A Flexible Methodology for Uncertainty-Quantified Monitoring of Abrasive Wear in Heavy Machinery Using Neural Networks and Phenomenology-Based Feature Engineering

Thomas Bate[1], Marcos E. Orchard[2], and Nicolás Tagle[3]

[1,2] *Department of Electric Engineering, Universidad de Chile, Santiago, Metropolitan Region, 8370451, Chile*
thomas.bate@ing.uchile.cl
morchard@u.uchile.cl

[3] *Minera Los Pelambres, Santiago, Metropolitan Region, 7550162, Chile*
ntagle@pelambres.cl

ABSTRACT

This paper introduces a cutting-edge methodology for the monitoring of abrasive wear, particularly focusing on SAG (Semi-Autogenous Grinding) mills liners. The lack of a regular inspection regime has historically led to opportunistic and thus, irregular wear measurements that are challenging to integrate into machine learning algorithms for condition-based maintenance. The study unveils a virtual sensor designed to estimate the mill liner's remaining thickness, aiming to offer daily updates and assist the maintenance team in determining the optimal timing for liner replacements without the need for halting operations. This approach is positioned as a strategic response to the critical need for efficient maintenance strategies, addressing the inherent challenges in real-world industrial settings where data quality may be poor and operational realities dominate. A significant aspect of this methodology is its emphasis on uncertainty quantification, vital for informed maintenance decision-making. This novel approach has been successfully applied to SAG mills at Minera Los Pelambres, showcasing its potential for broader applications across scenarios characterized by sporadic data collection. The results showcase an error of ±7.4254 mm of remaining thickness on the validation set, demonstrating the effectiveness of the methodology. The key contributions of this work lie in its ability to utilize low-quality data effectively and its low complexity, reducing barriers to implementing predictive health monitoring (PHM) algorithms. The successful implementation highlights the methodology's adaptability and flexibility, marking a significant advancement in the domain of maintenance strategy for the mining industry.

1. INTRODUCTION

SAG (Semi-Autogenous Grinding) mills are indispensable assets in the mining industry, serving as the cornerstone of ore processing operations. The significance of these mills cannot be overstated, as any downtime incurred due to maintenance activities can translate into substantial financial losses. The costliness of SAG mill stoppages underscores the critical need for effective maintenance strategies to ensure continuous operation and productivity.

Within the maintenance team at Minera Los Pelambres, there arose a strategic initiative aimed at reducing the duration of mill downtime attributed to inspections. To support this endeavor, the concept of the virtual remaining liner sensor was conceived. A virtual sensor, by definition, offers an approximation of a state based on other measurable variables or states, serving as an indirect measurement. In this instance, the objective was to estimate the remaining thickness in millimeters of a specific component of the mill liner online and with daily frequency updates. The overarching goal was to provide the maintenance team with a decision support tool to determine the optimal timing for liner replacement without necessitating mill shutdowns solely for inspections. One of the major challenges in implementing this system was the absence of a regular inspection schedule for the mill. Historically, inspections were performed opportunistically, aligned with planned mill shutdowns. This approach resulted in irregular wear measurements, complicating their utilization in ML algorithms designed to predict wear and determine the remaining liner thickness accurately.

The endeavor to minimize downtime due to maintenance activities has long been a focal point, with condition monitoring tasks representing approximately 13% (Kawahata, Schumacher, & Criss, 2016) of total mill downtime. Traditional approaches to address this challenge, such as discrete element method (DEM) (Wu, Che, & Hao, 2020) simulations, have proven to be exceedingly complex and costly to implement in productive environments, necessitating expensive software and extensive time investments to achieve realistic simulations.

This paper presents a methodology specifically tailored to address the challenges inherent in real-world industry environments, where data quality is often poor, and operational considerations are paramount. By focusing solely on wear monitoring issues prevalent in the industry, this methodology emphasizes the importance of incorporating operational insights into the model design to ensure effective utilization by maintenance teams. The majority of degradation monitoring algorithms are developed using synthetic data or data obtained under suitable acquisition settings, like the measuring method proposed in (Powell & Chandramohan, 2011) (appropriate and stable sampling rates, low measurement error). However, few academic works focus on solving real-world problems where data quality is poor, as the results are naturally less impressive than those generated in studies with high-quality laboratory data, enabling the use of state-of-the-art machine learning algorithms to achieve high precision (Li et al., 2022). The most significant contribution of this work is to provide a methodology that utilizes low-quality data to its fullest potential. The low complexity of the method reduces the barriers currently faced by the industry in implementing PHM algorithms.

