Journal of
Manufacturing and Materials Processing

MDPI

Article

A Step beyond Reliability in the Industry 4.0 Era: Operator-Leveraged Manufacturing

Alejandro Muro Belloso [1,*], Kerman López de Calle Etxabe [1], Eider Garate Perez [1,2] and Aitor Arnaiz [1]

[1] Tekniker, Basque Research and Technology Alliance (BRTA), Iñaki Goeanaga, 5, 20600 Eibar, Spain; kerman.lopezdecalle@tekniker.es (K.L.d.C.E.); eider.garate@tekniker.es (E.G.P.); aitor.arnaiz@tekniker.es (A.A.)

[2] Faculty of Informatics, University of the Basque Country (UPV/EHU), Manuel Lardizabal Pasealekua, 20018 Donostia, Spain

* Correspondence: alejandro.muro@tekniker.es

Abstract: Avoiding downtime is one of the major concerns of manufacturing industries. In the era of connected industry, acquiring data has become cheaper than ever; however, turning that data into actionable insights for operators is not always straightforward. In this work, we present a manufacturing scenario involving a circular blade rubber cutting machine, where the goal is to minimize downtime. Historical cutting data are available, and the aim is to provide the machine operators with an intuitive tool that helps them reduce this downtime. This work demonstrates how, in an Industry 4.0 environment, data can be leveraged to minimize downtime. To achieve this, different survival model approaches are compared, a Health Index (HI) is developed, and the model deployment is analysed, highlighting the importance of understanding the model as a dynamic system in which the operator plays a key role.

Keywords: Industry 4.0; reliability; health index; survival models; human feedback; human-centered

check for updates

Citation: Muro Belloso, A.; López de Calle Etxabe, K.; Garate Perez, E; Arnaiz, A. A Step beyond Reliability in the Industry 4.0 Era: Operator-Leveraged Manufacturing. *J. Manuf. Mater. Process.* **2024**, *8*, 215. https://doi.org/10.3390/jmmp8050215

Academic Editor: Jing Shi

Received: 24 July 2024
Revised: 4 September 2024
Accepted: 25 September 2024
Published: 28 September 2024

1. Introduction

Downtime, that is, the period of time during which critical machines or processes are unavailable, has a great impact in manufacturing production. An accurate computation of the downtime–related costs is difficult, but the following expense sources are typically recognised: direct costs, such as labour and material; indirect costs, such as the impact of stopping and restarting the production as well as the non–realized revenue. For instance, ref. [1] surveyed Swedish firms in 2014, and the costs considered could reach up to 23.9% from the total manufacturing cost ratio and 13.3% from planned production time.

At the same time, there is a clear trend in the industry towards increasing the number of installed sensors. According to the survey presented by Plant Engineering in 2021 [2], 36% of the surveyed plants planned to implement or increase investment in sensors and/or remote monitoring. It is expected that the inclusion of these sensors will provide a clear benefit in the following aspects: enable predictive maintenance, insights or/and analytics (according to 66% of the respondents), improve equipment effectiveness (according to 55% of respondents), and utilize data over instinct to improve uptime and meet production goals (according to 53% of the respondents). In such a scenario, data–driven models are effective tools that could turn the data recorded by the sensors into beneficial insights that could improve production and decrease downtime.

This work focuses on the downtime reduction of tread cutting blades installed in a tyre manufacturing plant. There is a lack of literature addressing the specific problem of a circular cutting system used on rubber. Most existing studies on circular saw modelling and wear detection focus on wood sawing [3], with fewer works available on cutting other materials, such as rocks [4] or metal [5]. In general, the works trying to determine the predominant factors on wood cutting and wear formation tend to be validated on a lab set–up, and the results regarding the most prominent factors affecting the cutting/wear formation are ambiguous [3]. At the same time, the works reporting wear evolution tend to show monotonic and almost linear wear increase in wood cutting [6,7] as well as rock [4] and aluminium [5] cutting. In the previous works, the main factor affecting the wear formation is either the material of the workpiece or the physical properties of the cutting tool. Considering the previous, in the work presented by [8] it is suggested that, in an industrial scenario, the most predominant aspect could be considered the material properties of the workpiece.

In industrial scenarios in which aging or wearing estimations are required and failure data is available, survival–type models, such as the Weibull model, that can be used for complex failure datasets [9] are quite extensive. However, given that research suggest that the reliability of a piece of equipment or system is greatly affected by operating conditions [10], those models can be further improved by using exogenous factors or covariates in case there is existing additional information that could be used for a more accurate modelling.

Regarding how to provide operators with asset life–expectancy information, the Health Index (HI) is a popular tool in the industry. HI is a practical approach that combines the results of operating observations, field inspections, and on–site and laboratory testing in order to provide an objective and quantitative index that informs the operator of the overall health of the asset [11,12]. Furthermore, HI is an excellent indicator in reflecting the results of optimal balance among capital investment, cost of the asset, and operating maintenance.

There are numerous works in the literature that discuss data models and how to develop them. However, few explore the aspects related to experimenting with these models when it comes to deploying them in industrial environments. Since the emergence of what is considered the precursor of MLOps [13], where some of the key aspects and bad practices in deploying and maintaining models were identified, the term MLOps has gained great relevance. However, few studies provide evidence of how to instantiate it, especially in an Industry 4.0 environment, as highlighted in [14].

In the current Industry 4.0 context, new research suggests that Industry 4.0, which is leveraged by big–data analytics, artificial intelligence, and digital twins, is limited by its emphasis on promoting the efficiency and flexibility of industry rather than focusing on industrial sustainability and workers' welfare [15]. Therefore, the next step beyond this paradigm requires bringing humans back into the process for collaboration and reintroducing the human touch to manufactured products while simultaneously focusing on sustainable manufacturing [16]. For this to happen, the role of co–intelligence—where humans assist robots by training, explaining, and sustaining, and robots assist humans by amplifying, interacting, and embodying—will be paramount [17].

This work presents an online approximation for reducing a circular cutting blade system's downtime. In doing so, an improved survival modelling approach that accounts for the operation variability is compared against other survival models. The models are built with existing sources of data and are designed to provide the final users, the operators, a Health Index of the cutting system, which empowers the operator while requiring their feedback. Instead of focusing on the improvement of the model with a fixed dataset, this work, situated in between the 4.0 and 5.0 industrial revolutions, provides evidence of the research of models that are already deployed and could benefit from the inclusion of new data sources by introducing the human in the loop and need to consider other aspects, such as retraining.

2. Materials and Methods

2.1. Study Setting

This investigation was conducted within an industrial environment, specifically, a tire plant that routinely cuts different rubber treads, for which it employs a rotating blade. Before initiating this study, the plant relied on a reactive maintenance strategy for managing the blades' lifespan. Under this regimen, blade replacements were only made when the blade became jammed or when defects, such as burrs, were observed on the cut treads.

2.2. Project Framework

The research forms part of the European AI–PROFICIENT project, aimed at integrating Artificial Intelligence (AI) technologies into manufacturing operations. The primary objective within this framework is to develop adaptive algorithms capable of estimating the wear status of critical components like the cutting blade and providing actionable insights to operational staff, while continually learning from the data gathered.

2.3. Development Phases

The project unfolded in two distinct phases:

- **Initial Phase:** During the early phase, pre–existing data collection systems were utilized to gather necessary operational data. This phase focused on leveraging available historical data to begin the algorithm development process.
- **Implementation Phase:** Subsequently, a dedicated platform was developed to facilitate real–time data acquisition from production machinery. This new system enabled a continuous stream of data, enhancing the immediacy and relevance of the information available for algorithm refinement. The deployment of this platform marked a transition to a more dynamic data environment, where algorithms could be applied directly within the operational context of the plant.

