generate damage unnoticed by the the monitoring system. Since no such strain was identified during training, it is however unlikely that normal operating cases causes such strain. For external events however, we know little about such effects. Also the sampling rate of the sensors also plays a part, since short peaks can cause damage, but pass completely unnoticed by the monitoring system. A good understanding of the strain frequency contents in point of interest is crucial for successful application of the models.

The process of generating the models from measured data




Figure 20. Simulated data vs. measured strain data for the unloading of the box. The model captures all distinct events during unloading, but underestimates peak height causing poor performance when calculating damage.

poses a limitation especially for newly developed vehicles. The cost and effort required to create a suitable amount of data, covering all relevant driving cases, is not to be underestimated. One way to mitigate this limitation would be to generate the models using data from high fidelity simulation models.

10. EXTENDING THE RESULTS TO PROGNOSTICS

This section covers question Q4 from the problem formulation, how to use the models to improve reliability of the vehicles. The models developed in this work are light-weight from a computational point of view. This enable real time calculations of the stress history in a large number of positions, and also damage accumulation calculations in real time. This can be utilized for prognostic purposes in a number of ways.

10.1. Levels of damage estimation

The simplest way of estimating current damage, is to assume all machines are used equally, and let damage be a function of engine hours only. Based on experience from other vehicles, an average damage per time unit is used to evaluate the current condition. The most obvious flaw with such a method, is that no care is taken with regard to machine usage. No individualization of the machines can be done.

A more advanced scheme would be to use drive cycle segmentation, and to base current accumulated damage on how much time the machine spends doing different tasks. An unloading cycle, for example, generates vastly more damage on some locations than driving without load. This would give a more accurate result on a per-vehicle basis, but still allow a lot of room for error with regards to how roughly the machine is used. In short we could differentiate vehicles that run long hauling cycles and few unloading cycles from vehicles running short hauling cycles and many unloading cycles.




Figure 21. Accumulated damage during unloading. A gross underestimation of damage is the result from poor peak-height estimation in the stress signal.

With access to virtual stress measurements on various locations on the machine, a new paradigm of individualized monitoring is possible. Knowing the stress history enables a system where calculated accumulated damage is available in real-time. Fig. 23 shows a complete cycle, where the simulated stress signal is shown on top of the measured stress signal. Different models are used for the different segments of the cycle. Access to such stress-time history lets us differentiate vehicles on a per-cycle basis. The following section shows a real data example of how the information can be used.

10.2. A concrete example of individualized monitoring

The following example is based on the real data from the measurements used in this work. Fig. 24 shows the accumulated damage at position s6 for each different segment for every cycle during two weeks of operation at site 1. The upper dots of each color show measured damage, and the lower dots show damage as calculated from the simulated stress signals. The total damage is underestimated by the model, but more importantly, the variations in damage rate are captured by the model. As seen, most of the total damage at position s6 originates from the unloading segment. The cycles appear grouped in short sections lasting about half a day, which are interpreted as the work shifts on site. Another pattern seen is that the slope of damage accumulation seems to alternate between a high and low value. This is particularly clear around day 10 as emphasized by two red lines corresponding to the slope. By manual investigation of the data during these unloading cycles, it could be shown how two different operators used very different techniques to unload the vehicle. The first was more aggressive, running all functions at maximum. The second was gentler, slowly lifting the box and reducing lifting speed before hitting the end stop. Such difference in usage has a large impact on the life of the vehicle, and by measuring the




Figure 22. Accumulated damage for a large number of unloading cycles, sorted after simulated damage. The linear model captures around half the actual damage, both for small and large damage cycles.




Figure 23. Stress for the joined segments of a full driving cycle. Simulated stress is shown as one color per segment type on top of the black measurement.

accumulated damage the difference can be found and acted upon.

10.3. Continuous feedback of machine life

By analyzing the damage generated during different segments, we can learn how a particular individual vehicle is used, and how damage is likely to evolve over time. The influence from operator behavior is already shown in the previous section. Similar differences are likely to be found from different factors in the mine, such as road condition and size of material transported. Other possible differences could be found in the vehicle configuration, such as air pressure in the tires.

As long as the current damage level is known, and the usage of the machine is not expected to change, a prediction of future

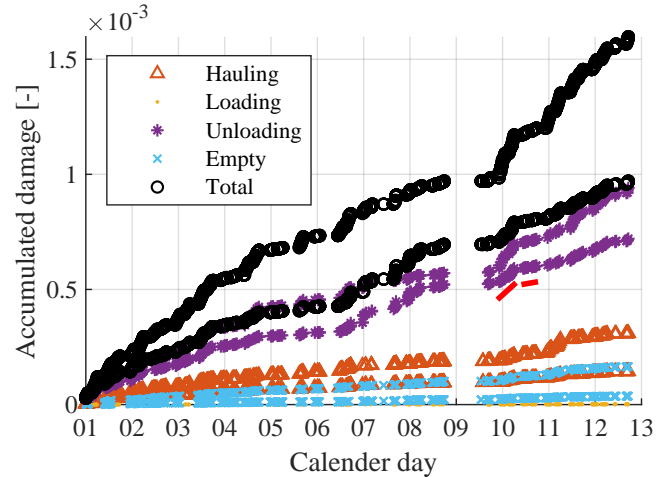


Figure 24. Damage accumulation for different segments for all cycles during a two week period. Upper dots of each color show measured damage and lower dots simulated damage, since the simulation models underestimated the damage. Dots are clustered in small half-day long groups representing the work shift. Interesting to note is the large differences in damage accumulation for different shifts, as shown by the red lines at day 10-11.

damage can thus be made with much better accuracy than using a time average. This prediction can in turn be used to schedule maintenance intervals, and in the long run also the economic life of the entire machine.