This challenge is predominantly practical rather than theoretical, as the methodology was conceived with real-world industrial scenarios in mind, the approach is quite easy to implement. Leveraging neural networks and feature engineering based on phenomenology, the proposed approach is successfully implemented to monitor the liners of SAG Mills at Minera Los Pelambres, providing a decision-support tool for the maintenance team. The methodology is designed to infer deviations from an average wear rate curve, utilizing features that represent stress factors on the mill derived from both historical lining data and real-time mill operation information. The predictive modeling aspect of the methodology employs neural networks, these networks offer accurate inferences of wear progression, allowing for proactive maintenance strategies and the timely identification of potential issues. Complementing the data-driven approach, the methodology incorporates feature engineering grounded in the phenomenology of abrasive wear. This ensures that the monitoring scheme is not solely reliant on learned patterns but also integrates domain knowledge, enhancing interpretability and generalization. A distinctive feature of the methodology is its focus on uncertainty quantification in wear monitoring during the online operation of the model, this is crucial for decision-making, providing the maintenance team with insights into the reliability of wear assessments and facilitating the prioritization of maintenance interventions.

Incorporating uncertainty quantification in industrial monitoring is essential for enhancing intelligent maintenance decision-making. This approach provides a probabilistic perspective on operational data, facilitating a more nuanced understanding of equipment behavior and maintenance needs. The key benefits include:

1. **Enhanced Decision-making**: Uncertainty quantification allows for informed, risk-aware decision-making. By understanding the range of possible outcomes and their probabilities, maintenance teams can make decisions that improve safety, operational efficiency, and financial performance.

2. **Optimized Maintenance Scheduling**: It aids in identifying the most opportune moments for maintenance actions, balancing preventive and corrective strategies. This optimization minimizes operational disruptions and costs while extending equipment lifespan.

3. **Risk Management**: Understanding model uncertainty helps in managing the risks associated with maintenance activities. Identifying high-risk scenarios enables prioritization of critical maintenance interventions, ensuring operational continuity and safety.

4. **Confidence in Predictive Models**: Quantifying uncertainties builds confidence in predictive maintenance models by transparently communicating their reliability. This transparency is crucial for trust among operational staff and stakeholders.

The successful implementation on SAG Mill liners at one of the world's largest copper mine validates the methodology's efficacy. Beyond its application to this specific context, the methodology is highlighted for its flexibility and adaptability to other scenarios. Notably, it accommodates few and irregular measurements of the asset's state, making it applicable in situations where data collection may be limited or sporadic.

However, unlike traditional approaches such as discrete element method (DEM) simulations which are complex and costly to implement, this methodology provides a straightforward and cost-effective alternative. By leveraging low-quality data and emphasizing uncertainty quantification, we address the practical challenges faced in real-world industrial settings. This approach not only simplifies the implementation of PHM algorithms but also ensures robust predictions even with sporadic data collection. Additionally, the incorporation of phenomenology-based feature engineering enhances the interpretability and reliability of the model, setting it apart from purely data-driven methods.

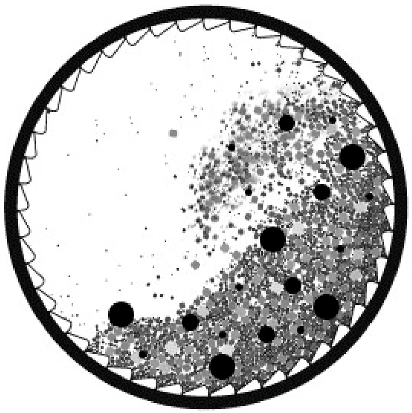

Figure 1. SAG Mill illustration.

2. SAG Mills Background

SAG mills consist of rotating drums containing metal balls that cascade and impact against the mineral (Figure 1), effectively grinding it. The collision between grinding balls and mineral particles fractures and further grinds the mineral, producing finer material. SAG mills are distinguished by their large diameter and short length compared to ball mills. The interior of the mill is lined with lifter plates to lift the material inside, facilitating the cascading material flow for grinding. Figure 1 illustrates this operation. Various lifter configurations and positions line the interior of the mill. Of particular interest is the discharge end cap, where worn or fractured components could lead to ball escape and downstream processing issues.

To assess the condition of the mill liners, a procedure known as a ”faro” is typically conducted. A faro is akin to a radiographic examination of the mill, providing precise measurements of the liner condition. However, these stoppages are costly. There has never been a consistent schedule for the faros, therefore the sampling rate is inconsistent, the only constant measurement made is at the end of the liners life when it is removed from the mill. The virtual sensor developed in this study aims to reduce the frequency of faro inspections, offering online monitoring capabilities to track the remaining liner thickness, particularly focusing on the discharge end cap, where liner failure poses significant operational risks.