Throughout both phases, the integration of AI into the plant's processes was aimed at transitioning from reactive to predictive maintenance strategies, thereby increasing operational reliability and efficiency.

2.4. Data Sources

2.4.1. Original Data—Prior to Deployment

Before the deployment of the platform, the data used for the modelling and analysis came from a combination of handwritten maintenance logs and the signals of the Supervisory Control and Data Acquisition (SCADA) system. Maintenance logs report the maintenance actions carried out during each shift (such as a change of the cutting blade) and were cross–matched with the cutting system signals from the SCADA that included, among other things, the number of cuts and the materials being cut. Knowing the approximate time when the blade was placed and removed, together with the cuts it performed, enabled the identification of the number of cuts achieved by each blade during its usable life. As a result, 227 blades were identified from the historic data, each having an ordered record of the cuts performed. Five of those records are depicted in the following Table 1.

Table 1. Snippet of the original data used in [18], with aggregated cuts for clarity.

bladeID	Cuts
8	58,448
10	93,150
11	117,759
12	95,981
13	24,206

2.4.2. New Data—Post Deployment

During the second part of the project, new measurements and interfaces were installed, which allowed the recording of new information for each blade used in the process. Table 2 displays an example of the new data.

Table 2. Snippet of the new data obtained after new sensors were installed.

bladeID	Cuts	Reason	Jams	waterJetCleaning
288	58496	EOL	0	0
289	83,090	QualityIssue	1	0
290	31,514	EOL	0	0
291	11,439	Jam	1	0
292	12,902	Jam	1	0

where, in the column indicating **reason** for replacement:

- EOL stands for End Of Life and indicates the blade has achieved its maximum number of usable cuts in practice.
- Jam is used to signal that after the blade jammed during the cutting process it was replaced due to it not being suitable for further cuts.
- Quality Issue indicates that the blade was replaced because the cutting process gave results of poor quality.

Furthermore, the columns **Jams** and **waterJetCleaning** further indicate how many times the blade jammed during its lifetime and the amount of cleanings via water jet it received.

2.5. Models

2.5.1. Weibull–Based Reliability Model

Traditional survival/reliability analysis is used to analyse the expected duration of time until a certain event occurs. In this case, that "time" is measured by the number of cuts made by a blade until its replacement. In such an approach, the survival function $S(Cuts)$ gives the probability that a blade is still useful (has not been replaced) after a certain number of cuts. The Weibull distribution is used, which is particularly well suited for modelling lifetime distributions. It is expressed mathematically as:

$$S(Cuts) = \exp\left(-(Cuts/\lambda)^k\right), \tag{1}$$

where λ and k are referred to as *scale* and *shape* parameters, respectively, and are obtained via Maximum Likelihood Estimation (MLE).

This approach allows for a probabilistic characterization of the blades' total cuts, which comes in handy when treating blades as a group but has its downsides when it comes to treating them individually. For instance, that the majority of blades do not reach 50,000 cuts does not mean that one particular blade will not be in need of replacement after 10,000 cuts.

2.5.2. Context–Aware Weibull Models

As we continue to make progress in the development of the project, more sensors are installed and more information on the blades life cycle can be obtained. More precisely, we can utilize information regarding the blade jamming or cleaning actions that had taken place during their lifetime. These actions are recoded together as *interventions*, creating up to three different categorical classes for 0, 1, or 2 (or more) *interventions*. These three classes will be used to further characterize the blades and model them accordingly. That is to say, the Weibull approach will be complemented by considering different models/Weibull distributions according to the *interventions* performed on each blade.

2.5.3. Survival Analysis via XGBoost

For the sake of comparing our approach to the more complex and commonly utilized state–of–the–art approaches, we will be applying Gradient Boosting techniques to the survival analysis via XGBoost (version 2.0.3) [19]. We will be using the Accelerated Failure Time (AFT) model, which assumes the log of the survival time is linearly related to the covariates with a random noise component. This approach allows the use of powerful tree–based learning methods like XGBoost for survival analysis, leveraging its capacity for handling various types of data and complex non–linear relationships within the data.

2.6. Health Index

The creation of a Health Index (HI) requires that we combine complex information regarding the condition of each blade in order to provide a single numerical value as an indication of the overall condition of the asset. The objective is that HI becomes a useful tool for the operator and helps them make decisions based on an estimation of the degradation of the blade.

The modelling approach suggested here follows the one presented in de Calle Etxabe Kerman et al. [8], which is based on survival functions obtained via Weibull fitting. Such approaches reflect the probability of a blade being alive, i.e., usable, after a certain amount of cuts. However, the interpretation of probabilistic information is not as straightforward for the final users (operators or maintenance managers), who are ultimately the ones in charge of the blade replacements. Therefore, it becomes a necessity to adapt the definition of Health Index to match data specifications and to be of use to the people dependent on it. This work proposes a way to turn a survival function into a proper Health Index, whose interpretation is rather straightforward, with values close to 1 meaning the asset is in good condition, whereas values close to 0 mean the asset should be replaced soon. Given the one–to–one relationship between the survival rate and *Cuts*, we could set programmed replacements once the blade manages to reach a certain number of cuts—a threshold of sorts. This is to say that, by fixing the *Cuts* to be made before a blade is replaced, we expect that the percentage of programmed replacements will tend to equate the survival rate of the batch of blades.

Note, however, that to define the best threshold, different costs associated with each type of blade replacement (programmed and unprogrammed) must be taken into account. As the exact quantification of these costs is difficult, mostly due to all the external opportunity costs associated with an unplanned change, various scenarios are compared through simulations.

Simulation

Different simulation scenarios are proposed in this section on account of the relative cost established for the different replacements,

$$\text{Relative Cost} = \frac{\text{Unprogrammed Replacement Cost}}{\text{Programmed Replacement Cost}}\,, \tag{2}$$

where Unprogrammed Replacement Cost is at least the same as Programmed Replacement Cost so that Relative Cost ≥ 1.

Simulations are carried out as illustrated in Figure 1 for a fixed number of Total Cuts and varying Threshold, defined for survival rates ranging from 0.10 to 0.90, with 0.10 steps. Note that this is the same as considering each of the nine 10–quantiles (or deciles), ranging from D_1 to D_9, since the survival function is related to the cumulative distribution function for the Weibull distribution:

$$F(Cuts) = 1 - S(Cuts)\,. \tag{3}$$

The simulation process depicted in Figure 1 can be summarised as follows:

- After initialising necessary simulation parameters, blades are sampled from the fitted Weibull distribution.
- Its *Cuts* are then compared to the Threshold, flagging a blade as an Unprogrammed Replacement if the value is not met. Otherwise, a Programmed Replacement is established.
- Simulation parameters are updated accordingly and a new blade is sampled (with replacement) until Cuts Done reaches the Total Cuts required.

- Finally, after obtaining the total number of Programmed Interventions and Unprogrammed Interventions, the final Relative Cost is computed according to (2):

$$\text{Total Relative Cost} = \text{PI} \cdot 1 + \text{UI} \cdot \text{Relative Cost}\,. \tag{4}$$

Due to the stochastic nature of the simulation, the simulation process is repeated several times and the final cost results are then averaged. Once the most appropriate cuts value is chosen as the Threshold (normally the one that minimises the costs) for *Cuts*, we take the following approach for creating the Health Index. Considering a constant wear–speed, the Health Index decreases linearly from 1 as the blade cuts and reaches 0 after performing the established number of cuts. The blade is changed thereafter.

