Liner Wear Prediction Using Bayesian Regression Models and Clustering

Jacob Van Den Broek[1], Melinda Hodkiewicz[2], and Adriano Polpo[3]

[1,3] *School of Physics, Mathematics and Computing, University of Western Australia, Perth, WA, 6009, Australia*
Jacob.Vandenbroek98@gmail.com
Adriano.Polpo@uwa.edu.au

[2] *School of Engineering, University of Western Australia, Perth, WA, 6009, Australia*
Melinda.Hodkiewicz@uwa.edu.au

Abstract

Chutes, bins, and hoppers are critical assets in bulk commodity handling. Sacrificial wear liners are employed to protect these assets from abrasive wear. An essential maintenance challenge is optimising the timing of liner replacements. Traditionally, episodic human inspections have been in place, but now, real-time wireless IoT sensing systems that measure liner thickness are being used. We propose a novel approach to estimate the remaining useful chute liner life. Instead of linear extrapolation based on individual sensor wear rates (commonly used in industry), we leverage a Clustered Bayesian Hierarchical Modeling (BHM). Two models are developed: Model 1 (Cluster Exemplar) uses parameters from the closest cluster exemplar, while Model 2 (Spatial and Temporal BHM) incorporates data from the active sensor, with prior distribution informed by Model 1. Data are drawn from a single hopper with 88 sensors, 20 of which reached their end-of-life threshold. Both Model 1 and Model 2 outperform the industry regression approach, significantly reducing overprediction. Notably, Model 2 predicts remaining useful life within 95% credible intervals and identifies anomalous sensor performance. This innovative use of historical and adjacent sensor data enhances wear degradation prediction, contributing valuable insights to the literature.

1. Introduction

Bulk commodities are a major contributor to the Australian economy with A$229 billion of products sold in 2023/24 (Thurtell, 2022). In the mining sector, the safe and efficient transfer of bulk ore is achieved through a network of assets and systems across a mine site. These include conveyor systems with multiple feed chutes, transfer chutes, bins, and hoppers. Due to the high-throughput and coarse nature of the ore, sacrificial wear liners are installed to protect these structures from abrasive wear. If these liners wear too far, there is damage to the structural steel and if they wear too quickly there can be unplanned downtime resulting in lost production. The replacement of worn liners is more cost effective than a repair of the structural asset. Premature failure of liners can lead to up to 50% more time dedicated to lining maintenance (Malone, Hu, Clinton, & Ore, 2013). Traditionally, liners are monitored, and data recorded by means of manual inspection, ultrasonic thickness testing (Padole, Joshi, & Engineer, 2002) or 3D laser scanning (Vaníček et al., 2012). These methods all require a shutdown to allow operators to enter a high-risk area to conduct inspections, the accuracy of which is dependent on the skill of that operator.

A safer alternative is an Internet-of-Things (IoT) sensor monitoring system for chute wear monitoring, such as WearSense™ developed by Metso, a global technology and services company in mining and aggregates. In the WearSense design the sensor is housed in a fastener, which bolts the liner to the chute wall. The sensor design, incorporates a probe that penetrates to the wear surface of the liner and is shown in Figure 1. The probe wears with the liner as material flows through the chute, recording thickness as longitudinal point estimates. A schematic of the installation of a sensor is also shown in Figure 1. Sensor output is transmitted to a cloud-based data warehouse where the data is processed and displayed on a web user interface. At present, linear regression is applied to the data to predict future thickness values for each sensor and interpolation between the sensors to provide a visual map of the wear profile inside the chute. As the wear of liners is measured by sensors, when referring to the wear of a sensor in the text, we are referring to the degradation of a liner at the particular location where the sensor is installed.

Wear is a complex mechanism and traditionally, in order to

https://doi.org/10.36001/IJPHM.2025.v16i1.4266

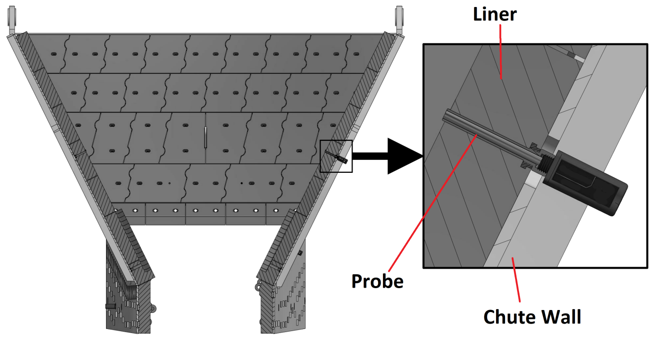

Figure 1. Schematic representation of a sensor installed on a single liner of a coarse ore hopper chute.

predict the lifespan of liners, it is necessary to understand the chute structural design, ore type, operating and environmental factors to develop wear models for each individual chute.

Previous methods to model wear in chutes have involved numerical and experimental studies. The application of experimental studies (W. Chen, Biswas, Roberts, O'Shea, & Williams, 2017; Hawk, Wilson, Tylczak, & Doğan, 1999) which simulate wear in controlled environments using experimental apparatus is limiting as fluctuations in the ore properties and external factors are not accounted for, therefore making them unsuitable for complex, real-world situations (Petrica, Badisch, & Peinsitt, 2013). Advancements in numerical methods have enabled modelling of spatial wear under operational conditions using discrete element modelling (DEM) and finite element modelling (FEM) methods (Forsström & Jonsén, 2016; Xu, Luo, & Zhao, 2018). The accuracy and rapid prediction of wear using numerical methods is limited due to expensive computation time and the assumptions of ideal modelling conditions (Ou & Chen, 2022).

Wear in chute liners can be visualised as a set of 3D images over time. Modern 3D scanning produces a map of the chute but is expensive to do requiring production to stop and the chute to be clear. When sensors are used we only have observations at discrete points in space representing a fraction of the surface area of the chute. When engineers analyse this data they have to interpolate between the sensors (in space) and examine the thickness measures of one or more sensors in time. These processes impose a heavy cognitive load and involve subjective interpretation.