The availability of SAG mills is paramount in mineral processing, as every hour of downtime translates to substantial financial losses, valued in thousands of dollars. Optimization of maintenance activities is crucial to minimize mill stoppages, balancing the risk of failure with maintenance requirements.

Given the criticality of mill uptime, any condition monitoring initiative aiding in optimal maintenance scheduling adds significant value. The virtual sensor developed in this study aligns with this objective, providing the maintenance team with decision support tools to optimize faro inspection schedules without necessitating mill shutdowns. This approach not only reduces downtime but also mitigates the risk of operational disruptions downstream.

The current market offers various solutions for monitoring the liners of SAG mills, each with its own set of advantages and challenges. Many of these solutions rely on expensive hardware or require interventions directly on the mill cylinder (Dandotiya, Lundberg, & Wijaya, 2011). However, few effectively leverage historical data to optimize monitoring processes.

One significant challenge is the difficulty of integrating additional monitoring equipment into large-scale machinery already in production. SAG mills are massive industrial units critical to the mineral processing chain, and any modifications or additions to these machines must be carefully implemented to avoid disrupting operations. Installing new monitoring devices often involves intricate engineering work and may require halting production for extended periods, leading to significant downtime and revenue loss for mining companies.

Moreover, the harsh operating conditions within SAG mills present further challenges. These mills operate in environments characterized by high temperatures, dust, and vibrations, which can adversely affect the performance and longevity of monitoring equipment. Ensuring the reliability and durability of monitoring devices under such conditions is essential but often requires additional investments in ruggedized hardware and protective enclosures. Mill liners are located deep within the mill cylinder, necessitating specialized equipment and skilled personnel for installation and maintenance tasks. Any monitoring solution that requires frequent access to the liners may incur significant logistical challenges and operational disruptions.

In light of these difficulties, there is a growing need for innovative monitoring solutions that can leverage existing data infrastructure and minimize disruptions to mill operations. Solutions that harness historical data and employ non-intrusive monitoring techniques offer promise in this regard, providing valuable insights into liner wear patterns while minimizing the need for costly hardware installations and production stoppages.

At Minera Los Pelambres, where three SAG mills—SAG1, SAG2, and SAG3—are operational, the focus is primarily on SAG1 and SAG2 due to their comprehensive data records and identical machinery specifications. The successful replication of the methodology for SAG3 underscores its potential scalability and applicability across multiple mill units, albeit not covered in this document.

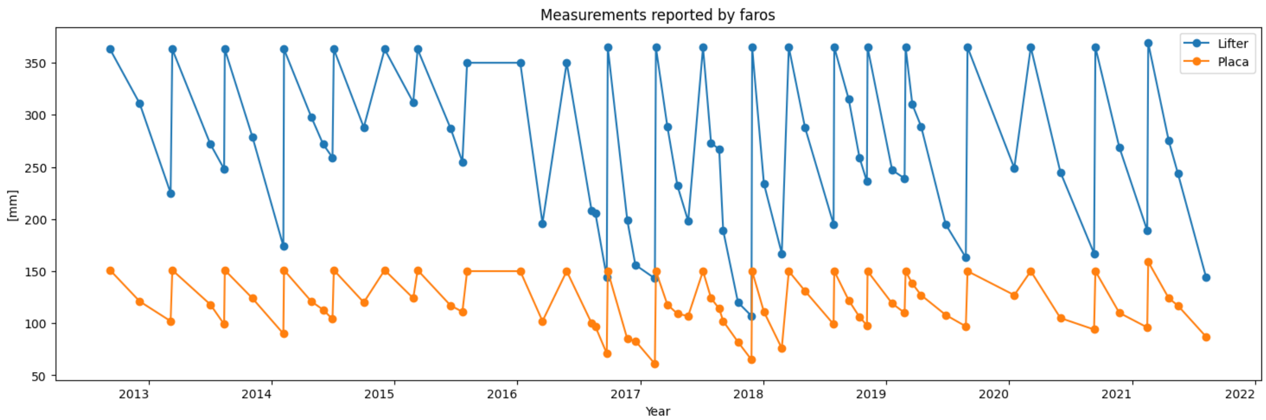

Figure 2. Historic remaining liner mesurements SAG1.