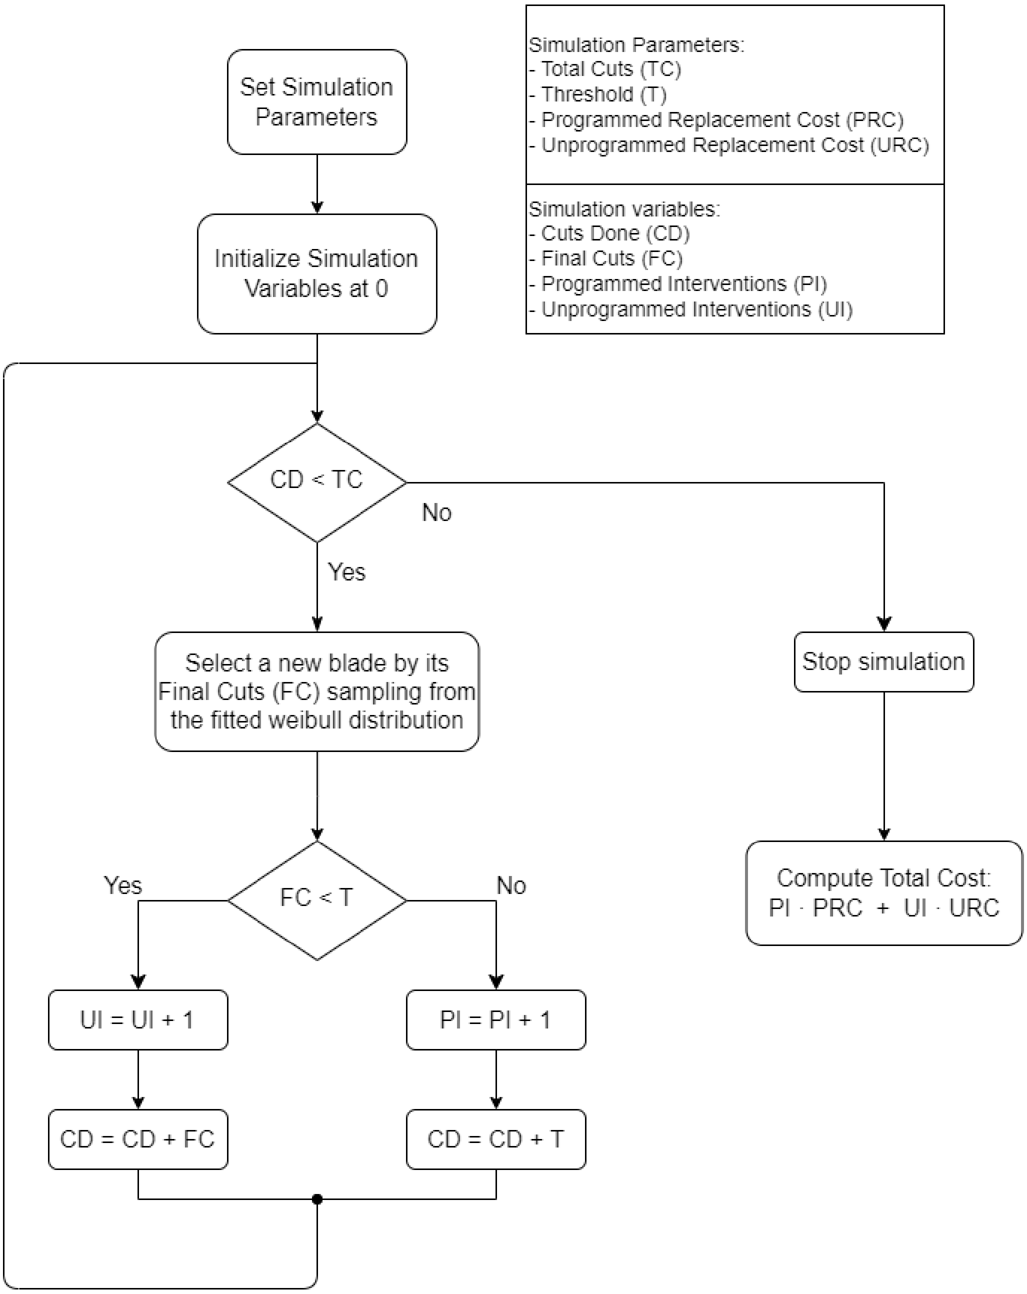

Figure 1. Diagram depicting the simulation used in our reasoning.

We can estimate the Health Index of the blades by considering a constant wear–speed. This way, the Health Index decreases linearly from 1 as the blade cuts and reaches 0 after performing the established number of cuts (the blade is changed thereafter). Problems arise when the blade reaches its end of life before that point (after all, it is modelled based on survival probability) because HI never gets to 0. A work–around solution to this problem may consist of simply including certain uncertainty intervals for the HI.

2.7. Platform Deployment

While developing the algorithmic part, an industrial IoT platform was developed in order to facilitate the acquisition of data from the machines and to ease the deployment of the developed algorithms.

In an attempt to follow best practices for model deployment, MLOps principles were considered. In that regard, special attention was paid to the following aspects:

- Automation: Automation is one of the keys of MLOps. It involves avoiding manual interventions in the deployment and operation processes as much as possible. This not only speeds up workflows but also reduces the risk of human error and enhances consistency across different environments.
- Version control: Code has been versioned using a Git versioning system. This enables greater control over changes and the ability to track and roll back to specific versions if necessary, enhancing the manageability and security of codebases.
- Continuous training: Given that new data are being created every day, the model might no longer act as expected. In that regard, the system can improve over time by retraining; this way, potential drift on data can be tackled without major consequences.
- Infrastructure as a code: For the sake of reproducibility, the deployment of the infrastructure has avoided manual configuration steps as much as possible. This has been possible by using docker containers that enable defining the configurations and software requirements through code, which should be preferred over manual steps written on the documentation.
- Scalability: Scalability has been ensured by using elastic solutions such as cloud–based databases and storage, which can dynamically adjust to the load demands. This is advantageous even if the current model size and data volume are modest, as it prepares the system for future growth without requiring re–engineering. This approach ensures that the infrastructure can handle increases in data input and user demand smoothly, facilitating a seamless scaling–up process when needed.

In addition, some human–centred considerations have been taken into account in order to ensure a conflict–less integration of the system in the plant. In that regard, the following considerations have been taken:

- Minimizing operator overhead: An attempt to require minimal intervention from the operator has been made to avoid solutions that would cause more effort than relief. Nevertheless, it must be kept in mind that, before the deployment of the platform, no warning was given by the cutting system; it would work until failure, leading to a stoppage of the cutting process.
- Transparency/Understandability: The models have been designed to be as interpretable as possible, favouring simple–to–understand models over complex black boxes. This approach helps operators understand and validate the outputs, increasing the confidence when using the system.
- Engagement and feedback loops: Specific HMI's have been installed to provide the operator a means to retrain the algorithms by providing feedback.

The whole architecture of the deployed systems can be seen in Figure 2.