The degradation of any physical process occurs over time and there are uncertainties associated when observing these processes. Well-known stochastic processes such as the gamma, inverse Gaussian and most commonly the Wiener process are used to describe the uncertainty pertaining to these processes. In predicting the remaining useful life of cylinder liners in ship engines (Hermann & Ruggeri, 2017), the authors present a stochastic model for degradation which accounts for abrasive wear and corrosive wear through a differential equation, considering a jump process and Wiener process respectively. Inference and forecasts were made in a Bayesian framework, computing posterior quantities through Markov Chain Monte Carlo sampling and satisfactory results were obtained. The authors were able to predict a probability for a thickness value at specified time intervals. The data used was simulated for 30 cylinders, each of which assumed to be realisations of the same process, therefore meaning the same model and parameters were used for each cylinder. In reality, each cylinder should not be regarded as realisations of the same process as they would be placed in different environments. This is also an issue for the purpose of this study and is discussed further in the hierarchical modelling section. Similarly (Guérin et al., 2010; Wang, Lin, Wang, He, & Zhang, 2016), modelled disc brake wear and the remaining useful life of a axial piston pump as Wiener processes. A pitfall of the Wiener process for a process such as wear is that it does not consider the monotonically decreasing property of wear (Limon & Yadav, 2021).

When examining thickness data from one sensor in a chute, it is expected that other sensors in the chute will have similar rates of wear, albeit with some local variations to account for place in the chute. This 'other' information is valuable and should be included in a model for predicting future thickness and can be used to identify if one sensor is exhibiting anomalous behaviour. Statistical methods based on Bayesian hierarchical modelling have been developed to incorporate additional information in the form of dependency between sensors into the models for a single sensor. Our interest here is in 1) prediction, making inference on a hidden state value at a time beyond the current time, and in 2) filtering, making inference on the hidden state at the current time based on the current and all past data from that sensor and other sensors (Wikle, Zammit-Mangion, & Cressie, 2019). A Bayesian hierarchical model (BHM) has parameters with prior distributions at the bottom level of the hierarchy. For further details on BHM the reader is referred to (van de Schoot et al., 2021).

In reliability, the BHM has been shown to be a useful technique for modelling collective and individual conditions of assets simultaneously compared to modelling all assets together. In 1996 (J. Chen & Singpurwalla, 1996) utilised Bayesian hierarchical models when estimating reliability of emergency diesel generators in separate nuclear power plants. In 2015, a BHM was established to estimate the remaining useful life of aerospace gas turbines operating in various conditions and compared it to a non-BHM, finding that the BHM was superior in typical (heterogeneous) scenarios (Zaidan, Harrison, Mills, & Fleming, 2015). In general, hierarchical modelling remains questionable when the first layer assumes homogeneity between and within the similar assets, when realistically different assets may have different environments.

In summary, automating processes for condition monitoring and prognosis in industry has become increasingly important through the IoT revolution as the cost and safety benefits of a successfully implemented system are indisputable. A great deal of research has attempted to accurately model degradation in asset operations. In modelling chute lining wear, a hybrid method, combining Bayesian hierarchical modelling, and learning from historical sensors through clustering is proposed. It has been shown that when sufficient data is presented, collaborative learning across assets will improve the performance of predictions. The modelling approach provides a framework that can be applied to other systems or industries in which similar degradation processes occur.

Focusing on a simple solution, easy to be developed and used by companies, we develop a novel method using the Bayesian Hierarchical Model (BHM) to estimate a linear regression model for the degradation of each sensor. Then, we apply a cluster analysis to identify a group of degradation patterns of the sensor to use later as an informative prior for a new sensor. Our new method considers the current available data of the new sensor to match the previously identified patterns and builds a model specific to the sensor. Also, the proposed method enables us to determine whether a sensor in a specific location is showing anomalous behaviour by comparing with other sensors in the chute.

The challenges are that degradation data for some sensors do not follow a linear trend; there are data available for only 20 sensors; and the only information available to be included in a model is time and thickness. Extra information which could describe the factors affecting wear, such as tonnage, angle of impact, and the ore properties is not available.

2. DATA AND METHODS

2.1. Data Description

The data includes raw telemetry for liner thickness readings from 968 sensors across 17 assets at 9 mine sites. This raw data is noisy and features non-monotonous decreases, irregular observation intervals, sensor failures, outliers, corrections, increasing values, and negative values. Only 3 chutes at 2 sites have historical data showing a liner thickness of 0mm (RUL-0), indicating the end of the liner's useful life. Upon replacement, these sensors are decommissioned or replaced. Table 1 presents the number of RUL-0 sensors. Full wear profile data is vital for training the model to accurately describe the wear process. Hopper 1 contains the most complete RUL-0 data from January 2020 to February 2022, is the focus of this study.

2.2. Data preparation

A SQL script was written to extract raw data from a data warehouse. This involved creating joins between 13 tables, producing a table that contains Site, Asset, Sensor, Liner material, Liner thickness, Timestamp, Commission date and Decommission date as columns. This table was loaded into RStudio (RStudio Team, 2020) as a data frame and all subsequent processing and analysis is conducted in R (version 4.1.2) (R Core Team, 2022) using the RStudio IDE.

Table 1. Total sensor counts and RUL-0 sensor counts for assets which contain historical data.

Asset	No. Sensors	No. RUL-0 Sensors
Hopper 1	88	20
Hopper 2	43	7
Train Loadout	101	11

To manage incorrect measurements and ensure monotonic decreasing sensor values. A filtering algorithm was developed which applied the following business rules in parallel:

- Remove all invalid points. This includes; all thickness measurements greater than the initial liner thickness, less than zero or NULL values.
- Remove large changes ($> \ (thickness_i \div 10) + 1$) in thickness, unless the elapsed time is greater than a user-defined minimum period (default 5 days). This removes spurious measurements and allows for the wear that may occur if a sensor is offline for an extended period of time.
- Remove any increases in thickness. The thickness of the wear plate can not increase.
- Liner thickness can remain constant over extended periods of time and therefore there are duplicate measurements. We are only interested in wear changes and the current value. Remove duplicate values; only keep the first recorded point (in time) at each recorded thickness.

A before and after of the filtering process is shown in Figure 2.

Modelling and evaluation requires the sensors which have seen their liner reach the end of its remaining useful life. Therefore the data set was filtered to contain only RUL-0 sensors. A RUL-0 sensor is any sensor that satisfies the following rules:

- Has at least one observation of 0mm thickness or below.
- Has at least 5 recorded thickness values.
- Has monotonically decreasing thickness (ie. no change in thickness greater than 20% of initial thickness)

The resulting clean data, which is a set of 20 sensors in the chute was randomly split into a 70% historical and 30% active set, meaning there is data for 14 sensors in the historical set and data for 6 sensors in the active set. Each of the active sensors data sets are further split into 10% subsets of observed data to assess each models prediction error over different levels of wear.