3. Data and Processing

The methodology involves two primary data sources: operational tags related to the mill's functioning and mineralogy, and liner wear measurements reported through inspections. Operational tags, which are time series data from sensors or states, can be extracted at various granularities, they are stored via Pi Systems, a platform for operational data management developed by OSIsoft. It plays a crucial role in this project by serving as the source of operational data. This platform captures, stores, analyzes, and visualizes real-time data from industrial processes. Pi System has its own data cleaning protocols which are critical to ensuring data quality and reliability before analysis, and these protocols were considered in the project's data processing strategy. Given that Pi Systems are commonly used across various industries to manage data, it is important to delve into how data is handled and processed within the methodology. The first step in data processing involves cleaning and imputing these tags. The main issues addressed in the data are outliers, non-numeric values, and missing values. The cleaning process categorizes tags into four types:

- Tags representing percentages: Values above 100 or below 0 are set as NaN (empty value).
- Tags for positive variables with distributions similar to normal: Values below zero and those above the 99th percentile distribution (outliers) are set as NaN.
- Binary variable tags, which include two variables:
  - Rotation direction: Non-numeric values are present, with two relevant states indicating clockwise and counterclockwise rotation. Clockwise is replaced with 1, counterclockwise with -1, and other messages with NaN, to then interpolate them with the closest value.
  - Mill state: Relevant messages indicate whether the mill is stopped (0) or operating (1). Other messages are set to 0.
- Tags that do not require cleaning.

Non-numeric values, often error messages from the tag storage system, are addressed next. Messages indicating a value above/below defined ranges are replaced with the tag's post-cleaning maximum/minimum value. Remaining NaN values are imputed linearly. The tags are cleaned on an hourly basis, including:

- Grinding hardness
- Feed water
- Load cell
- Stator current
- Noise detector
- F80
- Rotation
- Granulometry (of the incoming mineral) 100, 200, 325, 48, 65, 125 Inches
- Filling level
- Solids percentage
- Power
- Pressure
- Noise
- Speed

Following the removal of non-numeric values, data transformation begins, incorporating liner wear measurements. The term 'campaign' refers to the lifespan of the lining. A campaign begins when the lining is installed and ends when it is retired. These measurements are taken with the mill stopped and vary across campaigns, with up to 5 intermediate measurements in some campaigns and a median of 3. The above Figure (2) shows SAG1 grate wear monitoring over time, with two curves representing different grate positions as reported by faros. Points on the curves are measurements, and lines

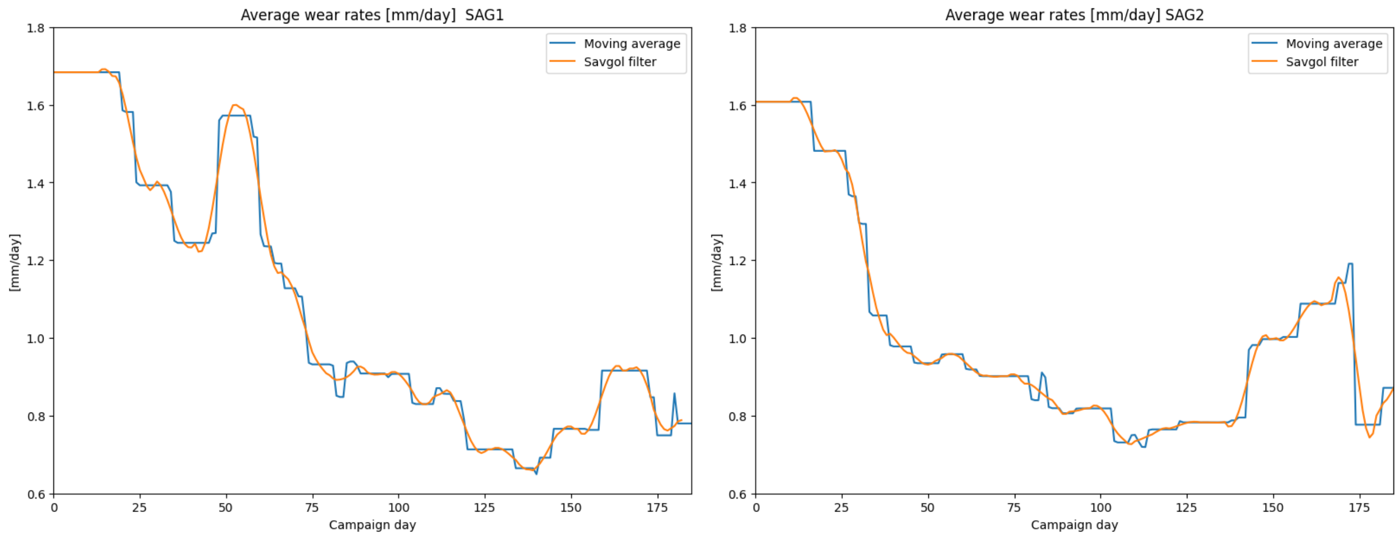

Figure 3. Average wear curves for the lifter position.

represent linear interpolations between points. This highlights the irregularity and varying wear rates in SAG1's history, a pattern consistent across other mills. Almost all the campaigns have a measurement at the start of the lining's lifespan and one near to its end.