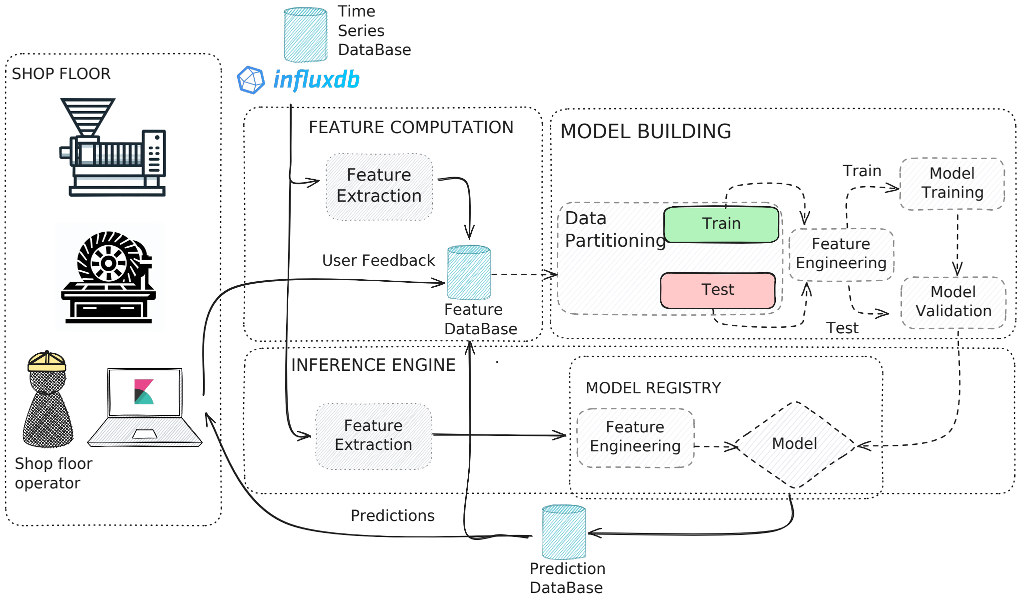

Figure 2. Depiction of the platform used for model deployment.

2.7.1. Human Feedback Interface

Following the human feedback approach suggested by [20], two sources of feedback for retraining were used: Explicit and implicit. The implicit feedback is gathered automatically as it is possible to know the amount of cuts a blade performs and, hence, it is possible to adjust the model outputs based on the new data. Regarding the explicit feedback, it is known that the reasons for changing a blade could be motivated by scheduled maintenance and not due to wearing. In addition, after several conversations with the operators it is possible to identify that some interventions that take place on the cutting system are known to impact the blades' life: when the blades are cleaned by water jets or an excessively thick work piece is being cut, the cutting system sometimes jams and is not able to cut properly. In this context, a new HMI has been installed close the local PC at the cutting system. This HMI (visible in Figure 3) pops up every time the security lid of the cutting system is opened, and it allows the operator to provide additional information regarding the motivation for opening the lid. In this way, it is possible to identify harmful interventions during a blade's life and, in addition, the causes for changing the blade.

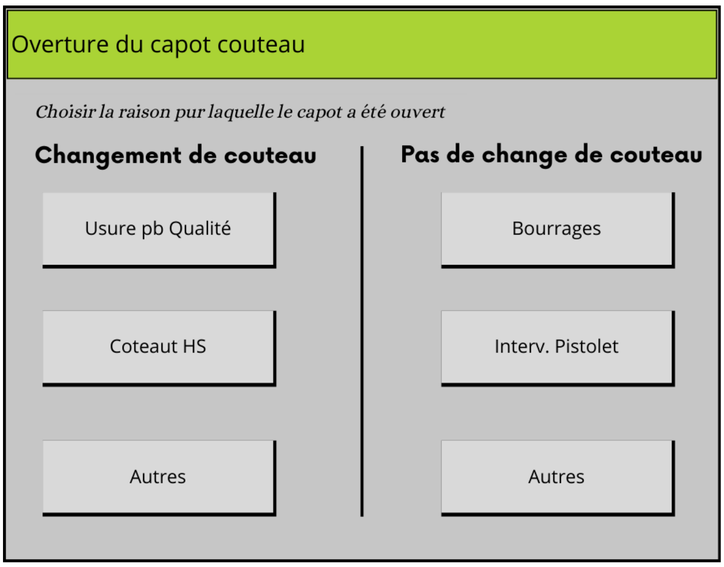

Figure 3. Interface installed at shop floor.

2.7.2. Feature Extraction

The signals obtained from the cutting system and the local PC are sent to the the Influx database. This includes the cut counter (which restarts every time the operator marks a change of blade on the HMI) and boolean signals that indicate the activation of the different buttons on the HMI. Once the signals are stored on the Influx time–series database, it is possible for the rest of the modules of the platform to access, almost in real time, the data gathered on the plant.

The raw readings, however, are not informative enough just by themselves. In that sense, a feature extraction module is needed; this module takes the raw signals since the previous blade replacement (that is, since the current blade was installed) and computes the different interventions (jams or water jets) and the cuts carried out up to the current time.

This processed information can be later used either for updating the features database in case the blade has been changed or to execute model predictions if knowing the Health Index is desired.

2.7.3. Prediction Service

The prediction service takes the features computed by the feature extraction module and, using the model, computes an estimation of the current health of the blade. In addition, it provides estimations of the evolution of the blade's wear for the next cuts. This information is stored on the prediction database, which is queried by a Kibana dashboard (see Figure 4) that can be accessed by the operator or maintenance manager. The dashboard displays the health index as well as the expected evolution of the index for the future.

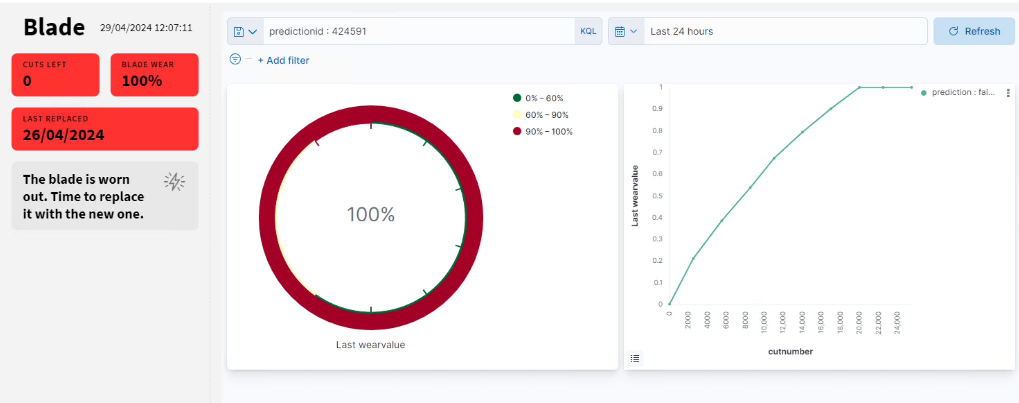

Figure 4. Kibana dashboard used by the operator to monitor blade wear status.

2.7.4. Executions

All the services are governed by a single algorithm that is executed every 30 s. In every execution, the algorithm carries out a different role: if a blade change is detected, it will update the feature table and launch a retraining of the model. If no blade change was detected, the latest model will be used to produce some inference.

3. Results & Discussion

3.1. Prior to Deployment

3.1.1. Modelling

As far as the data and information used to model the blades used in the cutting process, we can only rely on the total number of cuts each blade achieves before replacement. Therefore, no complex model suits our needs and we turn our attention to the more simple yet effective reliability models. More precisely, we use the Weibull distribution to fit the final cuts of our blades. Cuts are then related to the probability that a blade is still usable, meaning that a blade has a probability of reaching a certain number of cuts before "dying" (i.e., being replaced). After properly fitting our historical blade data, the cumulative distribution function, $F(Cuts)$, is related to the survival probability function, $S(Cuts)$, according to

$$F(Cuts) = 1\text{˘}S(Cuts)\,. \tag{5}$$

In turn, quantiles q are related to probabilities as $q = 1 -$ survival probability. Taking a large number of blades, this implies that we expect that a percentage of blades equal to the q quantile reached a number of cuts at least equal to the one with a survival probability of $1 - q$. This insight allows us to define a number of cuts that will serve as a threshold to decide on blade replacement. However, a proper choice remains to be explored.

3.1.2. Health Index

The main objective of our simulations is to choose the quantile (i.e., $1 -$ survival probability) for which a proper threshold of cuts can be obtained from our Weibull fit. This value is chosen by minimizing the total relative cost according to a defined relationship between programmed and unprogrammed costs. However, this relation is quite complex to define in practice and thus it is not available for us. Instead, we will take a conservative approach and choose the 5th decile (D_5 or $q = 0.5$) as it can be related with a reasonable survival probability of 50% and reasonable total relative cost from Figure 5.