An even more important application for monitoring the damage continuously, is when machines become more and more autonomous. With an operator on-board, usage is adjusted to local conditions. For an autonomous truck, it is not obvious how to choose optimal speed with respect to both production and life of equipment. Algorithms such as the ones presented in this work provides an important piece to designing such systems.

11. CONCLUSION

We have shown that available on-board sensors can be used to estimate damage on a mine truck frame. The investigation suggests that it is not possible to use any simple statistics on input signals to predict damage well on a per-vehicle basis. The usage of linear models show to be more reliable in capturing damage. The best result is obtained for a constant driving scenario, i.e., when the mass and shape of the vehicle do not change. To handle variations in mass and configuration of the vehicle, a segmentation of the usage cycle is proposed. When the data is divided in segments such as hauling, driving empty, loading, and unloading, the usage of separate models for each segment results in a better fit to verification data. Using this approach, linear models give a sufficient fit for all driving segments, at least to compare relative damage between different vehicles and usage.

When it comes to loading and unloading, linear models fail to get the static levels of stress correct. Nevertheless, a linear model for unloading is able to capture important differences in unloading behavior by the operators, and could be used in a relative manner. Finally, a number of ways to extend the results into prognostics are presented. By continuously knowing the accumulated damage of the vehicle, the task of predicting future damage can be individualized both to a specific machine, but also down to a specific shift or cycle. This enables improvements both for manually operated machines, but also in the case when mining machinery operates autonomously.

ACKNOWLEDGMENT

This work was partially supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP). We would like to thank Jari Hyvärinen, Ali Beg and Markus Bagge from Epiroc Rock Drills AB for their work on planning, collecting and cleaning the dataset used in this work, as well as for valuable insights.

REFERENCES

- Abelkis, P., & Potter, J. M. (1979). *Service fatigue loads monitoring, simulation, and analysis* (Vol. 671). ASTM International.
- ASTM E 1049-85 (Reapproved 1997). (1999). *Standard practices for cycle counting in fatigue analysis* (Vol. Vol. 03.01; Tech. Rep.). Philadelphia.
- Byggavdelningen, B., Göransson, L., & Åkerlund, S. (1999). *Boverkets handbok om stålkonstruktioner - bsk 99*. Boverket.
- Gurgenci, H., & Guan, Z. (2001). Mobile plant maintenance and the dutymeter concept. *Journal of Quality in Maintenance Engineering*, 7(4), 275-286.
- Jakobsson, E., Frisk, E., Pettersson, R., & Krysanter, M. (2017). Data driven modeling and estimation of accumulated damage in mining vehicles using on-board sensors. PHM Society.
- Japkowicz, N., & Stephen, S. (2002). The class imbalance problem: A systematic study. *Intelligent data analysis*, 6(5), 429–449.
- Koistinen, A. H., & Juuso, E. K. (2015). Stress monitoring of underground load haul dumper front axle with intelligent indices. In *4th workshop on mining, mineral and metal processing* (Vol. 3000, pp. 25–27).
- Ljung, L. (Ed.). (1999). *System identification (2nd ed.): Theory for the user*. Upper Saddle River, NJ, USA: Prentice Hall PTR.
- Molent, L., & Aktepe, B. (2000). Review of fatigue monitoring of agile military aircraft. *Fatigue & Fracture of Engineering Materials & Structures*, 23(9), 767–785.
- Molent, L., Barter, S., & Foster, W. (2012). Verification of an individual aircraft fatigue monitoring system. *International Journal of Fatigue*, 43, 128–133.
- Musallam, M., & Johnson, C. M. (2012, Dec). An efficient implementation of the rainflow counting algorithm for life consumption estimation. *IEEE Transactions on Reliability*, 61(4), 978-986.
- Newland, D. E. (2012). *An introduction to random vibrations, spectral & wavelet analysis*. Courier Corporation.
- Pais, M. J., & Kim, N. H. (2015). Predicting fatigue crack growth under variable amplitude loadings with usage monitoring data. *Advances in Mechanical Engineering*, 7(12), 1687814015619135.
- Palmgren, A. (1924). Die lebensdauer von kugellagern. *Zeitschrift des Vereins Deutscher Ingenieure*, 68(14), 339–341.
- Staszewski, W., Boller, C., & Tomlinson, G. R. (2004). *Health monitoring of aerospace structures: smart sensor technologies and signal processing*. John Wiley & Sons.
- Stephens, R., Fatemi, A., Stephens, R., & Fuchs, H. (2000). *Metal fatigue in engineering*. John Wiley & Sons.
- Zio, E. (2016). Some challenges and opportunities in reliability engineering. *IEEE Transactions on Reliability*, 65(4), 1769–1782.