Within each campaign, segments between contiguous measurements are defined to calculate consumed thickness and total operational hours, yielding a wear rate in [mm/hour], which is converted to [mm/day]. The objective is to update the liner's remaining millimeters daily. To achieve this, average wear rate curves for SAG1 and SAG2 are calculated (Figure 3). Operational days within each campaign are assigned a wear rate, and using all campaigns, daily wear rates are averaged to produce the curves shown in the next figure. A Savgol filter is applied to obtain the final average wear curve used throughout the study.

The difference between the curves is mainly due to SAG2 receiving more recirculated material, which is less abrasive. Finally, new variables are generated, including accumulated minaral-flows, moving averages, time window dispersions, and others detailed below, aggregated daily and indicating the mill's operational percentage per day.

- Accumulated mineral-flow.
- Clockwise mineral-flow.
- Counter-clockwise mineral-flow.
- Accumulated counter-clockwise flow.
- Accumulated clockwise flow.
- Velocity dispersion over a 72-hour window.
- Accumulated velocity dispersion sum.
- Accumulated power (electric consumption).
- Load cell weight moving average.
- Load cell weight dispersion over a 72-hour window.
- Accumulated load cell weight dispersion sum.
- Noise power moving average.
- Noise power dispersion over a 72-hour window.
- Accumulated noise power dispersion sum.
- Operational day of the campaign.
- Day of the campaign.
- Percentage of clockwise operation time during the campaign.

These enhancements prepare the variables for model input, with accumulations specific to each campaign.

4. Proposed Methodology

The initial decision was to utilize a unified model for both SAG1 and SAG2, justified by the fact that they are essentially the same machinery, albeit with some operational differences. All input variables are aggregated on a daily level, aiming for the model to approximate daily wear of the mill, subsequently accumulating it throughout the campaign for real-time wear monitoring.

After defining the model's input variables, the next step was addressing its output. Resulting from the daily aggregation of data, each data row contains the daily average of tags and created variables, a daily mill utilization percentage, and the interpolated wear rate for that specific day. However, the model does not output this daily wear rate directly; instead, it uses the deviation from the interpolated wear rate compared to the previously mentioned average rates. The creation of the output for the regresor and the histogram of this deviations is