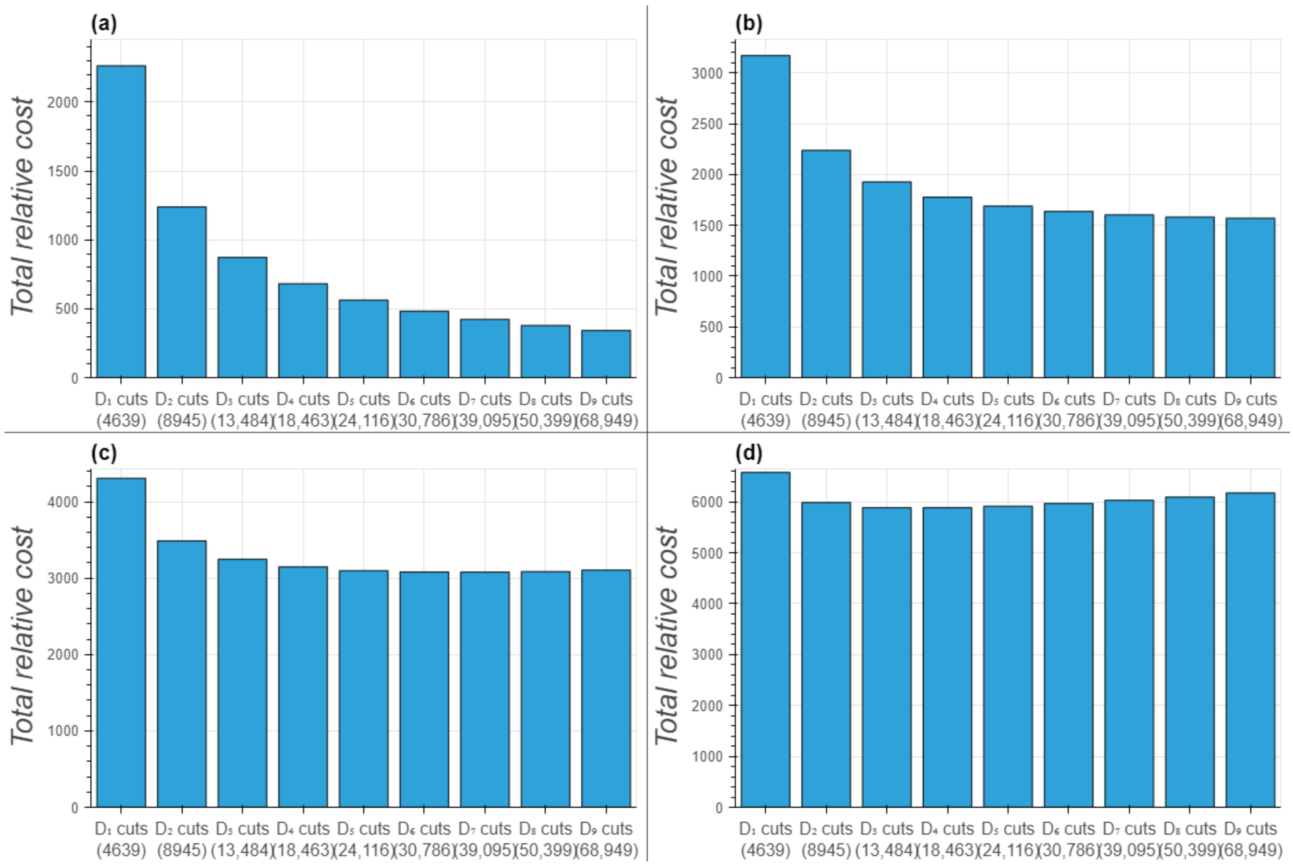

Figure 5. Results of the simulation with 100 different seeds and 10,000,000 cuts. Different simulation scenarios are presented according to the relationship of unprogrammed and programmed costs: (**a**) unprogrammed cost = programmed cost. (**b**) unprogrammed cost = 5 × programmed cost. (**c**) unprogrammed cost = 10 × programmed cost. (**d**) unprogrammed cost = 20 × programmed cost. The different deciles are used in the horizontal axis.

For higher values of the quotient between costs, a threshold for minimising the total cost will become more conservative so as to prevent as many blades as possible from unprogrammed replacements. We can also see that, if there were no difference in the relation between unprogrammed and programmed costs, as seen in the top left graph, the most sensible choice would be just to let each blade's life run its course since this ensures that we get the most cuts out of each one of them. However, this is rarely the case, since most of the time unprogrammed changes incur time delays related to stopping the production chain, replacement times, or other reasons.

From a suitable choice of quantile, the threshold cuts are obtained and a practical Health Index to asses the remaining usable life of the blades is created considering that every blade is "completely healthy" at the beginning with 0 cuts and is due for replacement, i.e., no remaining usable life, when the threshold is met. Thus, we opt for for linear estimation of the blades' status, as indicated in Figure 6, for the same quantiles discussed previously. Note that D_5 is depicted with increased width as it is the one we chose in practice.

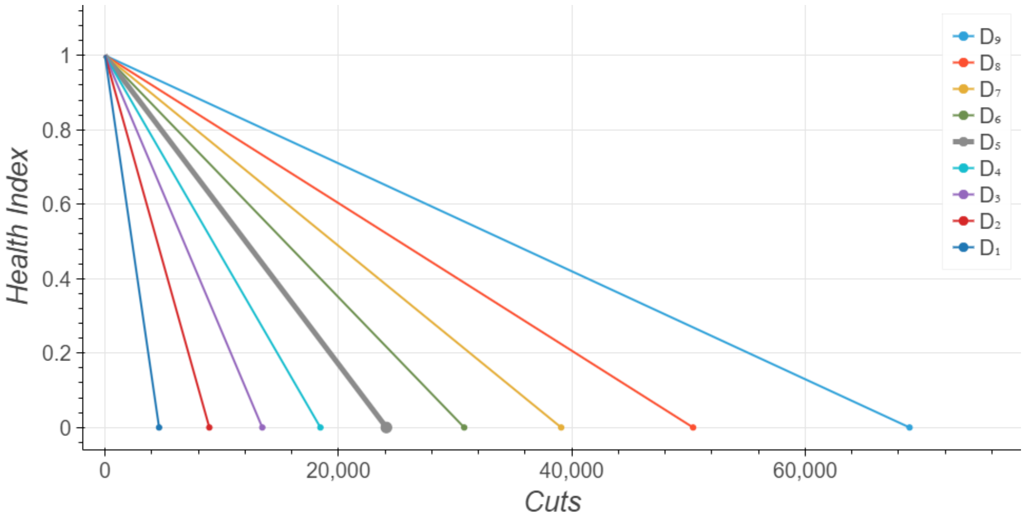

Figure 6. Health index results

3.2. Post Deployment

3.2.1. Data Quality

As the project advanced in development, more information could be extracted from the blades. This is mostly related to the feedback provided by the operator on the shop floor, who indicates the motivation for the replacement of the cutting blades. With this piece of additional information, it is now possible to categorise the reason for replacement, and new cut distributions follow, as seen in Figure 7.