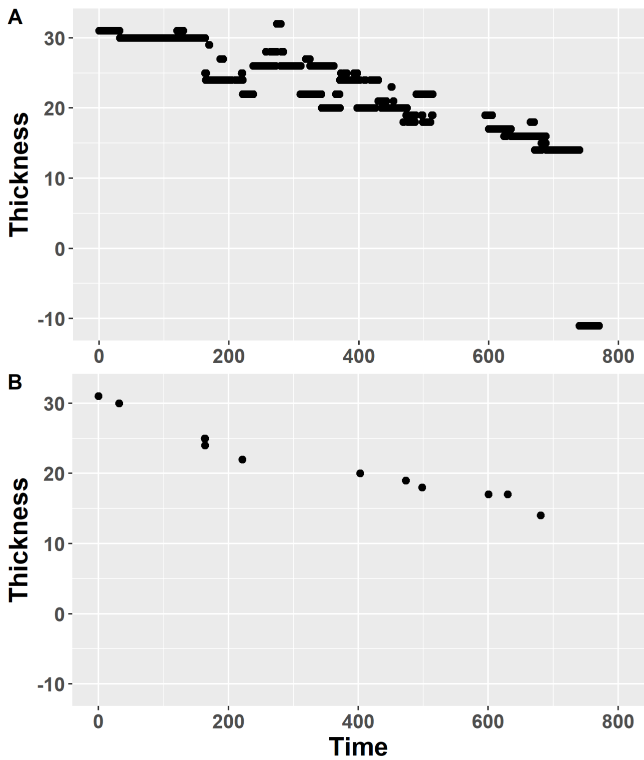

Figure 2. Thickness data over time for a single sensor before (A) and after (B) filtering script has been applied.

2.3. Methods

We use Bayesian Hierarchical modeling (BHM) with clustering to incorporate historical data from sensors. Many sensors can be in the same asset, so we do not assume that their degradation is independent. Instead, we estimate a linear regression for each sensor and use a hierarchical prior to model their dependence. We create clusters based on the estimated degradation parameters (intercept, which is related to the size of the sensor before any degradation, and slope, which is related to the current wear rate) to use as a prior for predicting the wear in new liners. Figure 3 shows the steps in the modeling process. First, we prepare the data and then split it into "historical" and "active" streams. The historical stream is used to create the clusters, while the active stream contains the sensor (or group of sensors) we are interested in. The active stream is further divided into 9 subsets of observed data. We use parameter estimates from the historical data to perform affinity propagation clustering. Each cluster is assigned a characteristic value (exemplar), which is used as a prior to the BHM.

Then, we have

1. Model 1: A BHM using the parameters of the closest related sensor from the 'historical' set. See Section 2.4.1 for more details
2. Model 2: An informed Bayesian model which uses the parameters of Model 1 as prior information and the current observed data for the sensor being modelled. See Section 2.4.3 for more details

The results of these models are compared with a linear regression for an individual sensor (the model commonly used in industry).

2.4. Bayesian modelling

Models are created in RStan (Stan Development Team, 2022), a probabilistic programming language for full Bayesian statistical inference with Markov Chain Monte Carlo.

2.4.1. Parameter estimation from historical sensors

A hierarchical Bayesian linear regression model was created to extract the parameter estimates for each sensor in the historical set. The likelihood function is defined as:

Let $\mathbf{Y} = \{Y_1, \ldots, Y_n\}$ be a random vector of response variables, and $\mathbf{y} = \{y_1, \ldots, y_n\}$ its n observed values.

Consider a design matrix $\mathbf{x} = \{\mathbf{1}, \mathbf{x}_1, \ldots, \mathbf{x}_k\}$, where $\mathbf{1}$ is a $n \times 1$ vector of ones, and $\mathbf{x}_j = \{x_{j1}, \ldots, x_{jn}\}^T$ is a $n \times 1$ vector of the j^{th} covariate (feature/independent variable). That is, $\mathbf{x}$ is a $n \times (k+1)$ matrix. Also, define the i^{th} row of the matrix as $\mathbf{x}_i = \{1, x_{1i}, \ldots, x_{ki}\}$.

Finally, consider $\boldsymbol{\theta} = \{\boldsymbol{\beta}, \tau\}$, where $\boldsymbol{\beta} = \{\beta_0, \beta_1, \ldots, \beta_k\}$ is a $1 \times (k+1)$ vector of linear coefficients, $\tau > 0$ is the precision parameter ($\sigma^2 = \tau^{-1}$) of a Normal distribution, and $\boldsymbol{\theta}$ is the vector of parameters.

The linear regression model is described as

$$\mathrm{E}(Y_i \mid \mathbf{x}_i, \boldsymbol{\theta}) = \boldsymbol{\beta}\mathbf{x}_i \tag{1}$$

and,

$$Y_i \mid \mathbf{x}_i, \boldsymbol{\theta} \sim \mathrm{N}(\mu_i = \boldsymbol{\beta}\mathbf{x}_i, \tau^{-1}) \tag{2}$$

where $\mathrm{N}(\mu_i, \tau^{-1})$ is the Normal distribution with mean μ_i and variance τ^{-1}, and $Y_i \mid \mathbf{x}_i, \boldsymbol{\theta}$ is independent of $Y_\ell \mid \mathbf{x}_\ell, \boldsymbol{\theta}$, $\forall i \neq \ell$.

For the hierarchical prior, we considered

$$\begin{aligned}
\boldsymbol{\mu} &\sim \mathrm{N}(0, \boldsymbol{I}), && (3)\\
\alpha &\sim \mathrm{Gamma}(0.01, 0.1) && (4)\\
\gamma &\sim \mathrm{Gamma}(0.01, 0.1) && (5)\\
\tau \mid \alpha, \gamma &\sim \mathrm{Gamma}\left(\alpha, \gamma\right), && (6)\\
\boldsymbol{\beta} \mid \tau, \mu, \Lambda &\sim \mathrm{N}(\boldsymbol{\mu}, (\tau\boldsymbol{\Lambda})^{-1}), && (7)
\end{aligned}$$

where $\boldsymbol{I}$ is a 2×2 identity matrix, and $\boldsymbol{\Lambda} = 1000\boldsymbol{I}$. Note that, $\mathrm{E}(\alpha) = \mathrm{E}(\gamma) = 0.001$, and $\mathrm{E}(\tau \mid \alpha, \gamma) = 0.000001$. Similarly, $\mathrm{E}(\beta_0 \mid \tau, \mu, \Lambda) = \mathrm{E}(\beta_1 \mid \tau, \mu, \Lambda) = 0$, and $\mathrm{Var}(\beta_0 \mid \tau, \mu, \Lambda) = \mathrm{Var}(\beta_1 \mid \tau, \mu, \Lambda) = 1000$. That is, we are considering diffuse priors for the parameters models