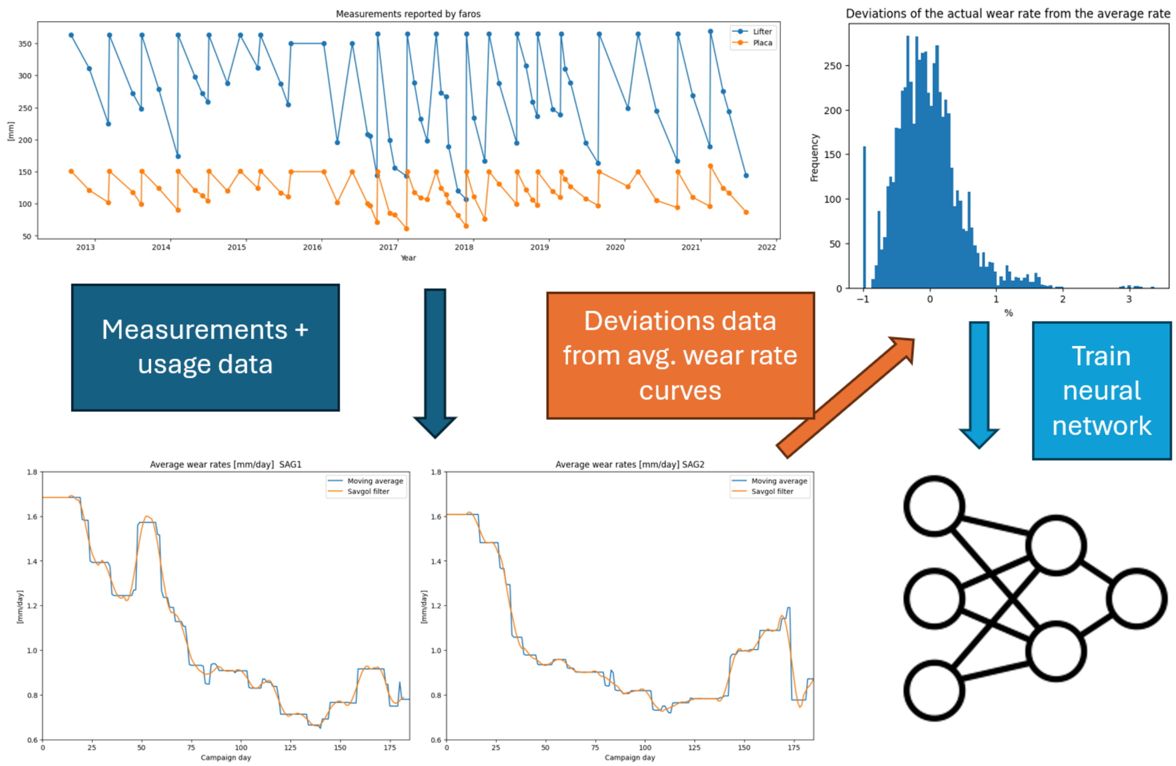

Figure 4. Computation of the deviations from the average wear rate curve.

illustrated in Figure (4). The wear expression in Equation (1) is derived based on the cumulative effect of daily operational conditions on the liner's wear rate. Specifically, we start with the basic principle that the wear on any given day is influenced by both the average wear rate for that day and deviations due to specific operational conditions. Mathematically, this can be expressed as:

$$s(T) = \sum_{i=1}^{T} (\delta_i + \delta_{model}(x_i)) * \alpha_i \tag{1}$$

where T denotes the current day, s represents the lining state measured in millimeters, δ_i is the average wear rate for operational day i, δ_{model} is the model's output, x_i is the model input, and α_i is the day i utilization percentage.

The training set was determined by selecting campaigns with a significant number of measurements to create average wear curves. Considering the irregularity of measurements per campaign and their impact on model training, campaigns with only one intermediate measurement, typically towards the campaign's end, were excluded from the training set (and included in the validation set) due to their limited information contribution about the wear pattern. Thus, the training set comprised 36 campaigns, with a testing set of 8 campaigns chosen for their recency (in order to test the methodology with the most recent campaigns) and lack of more than one measurement.

To train the model, only days from the training campaigns with at least 21 operational hours were used, addressing the distinct data distribution during mill stoppages or startups. The model (in production and testing) was fed all days regardless of operational hours, adjusting outputs by the calculated utilization percentage to prevent unexpected results from low-operation days.

Several types of regressors were tested, including linear regression, decision trees, and support vector machines; however, the neural network was chosen due to its superior interpolation capability and its effectiveness in handling the complexities and variations present in the historical wear data. A Multi-Layer Perceptron (MLP) neural network with three hidden layers and slight dropout was trained using the training set, aiming to minimize the error between the prediction and the actual deviation from the average wear rate curve. The network's performance was then tested against the validation campaigns, focusing on minimizing the projection error, defined as:

$$e_{proy} = \sqrt{\sum_{k=0}^{N} (s^k(T) - s^k_{real}(T))^2} \tag{2}$$

where k indexes the validation campaigns, and $s^k_{real}(T)$ represents the actual liner measurements for that campaign in millimeters.

The network's output represents a percentage deviation, which is then converted to millimeters of liner wear before being adjusted by the daily utilization percentage, ensuring the model accurately reflects operational impact on wear. The distinction between the metric training the neural network and the projection error metric highlights the goal of accurately modeling the actual mill state through the accumulation of neural network results.

5. RESULTS

Recalling from the previous section, the model was developed for the component known as the grate, which, as shown in Figure 2, is monitored at two positions on the grate, named lifter and plate. Therefore, there are two models, one for each position, and the results for both models, which follow the exact procedures described earlier, will be reported. The best model generated for the plate achieved a projection error of 7.4254 mm, whereas the lifter model had an error of 8.701 mm. Below is the table highlighting the projection errors for both models, two example campaigns (Figures 5 and 6) are given in order to illustrate the performance of both models, comparing their results with the faros and with the curve generated by integrating the average wear rates (also weighted by the daily utilization rate), which will serve as a reference. The projection error obtained for those two validation campaigns is also reported.

Table 1. Reported Projection Errors

Model Position	Projection Error (mm)
Plate	7.4254
Lifter	8.701

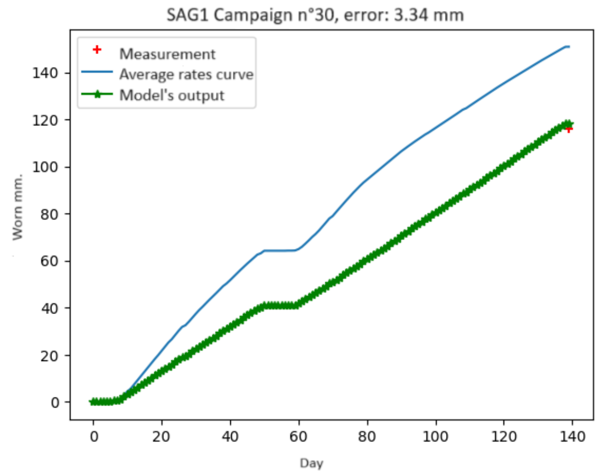

Figure 5. Model evaluation example with campaign 30 (validation) SAG.