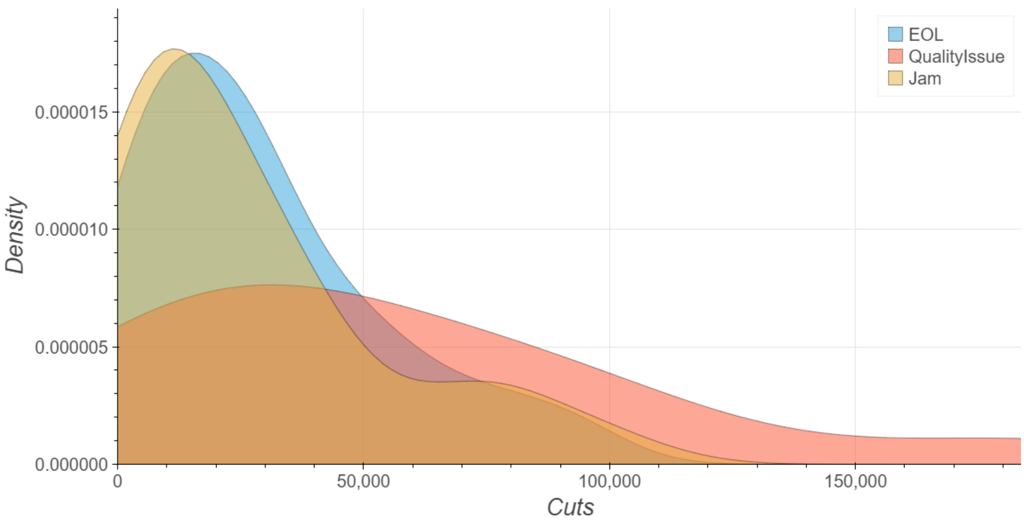

Figure 7. New data densities according to blade–change motivation.

In addition, every time the security lid is opened, operators can also mention the reason for intervening in the cutting system. This information is used to compute the amount of jams or cleaning via water jet that occurred during the blade's lifetime. By engineering these values and recoding them as factors, a new variable called *interventions* is created according to

$$interventions = jams + waterJetCleaning = \begin{cases} 0, \\ 1, \\ \geq 1 \end{cases} \tag{6}$$

and we end up with three different classes to consider. We chose a fixed number of three classes in order to keep a reasonable amount of registers in each one of them for further modelling. The number of blades within each category can be seen in Figure 8.

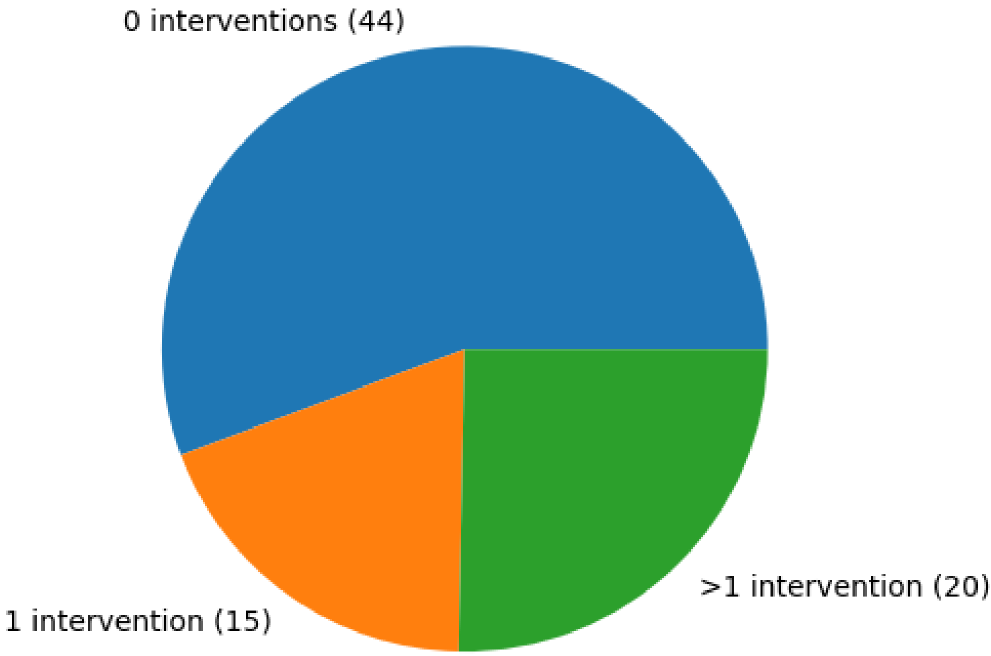

Figure 8. Pie chart regarding the amount of interventions for the blades.

According to Figure 8, from the new set of 79 blades recorded after the deployment of the platform, 55.3% have no interventions during their lifetime, 19% have a single intervention, and the remaining 25.3% have two or more interventions during their lifetime.

3.2.2. Modelling

The creation of different classes of intervention performed on our blades allows us to propose different modelling approaches:

- Simple Weibull fitting, in which the interventions are disregarded and only the total cuts are considered. This is our initial modelling approach and will serve as a baseline.
- Context–Aware Weibull, in which each class is modelled independently, setting three different fits to the total cuts of the blades with the same amount of interventions.
- XGBoost, taking the amount of interventions as a feature used in the predicting task.

For the experimentation, we perform 3–fold cross validation repeated 10 times, and the metrics used to compare the different approaches are the Root Mean Squared Error and Mean Absolute Error, as shown in Figure 9.

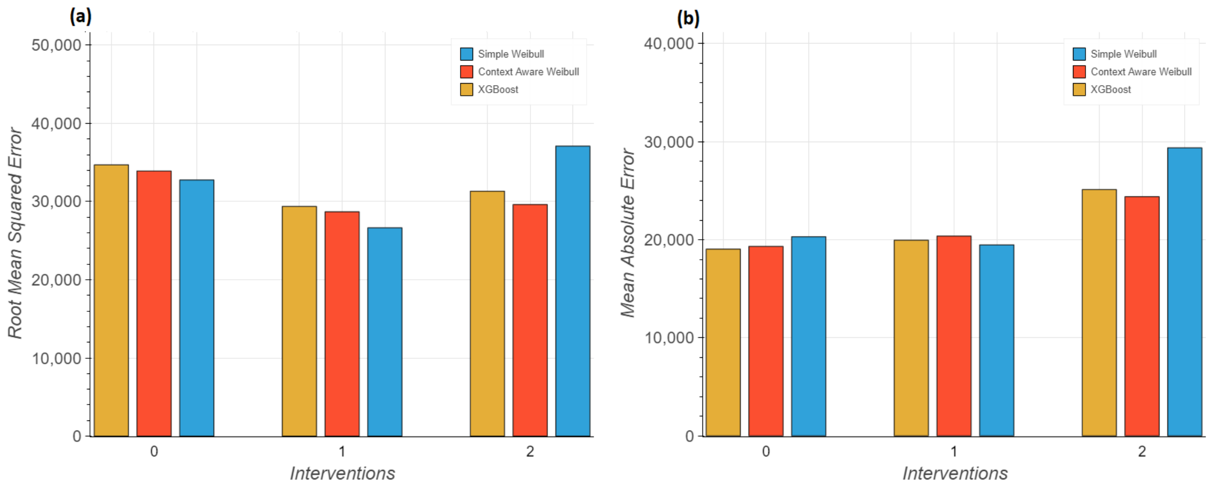

Figure 9. Cross -validation results for Root Mean Squared Error (**a**) and Mean Absolute Error (**b**).

It is noteworthy to mention that the results are not conclusive in the sense that, depending on which metric it is used, it might seem one modelling approach might be more accurate than the other. For instance, considering RMSE, a Simple Weibull model might perform better than the other modelling strategies and only provide a worse estimation for the cases where more than one intervention occurs. If, on the contrary, the MAE is considered, XGBoost or Context–aware Weibull might be more accurate approaches. This might be caused by the existence of outliers, which would be more penalized by RMSE or also by the fact that the variation on the number of cuts at the end of life is too big.

Regarding the computation cost, XGBoost requires significantly longer times for training the models, as presented in Figure 10, while both Weibull models require a similar amount of time. Nevertheless, given that the amount of blades that are changed daily is small, this computational cost is negligible as the dataset is not expected to increase exponentially in the future.