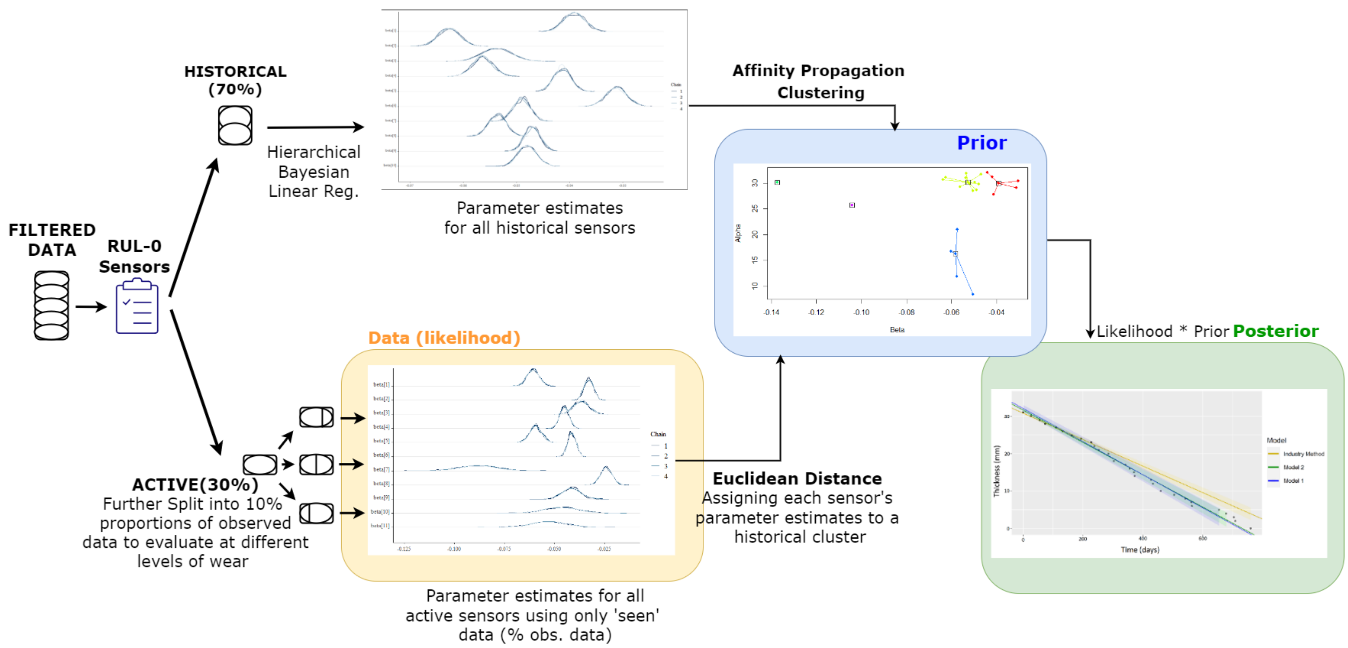

Figure 3. An overview of the approach used, depicting how data flows through filtering, clustering and modelling stages.

$\{\boldsymbol{\beta}, \tau\}$.

Sampling was initialised with 4 chains, a max sampling length of 2000 iterations and a warm up period of 1000 iterations. Utilising parallel processing, and no-u-turn (NUTS) sampling, the Markov Chain converged and a sample from the posterior distribution of the parameters was obtained.

2.4.2. Clustering the historical parameter estimates

As known, there is a relationship between different sensors based on their location in the chute. Sensors co-located on the same liner are expected to wear at similar rates whereas sensor at different ends of the chute will have different wear pattern. This suggests a spatial model would be useful. However, for industry, we need a simple and easy-to-interpret model to support the estimate of degradation prediction to unseen liners. To explore this spatial relationship, we have opted to first build a clustering analysis from historical data. The clustering analysis is used to group sensors for a liner with similar degradation patterns. Using the beginning of life of the new sensor (unseen liners), we find the historical cluster most similar to it. From the selected cluster, we extract the degradation parameters and use them to produce informative priors for the liner wear prediction.

The posterior means of the parameter estimates (β_0, β_1) for each sensor were standardised so that values represent the number of standard deviations above or below the mean. These values are then clustered using the affinity propagation method with Laplace kernel hyperparameterisation (Micchelli, 1984; Fitzgerald, Micchelli, & Pinkus, 1995). The affinity propagation clustering method was selected as the number of clusters is not known and needs to be informed by the data provided. The algorithm returns a number of exemplars, one for each cluster. An exemplar is a representative point for all points within the cluster. Further information on this calculation is in (Frey & Dueck, 2007).

2.4.3. Modelling and clustering active sensors

With Model 1 (Section 2.4.1) parameters estimated, and the clusters defined (Section 2.4.2), the modelling of active sensors is performed in a three stages. The first stage is to build the model for the currently available data of the active sensor. Giving that we are working with a degradation problem, our failure is when the sensor achieved a size of zero. Then, we know it is unlikely to have a failure at the beginning of life. For that reason, we can collect data from the active sensor at the beginning of life to improve the prediction estimation of the useful life. The second stage is to assign the parameter estimates of stage 1 to a cluster (from Model 1 estimates and the historical data). The third stage is to repeat estimation for the active sensors, now including the information from clustering as prior information.

1st stage. 1. Building the model for the active sensors.
In this stage, active sensors are modelled with the same likelihood and prior specification as the model for historical sensors (linear BHM with diffuse priors), and parameter estimates are determined via the same sampling process. See Section 2.4.1.

2nd stage. Assigning the active sensors parameters to a 'historical' cluster and building Model 1.

The active sensor parameter estimates are standardised and assigned to a 'historical' cluster by taking the minimum Euclidean distance to the exemplars. The intercept and slope parameters of the exemplar of the cluster to which the active sensor is assigned are converted back to the original scale and serve as the fixed parameter values for a linear regression for Model 1 (the closest related sensor). This model reciprocates the wear profile of an exemplar as determined by clustering. Note that, for Model 1, the inference is made based only on these parameters. There is no new data added from the active sensor.