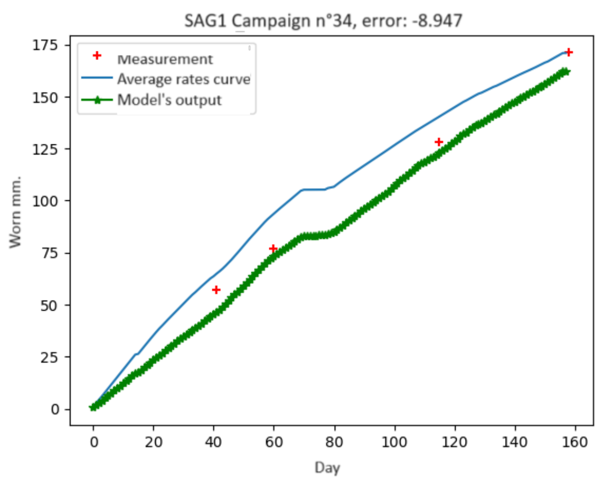

Figure 6. Model evaluation example with campaign 34 (validation) SAG.

5.1. Analysis

The study demonstrates the viability and effectiveness of modeling SAG1 and SAG2 operations jointly. Given that they are identical machines whose variables operate within the same ranges despite differences in their operational patterns, a unified model approach fosters a more robust solution. This robustness stems from training a neural network with data from both mills, offering a larger dataset per model than would be available if two separate models were trained for each mill, also avoiding over-fitting on a single mill's typical operation. This approach not only improves the model's accuracy but also its general applicability across identical machinery.

A significant insight from this work is the advantage of establishing an 'average' operational reference for each asset, as exemplified by the average wear curve against which each mill's wear is calculated. This methodology allows for the development of a model that calculates deviations from 'normal' operation, ensuring predictions remain within reasonable bounds. The histogram of deviations (used as output of the regresor) showed in Figure 4 confirms that there are no significant deviations from the average, highlighting the model's reliability in providing plausible calculations from an operational perspective. This trait is critical for maintaining operator trust in the model, a confidence that could be undermined by implausible model outputs.

The project faced considerable challenges due to the scarcity and poor quality of intermediate wear measurement data. Effective data handling and processing were crucial for maximizing the utility of the available information. Campaigns with only one measurement were primarily useful for validating the model's wear projections and offered limited value for training purposes.

The obtained results are satisfactory, yet they lack a crucial aspect to function online: integrating recent measurements to adjust predictions and quantify uncertainty.

5.2. Online Operation and Uncertainty Quantification

Implementing real-time functionality and accounting for measurement updates in the model necessitates a dynamic approach to incorporate new measurements from inspections, a crucial enhancement given the model's reliance on up-to-date information. A particle filter, a key algorithm in this study, offers an effective tool for state estimation in online settings when incorporating real-time measurements. This Bayesian recursive estimator employs discrete particles to approximate the posterior distribution of the estimated state, making it suitable for online state estimation with available measurements and a system model correlating model states with measurements. It involves initialization, prediction, and correction steps, recursively calculating state estimates.

In the context of this work, a simplified particle filter was implemented as follows:

1. **Initialization**: With an initial measurement always available, particles are sampled from a normal distribution centered on this initial measurement, with variance related to measurement error. Each particle's weight is initialized as $\frac{1}{N}$, where N is the number of particles.

2. **Model Prediction**: Particles follow the model's trajectory, with added noise to introduce variability among the particles.

3. **Measurement Update**: Upon receiving a measurement, the posterior state distribution is calculated using the particles, with new weights computed based on each particle's likelihood given the measurement. If a weight disparity condition is triggered, a resampling step occurs. The process returns to the previous step upon completion.

This particle-based approach generates a probability distribution of the state to be estimated, acknowledging and addressing the inherent uncertainty, thus offering a solution that manages the uncertainty associated with the state estimation process effectively. The variance of the particles in the particle filter was calculated based on the projection error, allowing the filter to produce calibrated uncertainty quantification.