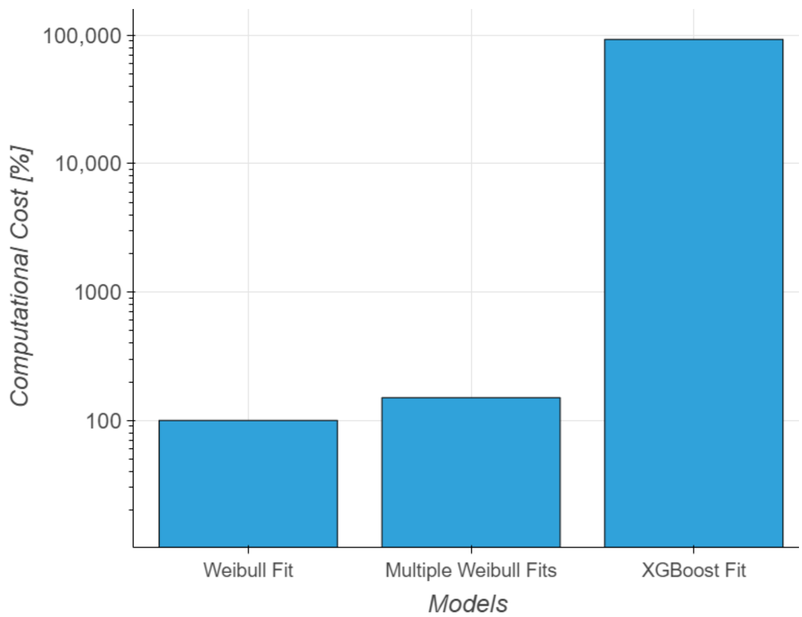

Figure 10. Time needed for computation relative to Simple Weibull fitting.

Although this insight may provide grounds for further model improvements, it must be considered that the number of new registered blades is small in comparison to the previous historic dataset (79 vs. 227). Therefore, this choice of best algorithm should be revised and reconsidered in the future once more data are made available.

Under such uncertain scenarios and for the sake of deployment simplicity, a decision was made to keep a regular Weibull model.

3.2.3. Retraining Criteria

Given that the deployed model is operating in a live environment, the need to improve over time has been considered. In that regard, the availability of new data suggests that, by retraining the model, the estimations should be improved. However, this raises another interesting question: the amount of data to be used for the retraining of the model. It is known that the distributions of the data tend to vary from the original distribution, this is called concept drift. To prevent consequent model deterioration, the choice of retraining window is not trivial. In an attempt to identify and establish the optimal amount of blades (or window size) used for fitting the model, a comparative analysis is carried out. Our main options consist of fixing the number of blades used for the Weibull fitting (i.e., using fixed moving windows over the incoming data) and performing the fit including every blade recorded so far, resulting in an increasing window collecting every blade. For comparison between the different window sizes, we refer to Figure 11, where the evolution of the goodness of fit over time is depicted.

Figure 11. R^2 evolution starting from 50 initial blades.

According to the results, including every blade as part of our retraining strategy yields the best and most stable results as far as Weibull fitting is concerned, as the fitting is kept more stable throughout. Other strategies that use smaller windows tend to fluctuate too much, which might imply there might be some sort of seasonality on the duration of the blades. Nevertheless, the overall fitting reached by just adding blades to the dataset used for training the models is superior, regardless of the drifts.

3.2.4. HI Deployment

Regarding the deployment of the model and the Health Index, the evolution of the Health Index for a specific blade is depicted in the following Figure 12. Even though, finally, regular Weibull is deployed, both Simple Weibull and Context–aware Weibull are compared in this figure for a particular blade for illustrative purposes.

Interestingly, the effect of interventions is the opposite to the one originally expected. Having interventions through a blade's life seems to prolong the life of the blade instead of reducing it. This is visible by the slopes of the the multi–modal HI, which decreases the slope whenever an intervention occurs. This means the blades that contain one or more interventions last longer than the ones that do not have any, which is counter intuitive and should be studied in detail.

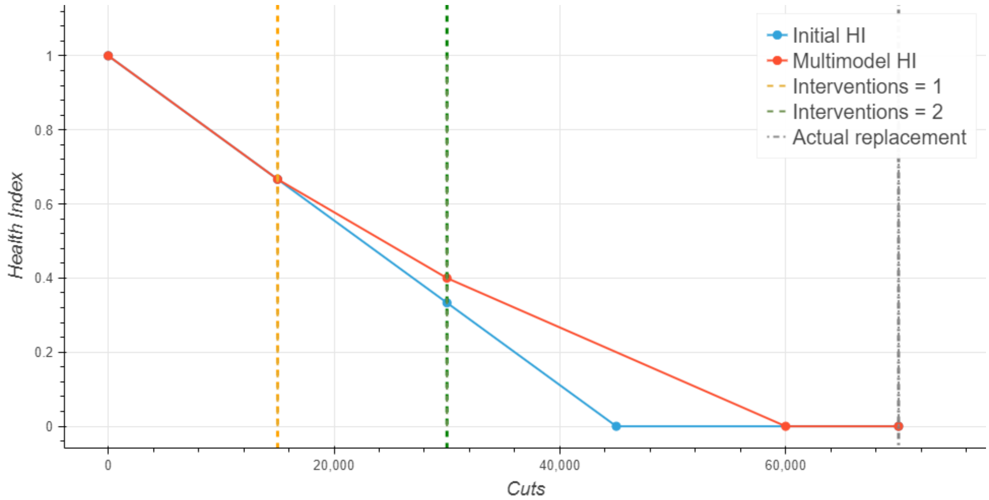

Figure 12. Depiction of the evolution of the Health Index in a blade with two interventions.

4. Conclusions

This work presents a research carried out in two stages in relation to the development of a cutting blade system Health Indicator. As in many of the works in the literature, the first modelling attempts are carried out on a fixed set of data. However, this work goes a step beyond by showing deployment considerations and constraints that are otherwise overlooked. Lying in between Industry 4.0 (with connected machines and data streams) and 5.0 (where the operator is placed in the centre), this work contributes to bridging the gap between static research and the development of models and the dynamic research that occurs post–deployment, a domain seldom explored in academia. Crucial challenges in the ongoing evolution of applied technologies are shown, such as the need for robust data streams and the role of the human within the system. This research work not only extends the theoretical framework but also ensures practical relevance in real–world applications. In addition, the simplicity of the approach, together with the modest data requirements (usage data), makes this approach easily transferable to similar scenarios.

A method to build a Health Index from a Weibull–like survival model is presented. This method considers the relation between planned/unplanned replacements and re–scales the outcome of the survival model to a user–friendly Health Index. Furthermore, the method is generic and could be used in other applications of a similar nature.

During deployment, it has been demonstrated that, beyond following the model software development best practices (typically gathered under the MLOps umbrella), the role of the operator is paramount. In this particular work, the operator is both the consumer and the provider of data. They benefit from the Health Index and, at the same time, provide the system with invaluable information regarding the interventions made to the blade as well as the blade replacements carried out. Lacking this information, the system would not be valid. For that purpose, the role of the human feedback system needs to be highlighted: It simplifies the work operation (by going from free–text to pop–up) and at the same time it improves the quality of the retrieved data, reaching 210 inputs given by the operators through the HMI at the time of the analysis. This is a simple demonstration on how the role of operators will evolve from being a consumer to being a consumer/prosumer under the human–centred AI paradigm.

With the available data at the time of the analysis, survival/reliability algorithms have been shown to be simple yet powerful methods to make the estimations, showing almost equivalent results to more complex models, including operations such as the XGBoost model.