3rd stage. Remodelling the active sensors with results from clustering as prior information to build Model 2.

For Model 2, the exemplar values of the cluster to which the active sensor is assigned, are converted back to the original scale and used as prior information in a Bayesian linear regression with the same likelihood function specified in the previous models. The specified priors are $\mathrm{N}(\beta_0, 1)$ and $\mathrm{N}(\beta_1, 0.001)$ where β_0 and β_1 in this case, are the intercept and slope parameters of the cluster exemplar. With 4 chains, a max sampling length of 2000 iterations and a warm up period of 1000 iterations, utilising parallel processing and no-u-turn (NUTS) sampling, the Markov Chain converged and a sample from the posterior distribution of the parameters was obtained for each active sensor.

2.4.4. Model evaluation

The two Bayesian models and linear regression model are evaluated using 5 random subsets of historical and active sensors, further evaluated at each 10% interval of observed data. For each subset and observed/unobserved proportion, mean absolute prediction error in time at 0mm thickness is used as the evaluation metric.

3. RESULTS AND DISCUSSION

3.1. Clustering

The results from clustering are shown in Figure 4, This figure shows that for this chute, the historical sensor data can be separated into 4 clusters, as determined by the affinity propagation method. The red cluster contains the faster wearing sensors with slope parameters between -0.055 and -0.065 (mm/day). The green cluster contains the slower wearing sensors with slope parameters between -0.045 and -0.035 (mm/day). The blue and purple clusters represent sensors that are wearing at a moderate rate between -0.056 and -0.045 (mm/day) with the purple cluster separated by its intercept parameter with values less than half of all the other sensors, identifying an anomaly in the sensor behaviour. Information about how the sensors are wearing in space can be learned by

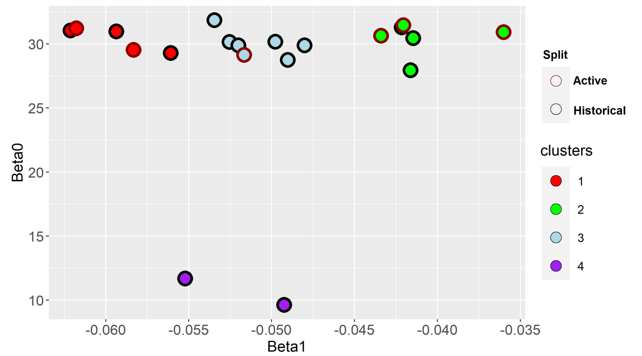

Figure 4. Clustering results after active sensors (using 30% observed data) have been assigned to a cluster. The x-axis represents the posterior mean of the β_1 estimate and the y-axis represents the posterior mean of the β_0 estimate. Each point represents both of the parameter estimates for a single sensor, colour-coded by their assigned cluster. All sensors inside a single chute are shown on the same figure. The colour of the outline of each point identifies whether the sensor belongs to the historical (black) or active (red) set.

relating these clustering results back to the sensors position on the chute.

3.2. Data quality/anomaly detection

The clustering results from the linear model using the historical data are mapped to a fold-out schematic diagram of Hopper 1. The results are presented in Figure 5.

Analysis of Figure 5 suggests a partition between the fast-wearing sections (cluster 1 slope parameter value = -0.059) and the comparatively slower wearing sections (clusters 2 and 3 slope parameter values = -0.041 and -0.05) of the asset. These intermediary results can provide value to plant operators without having to manually inspect the chute. This approach can be used to scan many sensors at once and identify areas of high or low wear and make informed operational decisions. In this case, using the direction of ore flow as reference; the faster wearing areas are towards the back of the chute. In addition to this, operators may be able to identify failed sensors or anomalies in the data. For example, cluster 4 displays an anomaly in the value of intercept parameter (11.67). In consultation with subject matter experts, it was revealed that this sensor was installed on a partially worn chute. Had any sensors recorded a dramatic variation in the slope parameter, either small or large, this could indicate a failure mode in the sensor in that it may have stopped recording or physically broken. These findings directly achieve goal (2) of the analysis by providing operators a low cognitive load method for visualisation. Enabling them to quickly inspect all sensors within the chute for anomalous behaviour and make informed operational decisions,