For instance, in campaign 34 of SAG1 (Figure 7), the particle filter had minimal impact due to the model's consistent accuracy. However, in campaign 27 (Figure 8) of SAG2, a significant deviation was corrected by the filter upon the third measurement, thereby improving model performance towards the campaign's end.

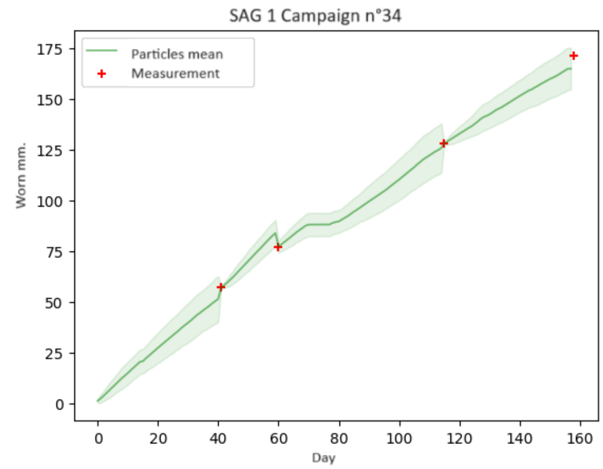

Figure 7. On-line operation of the model with the particle filter, SAG1 campaign 34.

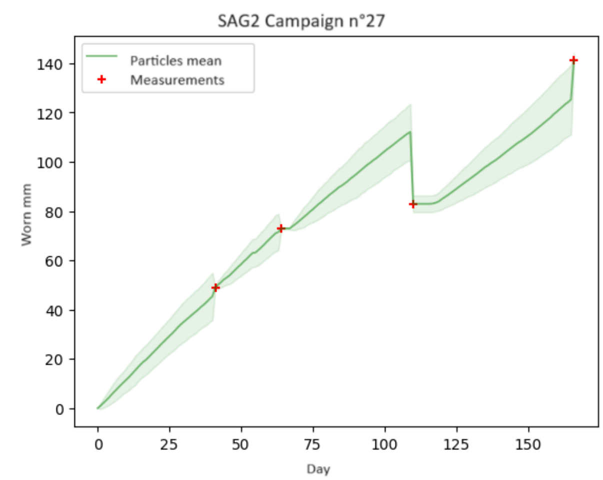

Figure 8. On-line operation of the model with the particle filter, SAG2 campaign 27

6. Conclusions

This paper presents an innovative methodology for abrasive wear monitoring in SAG (Semi-Autogenous Grinding) mills, addressing the challenge of irregular wear measurements due to the lack of a regular inspection regime. The introduction of a virtual sensor aims to estimate the liner's remaining thickness, providing daily updates to assist the maintenance team in scheduling liner replacements efficiently. This method proves critical in enhancing maintenance strategies, particularly in environments where data quality may be compromised and operational realities prevail. A key feature of this approach is the emphasis on uncertainty quantification, which is crucial for informed maintenance decision-making.

The successful application of this methodology to SAG mills at Minera Los Pelambres demonstrates its effectiveness and potential for broader adoption. Achieving an error of ±7.4254 mm of remaining thickness for the plate position and ±8.701 for the lifter in the validation set underscores the models precision. The methodology's ability to utilize low-quality data and its simplicity are among its most valuable contributions, reducing the barriers to implementing predictive health monitoring (PHM) algorithms and marking a significant advancement in maintenance strategies for the mining industry.

ACKNOWLEDGMENT

This work has been partially supported by FONDECYT Chile Grant Nr. 1210031, and the Advanced Center for Electrical and Electronic Engineering, AC3E, Basal Project FB0008, ANID.

REFERENCES

Dandotiya, R., Lundberg, J., & Wijaya, A. R. (2011). Evaluation of abrasive wear measurement devices of mill liners.. Retrieved from `https://api.semanticscholar.org/CorpusID:1596`

Kawahata, K., Schumacher, P., & Criss, K. (2016, 07). Large-scale mine production scheduling optimisation with mill blending constraints at newmont's twin creeks operation:. *Mining Technology*, *125*, 1-5. doi: 10.1080/14749009.2016.1212510

Li, K., Chen, M., Lin, Y., Li, Z., Jia, X., & Li, B. (2022). A novel adversarial domain adaptation transfer learning method for tool wear state prediction. *Knowl. Based Syst.*, *254*, 109537.

Powell, M., & Chandramohan, R. (2011, 01). A structured approach to modelling sag mill liner wear - monitoring wear.

Wu, W., Che, H., & Hao, Q. (2020, 11). Research on non-uniform wear of liner in sag mill. *Processes*, *8*, 1543. doi: 10.3390/pr8121543