Nevertheless, having to deal with a deployed model opens new research needs. The problem is no longer just fitting the data once, it is considering that this data increases over time and that its nature might change. In such a dynamic system, the analytical results obtained at certain points in time (such as the simple model having an equivalent fit to more complex fittings) might need to be revised in the future, when more data become available.

It is expected that the developed models will continue gathering more and more accurate data in the future. In that sense, revising the results of the best model and retraining strategies will be necessary in the future, as will considering the improvement of usage–based estimation with the inclusion of additional sensors to obtain a more detailed appraisal of the blades.

Author Contributions: Conceptualization, A.M.B., E.G.P. and K.L.d.C.E.; methodology, A.M.B. and K.L.d.C.E.; software, A.M.B.; validation, A.M.B. and K.L.d.C.E.; formal analysis, A.M.B., E.G.P. and K.L.d.C.E.; investigation, E.G.P., A.M.B. and K.L.d.C.E.; resources, A.A. and K.L.d.C.E.; data curation, A.M.B.; writing—–original draft preparation, A.M.B. and K.L.d.C.E.; writing—–review and editing, A.M.B. and K.L.d.C.E.; supervision, A.A.; project administration, A.A. and K.L.d.C.E.; funding acquisition, A.A. All authors have read and agreed to the published version of the manuscript.

Funding: This research has received funding from the European Union's Horizon 2020 research and innovation program AI–PROFICIENT under grant agreement No.957391. Additionally, this work has been supported by the Elkartek Program of the Basque Government (Spain) under the BEREZ–IA project, grant number KK–2023/00012. The content of this article represents the views of the author only and is his/her sole responsibility. Neither the European Commission nor the Basque Government accepts any responsibility for use that may be made of the information it contains.

Institutional Review Board Statement: Not applicable.

Data Availability Statement: The data supporting the findings of this study are derived from proprietary sources and contain sensitive information pertaining to industrial processes that cannot be disclosed. The nature of the data involves confidentiality agreements and intellectual property rights that prevent public sharing. Therefore, the datasets generated and analysed during the current study are not publicly available.

Acknowledgments: We thank Continental for their support and permission to publish the results of this research, which took place in the context of the AI–PROFICIENT project, funded by the European Union's Horizon 2020 research and innovation program. Continental's willingness to allow the dissemination of these findings demonstrates their commitment to advancing knowledge and innovation in the field.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Salonen, A.; Tabikh, M. Downtime costing—Attitudes in Swedish manufacturing industry. In Proceedings of the 10th World Congress on Engineering Asset Management (WCEAM 2015), Tampere, Finland, 28–30 September 2015 ; Springer: Berlin/Heidelberg, Germany, 2016; pp. 539–544.
2. Engineering, P. *Industrial Maintenance Report, March 2021*; Technical Report; Plant Engineering: Chiago, IL, USA, 2021.
3. Nasir, V.; Cool, J. A review on wood machining: Characterization, optimization, and monitoring of the sawing process. *Wood Mater. Sci. Eng.* **2020**, *15*, 1–16. [CrossRef]
4. Efe, T.; Akbay, D.; Demirdag, S.; Sengun, N.; Altindag, R.; Ugur, I.; Mutluturk, M. Investigation of The Cutting Performance of Some Carbonate Rocks with Circular Diamond Sawing Machine. In Proceedings of the 2014 ISRM International Symposium—8th Asian Rock Mechanics Symposium, Sapporo, Japan, 14–16 October 2014.
5. Kang, J.; Zhang, J.; Wang, K.; Zhang, D.; Bai, T.; Zhang, H.; Song, W. Analysis of Wear Mechanism and Sawing Performance of Carbide and PCD Circular Saw Blades in Machining Hard Aluminum Alloy. *Int. J. Refract. Met. Hard Mater.* **2023**, *116*, 106362. [CrossRef]
6. Ekevad, M.; Cristóvão, L.; Marklund, B. Wear of Teeth of Circular Saw Blades. *Wood Mater. Sci. Eng.* **2012**, *7*, 150–153. [CrossRef]
7. Kminiak, R.; Gašparík, M.; Kvietková, M. The Dependence of Surface Quality on Tool Wear of Circular Saw Blades during Transversal Sawing of Beech Wood. *BioResources* **2015**, *10*, 7123–7135. [CrossRef]
8. de Calle Etxabe Kerman, L.; Eider, G.P.; Aitor, A. Towards a Circular Rotating Blade Wear Assessment Digital Twin for Manufacturing Lines. *IFAC-PapersOnLine* **2022**, *55*, 561–566. [CrossRef]

9. Murthy, D.P.; Bulmer, M.; Eccleston, J.A. Weibull model selection for reliability modelling. *Reliab. Eng. Syst. Saf.* **2004**, *86*, 257–267. [CrossRef]
10. Prasad, P.; Rao, K. Reliability models of repairable systems considering the effect of operating conditions. In Proceedings of the Annual Reliability and Maintainability Symposium. 2002 Proceedings (Cat. No. 02CH37318), Seattle, WA, USA, 28–31 January 2002; IEEE: Seattle, WA, USA, 2002; pp. 503–510.
11. Zeinoddini-Meymand, H.; Vahidi, B. Health index calculation for power transformers using technical and economical parameters. *Iet Sci. Meas. Technol.* **2016**, *10*, 823–830. [CrossRef]
12. Azmi, A.; Jasni, J.; Azis, N.; Kadir, M.A. Evolution of transformer health index in the form of mathematical equation. *Renew. Sustain. Energy Rev.* **2017**, *76*, 687–700. [CrossRef]
13. Sculley, D.; Holt, G.; Golovin, D.; Davydov, E.; Phillips, T.; Ebner, D.; Chaudhary, V.; Young, M.; Crespo, J.F.; Dennison, D. Hidden technical debt in Machine learning systems. In Proceedings of the 28th International Conference on Neural Information Processing Systems—Volume 2 (NIPS'15), Montreal, QC, Canada, 7–12 December 2015; MIT Press: Cambridge, MA, USA, 2015; pp. 2503–2511.
14. Faubel, L.; Schmid, K.; Eichelberger, H. MLOps Challenges in Industry 4.0. *Comput. Sci.* **2023**, *4*, 828. [CrossRef]
15. Huang, S.; Wang, B.; Li, X.; Zheng, P.; Mourtzis, D.; Wang, L. Industry 5.0 and Society 5.0—Comparison, Complementation and Co-Evolution. *J. Manuf. Syst.* **2022**, *64*, 424–428. [CrossRef]
16. Raja Santhi, A.; Muthuswamy, P. Industry 5.0 or Industry 4.0S? Introduction to Industry 4.0 and a Peek into the Prospective Industry 5.0 Technologies. *Int. J. Interact. Des. Manuf. (IJIDEM)* **2023**, *17*, 947–979. [CrossRef]
17. Wang, L. A Futuristic Perspective on Human-Centric Assembly. *J. Manuf. Syst.* **2022**, *62*, 199–201. [CrossRef]
18. Kerman López de Calle, E.; Muro, A.; Eider Garate, P.; Arnaiz, A. Process variability aware Health Index for the optimal cutting blade replacement in industrial environments. *Procedia Comput. Sci.* **2024**, *232*, 3003–3012. [CrossRef]
19. Chen, T.; Guestrin, C. XGBoost: A Scalable Tree Boosting System. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, New York, NY, USA, 13–17 August 2016; KDD '16, pp. 785–794. [CrossRef]
20. Fernandez, I.; de Calle, K.L.; Garate, E.; Benzmuller, R.; Kessler, M.; Anderson, M. Human-Feedback for AI in Industry. In Proceedings of the CENTRIC 2023, The Sixteenth International Conference on Advances in Human-oriented and Personalized Mechanisms, Technologies, and Services, Valencia, Spain, 13–17 November 2023.