These results also confirm that spatial factors have some influ-

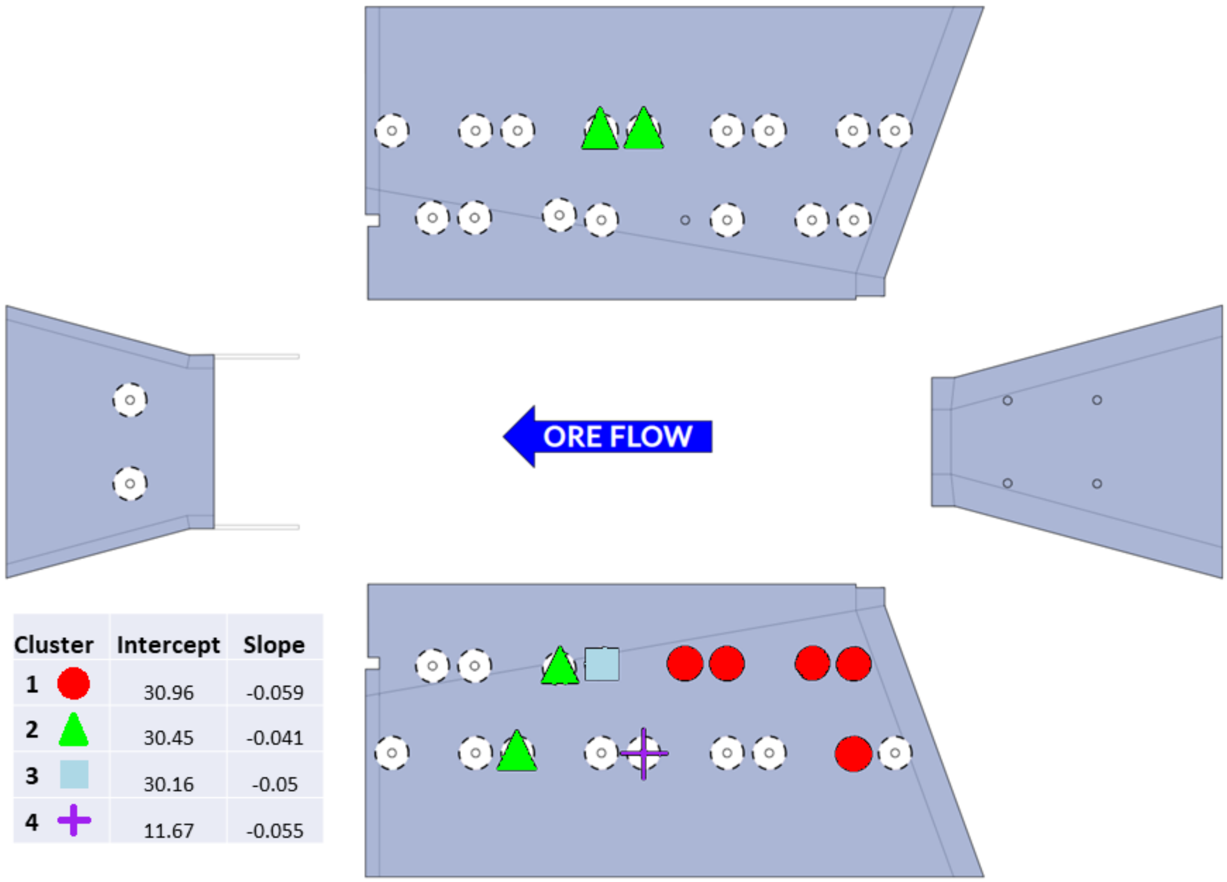

Figure 5. Clustering results related back to a fold-out schematic diagram of the chute interior (Hopper 1). Each polygon represents an interior panel of the asset. White icons represent non-RUL-0 sensors while the coloured icons represent sensors from the historic set with their respective clusters assigned.

ence on the wear rate of liners, as points which wear at similar rates tend to be clustered together. This finding agrees with Tobler's law that nearby things tend to be more similar than those far apart in both space and time (Tobler, 1970). Therefore confirming that positional information is important and should be included in the model.

3.3. Predictions for one sensor

In addressing goal (1) of this analysis, each model is able to predict a point in time when a specified sensor will reach a predetermined thickness value. However, only Model 2 is able to produce an estimate for uncertainty based on a 95% credible interval. Therefore Model 2 is the only model that achieves the goals of the analysis. We are unable to construct credible intervals for Model 1 as the parameter values are considered to be a fixed value, obtained directly from an exemplar of the historical set of sensors. Whilst potential methods such as confidence intervals and bootstrap sampling exist to measure the uncertainties of these models, their interpretation is not probabilistic in the parameter space, and therefore not as easy to interpret for engineers (Castle, Ham, Hodkiewicz, & Polpo, 2020).

Each model's fit to the unseen data and prediction for end of life can be visualized for each sensor. This is illustrated in Figure 6. The end of life is based on the point on the x axis where $y = 0$. We can see that only the future points at $y = 0$ lies within the credible intervals for Model 2 (green). The lines for Model 1 and Model 2 are very close and their estimates for end of life are only 43 days apart. On the other hand the linear regression sits significantly away from the data and over-predicts the remaining useful life substantially. For this sensor the $y = 0$ error of the linear regression is 97 days, while the errors for models 1 and 2 are 25 and 18 days respectively. While the data used for all 3 models is the same, Model 1 and Model 2 include the prior spatial, temporal, and engineering information from historical sensors. The predic-

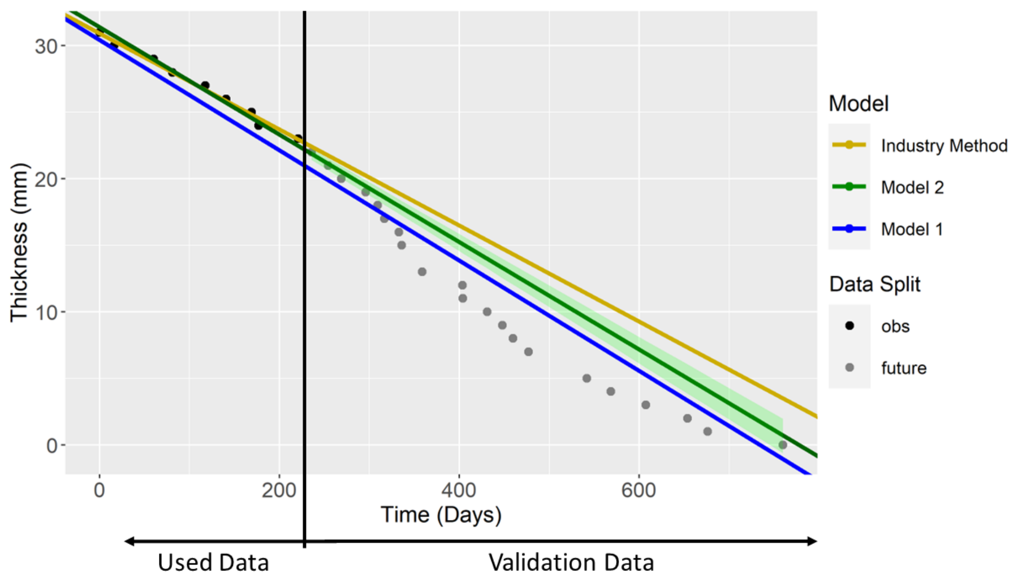

Figure 6. Predictions for all 3 models for a single sensor with 30% observed data of the active sensor. The linear regression model which is representative of the current industry method is depicted by the yellow line. Model 1 is representative of the closest related exemplar obtained from clustering and is depicted as the blue line. The green line depicts Model 2, which is representative of the prior informed model. Its 95% credible interval is shaded in green. The black points indicate the observed data of the active sensor. The grey points represent the true values that the sensor has recorded after the data was split. The grey points are not included in the model. The split is indicated by the vertical black line.

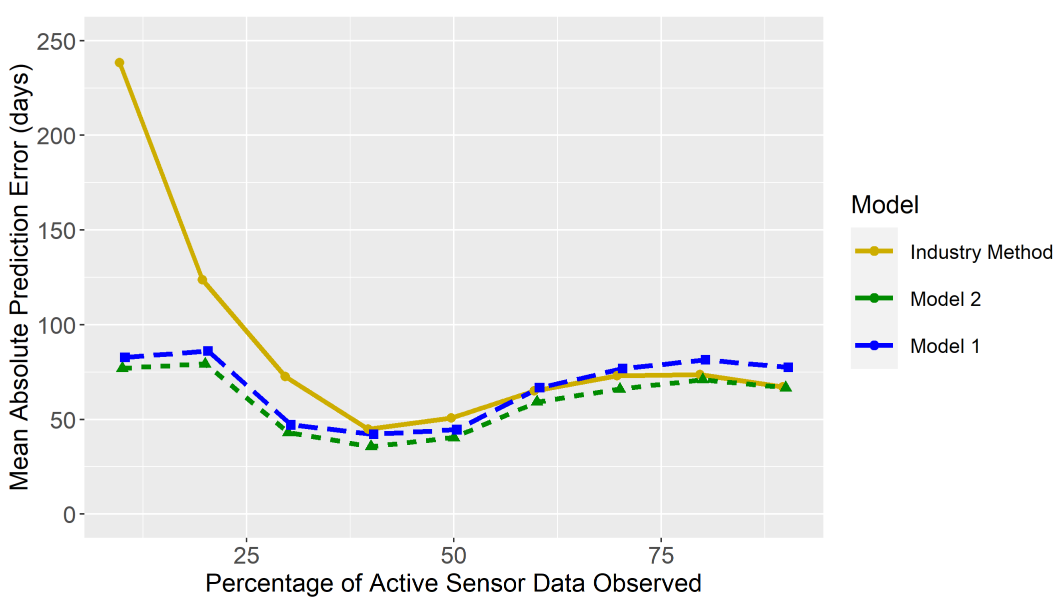

Figure 7. 5-subset mean absolute error at each 10% interval of observed data from the active set of Hopper 1. Mean prediction error is measured in absolute days from the ground truth of six test sensors across 5 different subsets.

tions for Model 1 and Model 2 are a clear improvement to the industry method for this sensor at this specific stage of wear, and indicate that the inclusion of prior information is beneficial. The performance of all models were evaluated in Section 3.4

3.4. Model evaluation

Figure 7 illustrates model errors across 5 different sets of 'historical' and 'active' sensors using the mean absolute error of the prediction at 0mm thickness as the evaluation metric for all models. We note that with limited access to test data the linear model performs poorly. As expected with more observed data, there is a general decrease in the prediction error. This trend is observed for models built on 10-50% of observed data. After which, the error stabilises.

Across all levels of observed data, Model 1 closely follows the trend of Model 2, with slightly improved predictions at every split of observed data. At 60% of observed data, the predictive performance of Model 1 intersects with the performance of the industry method, and performs slightly worse thereafter. Model 1 consistently outperforms both of the other models, for models built on 10-80% of observed active sensor data. At 90% of observed data the difference between the industry method and Model 2 is as expected, negligible. This is due to the fact that with more observed data, Model 2 is able to put more weight on the observed data, eventually converging with the industry model. Overall, Model 2 is the best performing model due to its ability to incorporate knowledge from historical sensors as well as the information observed from the wear from the sensor being modelled.

3.5. Limitations

An initial investigation into the literature and factors that affect wear identified that cumulative throughput (tonnes) is a much better explanatory variable than time (days). Using time as a explanatory variable is problematic as it is affected by operational decisions such as shutdowns or change in throughput. These directly influence the wear rate and can make a linear function redundant. The broader issue is that the companies who utilise services from service providers are unwilling to share their sensitive data with the service providers, even though they could benefit from an improved product. This could be solved by establishing a trusted body that can govern the agreements between companies and service providers to ensure that shared data remains safe and is not used for anything outside the intended purpose of the agreement.

3.6. Stan model efficiency

Stan uses a C++ compiler to compile the probabilistic written code meaning that once code has compiled, it executes relatively fast, however compilation times can be slow. The combined warm up and sampling times for the entire modelling process using 14 historical sensors and 6 active sensors is presented in Figure 8. There is a general increase in the computation time as more data is observed for the active sensors. The mean difference however, from 10% to 90% is 8.2 seconds and can therefore be considered negligible, Approximately 50% of the computation time is associated with the modelling of the historical sensors. In the real world, the historical sensor model does not need to be reconstructed each time we wish to inspect an active sensor/s. It can be saved, and therefore updated and loaded only when required. This finding confirms that the modelling process has the potential to be scaled up across many assets and sensors without impeding on the usability of the system in terms of computation time.

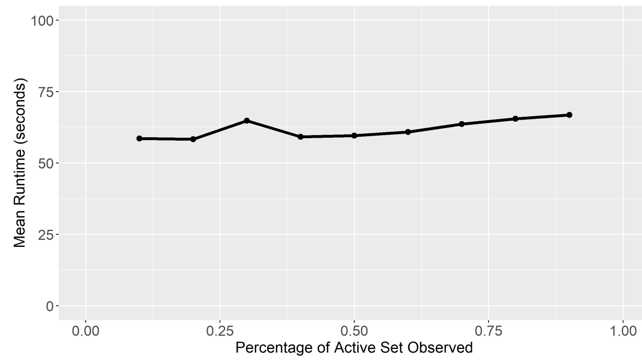

Figure 8. The mean run time of the entire modelling process using 14 historical sensors and 6 active sensors across all 9 splits of observed active sensor data. The mean is calculated across 5 different sets of historical and active sensors.

4. CONCLUSION

This paper demonstrates the successful development and testing of a remaining useful life prediction algorithm for chute liners. The pipeline uses data from historical sensors in a hierarchical Bayesian framework to capture the spatial, temporal and engineering factors contributing to liner wear and apply this knowledge to active sensors through clustering. The pipeline tested the performance of 2 models against the current method used in industry at different stages of wear. The results found that utilising information from historical sensors generally results in improved predictions, especially when less than 50% of wear has been observed, In addition, the predictions from Model 2 are given with 95% credible intervals. This presents an advantage to plant operators as they gain more oversight for scheduling planned maintenance activities. The clustering component of the pipeline also provides a way for plant operators and engineers to automatically scan a large number of sensors with a low cognitive load to identify anomalies in sensor behavior and identify sections of the chute experiencing high or low wear, compared to other sensors in the chute.

However before this system can be implemented in industry, when the quantity of data for other chutes becomes sufficient. Additional testing can be performed to confirm the viability of this pipeline for modelling wear in other chutes.

4.1. Future work

4.1.1. Including throughput data

On investigation of the profiles of wear for some sensors within each asset, it appears that at some points in time, many sensors deviate from a linear trajectory of wear. Possible reasons for this may be that when the liner reaches a certain thickness, its material properties change, resulting in a change in wear profile. Or perhaps more likely, operational decisions on site, such as shutdowns or changes in throughput may be the reason for the sudden change in wear profile across collections of sensors. For this reason it is proposed that cumulative throughput has higher correlation with the thickness than time.

4.1.2. Alternative modelling approaches

As the interpretability of the model is important for industry, the approach used linear functions to estimate wear. More complex functions could be used which address the non-linear patterns in the data for some sensors.

An improvement can be made to the clustering approach, which may influence the accuracy of the models. For Affinity Propagation and calculation of Euclidean distances between exemplars, only the mean is taken from the parameter distributions. We are therefore losing what is considered valuable information, contained in the distribution, about that parameter. To keep this information, hierarchical clustering with Kullback-Leibler divergences between sensor distributions can be used with the 'Ward' and 'silhouette' methods to determine the optimum number of clusters, and partition the sensors into their respective clusters. This has the advantage of incorporating the full posterior of the related historical sensor as prior information when building Model 2, rather than just the mean of the intercept and slope parameters.

ACKNOWLEDGMENT

The authors wish to thank Metso for funding Jacob Van Der Brock's project through the UWA CEED program, and acknowledge Piero Velletri, Alessando Canestrari and the wider WearSense team whose subject expertise, data availability and support for this paper has been invaluable.

REFERENCES

Castle, P., Ham, J., Hodkiewicz, M., & Polpo, A. (2020). Interpretable survival models for predictive maintenance. In *30th European Safety and Reliability Conference and 15th Probabilistic Safety Assessment and Management Conference* (pp. 3392–3399).

Chen, J., & Singpurwalla, N. D. (1996). The notion of "composite reliability" and its hierarchical bayes estimation. *Journal of the American Statistical Association*, *91*(436), 1474–1484.

Chen, W., Biswas, S., Roberts, A., O'Shea, J., & Williams, K. (2017). Abrasion wear resistance of wall lining materials in bins and chutes during iron ore mining. *International Journal of Mineral Processing*, *167*, 42–48.

Fitzgerald, C. H., Micchelli, C. A., & Pinkus, A. (1995). Functions that preserve families of positive semidefinite matrices. *Linear Algebra and its Applications*, *221*, 83–102.

Forsström, D., & Jonsén, P. (2016). Calibration and validation of a large scale abrasive wear model by coupling dem-fem: Local failure prediction from abrasive wear of tipper bodies during unloading of granular material. *Engineering Failure Analysis*, *66*, 274–283.

Frey, B. J., & Dueck, D. (2007). Clustering by passing messages between data points. *Science*, *315*, 972–977. doi: 10.1126/science.1136800

Guérin, F., Barreau, M., Demri, A., Cloupet, S., Hersant, J., & Hambli, R. (2010). Bayesian estimation of degradation model defined by a Wiener Process. In *Mathematical and statistical models and methods in reliability* (pp. 345–357). Springer.

Hawk, J., Wilson, R., Tylczak, J., & Doğan, Ö. (1999). Laboratory abrasive wear tests: investigation of test methods and alloy correlation. *Wear*, *225*, 1031–1042.

Hermann, S., & Ruggeri, F. (2017). Modeling wear in cylinder liners. *Quality and Reliability Engineering International*, *33*(4), 839–851.

Limon, S. M., & Yadav, O. P. (2021). Remaining useful life estimation considering prior accelerated degradation data and bayesian inference for multi-stress operating conditions. *International Journal of Mathematical, Engineering and Management Sciences*, *6*(1), 103.

Malone, G., Hu, X., Clinton, D., & Ore, B. B. I. (2013). Wear property and impact test rig design for comparing wear liners used in transfer chutes. In *CEED Seminar Proceedings* (pp. 1–6).

Micchelli, C. A. (1984). Interpolation of scattered data: distance matrices and conditionally positive definite functions. In *Approximation Theory and Spline Functions* (pp. 143–145). Springer.

Ou, T., & Chen, W. (2022). On accurate prediction of transfer chute wear using a digital wear sensor and discrete element modelling. *Powder Technology*, *407*, 117680.

Padole, P., Joshi, M., & Engineer, J. (2002). Application and implementation of residual life assessment techniques for coal handling plant. In *Nde2002 National Seminar of ISNT* (Vol. 5).

Petrica, M., Badisch, E., & Peinsitt, T. (2013). Abrasive wear mechanisms and their relation to rock properties. *Wear*, *308*(1-2), 86–94.

R Core Team. (2022). R: A language and environment for statistical computing [Computer software manual]. Vienna, Austria. Retrieved from `https://www.R-project.org/`

RStudio Team. (2020). Rstudio: Integrated development environment for R [Computer software manual]. Boston, MA. Retrieved from http://www.rstudio.com/

Stan Development Team. (2022). *RStan: the R interface to Stan.* Retrieved from https://mc-stan.org/ (R package version 2.21.5)

Thurtell, D. (2022). *Australian government resources and Energy Quarterly: December 2024 overview.* Retrieved from https://www.industry.gov.au/ (Accessed: 2024-07-22)

Tobler, W. R. (1970). A computer movie simulating urban growth in the detroit region. *Economic Geography*, *46*(1), 234–240.

van de Schoot, R., Depaoli, S., King, R., Kramer, B., Märtens, K., Tadesse, M. G., … others (2021). Bayesian statistics and modelling. *Nature Reviews Methods Primers*, *1*(1), 1–26.

Vaníček, M., Chamra, S., Jirásko, D., Macháček, J., Vaníček, I., Záleskỳ, J., … Soga, K. (2012). Methods of monitoring of Metro lining in Prague. In *Geotechnical Aspects of Underground Construction in Soft Ground* (pp. 265–274). CRC Press.

Wang, X., Lin, S., Wang, S., He, Z., & Zhang, C. (2016). Remaining useful life prediction based on the Wiener Process for an Aviation axial piston pump. *Chinese Journal of Aeronautics*, *29*(3), 779–788.

Wikle, C. K., Zammit-Mangion, A., & Cressie, N. (2019). *Spatio-temporal statistics with R*. Chapman and Hall/CRC.

Xu, L., Luo, K., & Zhao, Y. (2018). Numerical prediction of wear in SAG mills based on DEM simulations. *Powder Technology*, *329*, 353–363.

Zaidan, M. A., Harrison, R. F., Mills, A. R., & Fleming, P. J. (2015). Bayesian hierarchical models for aerospace gas turbine engine prognostics. *Expert Systems with Applications*, *42*(1), 539–553.