Joint Prescriptive Maintenance and Production Planning and Control Process Simulation for Extrusion System

Kevin Wesendrup[1] and Bernd Hellingrath[2]

[1,2] *European Research Center for Information Systems, Münster, North-Rhine-Westphalia, 48149, Germany*

kevin.wesendrup@ercis.uni-muenster.de
bernd.hellingrath@ ercis.uni-muenster.de

ABSTRACT

Production planning and control (PPC) is the mainstay of every manufacturer and ensures flawless production processes. However, PPC is jeopardized by breakdowns that can only be tackled with appropriate maintenance. In the past, static strategies, such as reactive and scheduled maintenance, have been used. Yet, with growing system complexity, Industry 4.0, and abundant sensor data, dynamic strategies through PHM have emerged. The most advanced maintenance strategy is prescriptive maintenance (PxM), which allows manufacturers not only to predict failures but also to establish condition-based production plans and controls. To this end, our study explores the integration of PxM with PPC. First, we propose a fault prediction model based on health indicators and future loads of a single-machine system. The proposed fault prediction is integrated into a joint PxM and PPC simulation model that compares the makespan of three joint PxM and PPC strategies inter se and versus reactive and scheduled maintenance. A simulation study using industrial data from an extrusion process evaluates the different strategies across different time horizons (one month to a year). The findings indicate that joint PxM and PPC outperform other strategies, providing significant time savings over traditional methods. Further, a sensitivity analysis is conducted to assess the robustness of the PxM strategies under varying levels of measurement noise, revealing potential challenges under high noise conditions. The study contributes to the field of PHM by providing insights into the effectiveness of joint PxM and PPC strategies and offering a comprehensive analysis of their performance under different conditions.

1. INTRODUCTION

Production planning and control (PPC) is the brain of any manufacturing company (Kiran, 2019). Its main objective is to ensure that products and services are produced efficiently, at the right cost, and in the quantities required to meet customer demand (Cadavid, Lamouri, Grabot, Pellerin, & Fortin, 2020). Production planning is concerned with scheduling or lot sizing, while production control monitors and regulates production capacities (Schmidt & Schäfers, 2017). However, breakdowns are detrimental to these production plans and controls (Zarte, Wunder, & Pechmann, 2017).

Maintenance is concerned with avoiding breakdowns and, in the best case, improving overall business performance (Bousdekis, Magoutas, Apostolou, & Mentzas, 2018), for which reactive and scheduled maintenance strategies were employed traditionally. The former maintains machines after failure, which does not eliminate losses of production capacity; the latter maintains machines regularly but is often overly strict, causing unnecessary maintenance activities (Liu, Chang, & Chen, 2023).

With the abundance of sensor and production machine data, condition-based maintenance and prognostics and health management have been established, which help to maintain machines adequately through health detection, prognostics, and decision-making (Guillén, Crespo, Macchi, & Gómez, 2016). While health detection and prognostics enable a predictive maintenance strategy, value is only generated through decision-making (Jia, Huang, Feng, Cai, & Lee, 2018), i.e., prescriptive maintenance (PxM). PxM is an evolution of predictive maintenance (Meissner, Meyer, & Wicke, 2021) and enables adjusting production plans and controls based on condition information (Elbasheer et al., 2022).

For instance, PxM can help PPC decision-makers by constructing optimal schedules based on order-specific degradation (Zhai, Gehring, & Reinhart, 2021), finding production quantities so that a failure during a lot is prevented (Zheng, Zhou, Gu, & Zhang, 2021), or control machine speeds to postpone breakdowns (Broek, Teunter, Jonge, & Veldman, 2021). However, while PxM gains popularity, the current literature is limited (Pinciroli, Baraldi, & Zio, 2023).

https://doi.org/10.36001/ijphm.2024.v15i2.3839

PxM decisions are generated through either optimization or simulation (Goby, Brandt, & Neumann, 2023). While the former is ideal for finding the best solution to well-defined problems with clear objectives and constraints, simulation is suitable for understanding complex systems and exploring different scenarios, especially when the underlying processes are not well-defined (Laguna & Marklund, 2018).

The available literature on PxM is focused on the optimization of singular PPC steps, such as scheduling, lot sizing, and production control. However, a) it is unclear how the different PxM strategies compare inter se and how they fare against 'traditional' strategies, such as reactive or scheduled maintenance; b) existing studies are based on optimization and assume perfect knowledge of the underlying system, e.g., by assuming the failure behavior of a machine or product-specific degradation values are known (Bousdekis et al., 2018); and c) most PxM studies are not validated with industrial data (Zhai et al., 2021).

Therefore, this study's research goal is as follows:

Develop a simulation that compares joint prescriptive maintenance and production planning and control strategies and showcases whether they outperform other maintenance strategies using industrial data.

The work is divided into multiple sections to address this goal. The following section describes a literature review of related works and outlines the addressed research gap in more detail. Section 3 describes the integrated PxM and PPC problem. Section 4 presents our joint PxM and PPC process simulation model. The penultimate section demonstrates the instantiation of our model in a real industrial use case of an extrusion process. Section 6 concludes the work, highlights limitations, and derives a research agenda.

2. LITERATURE REVIEW

We retrieved related literature that employs PxM for at least one PPC-related decision. We found that most works focus on optimizing decisions related to joint PxM and production scheduling, lot sizing, or control, as shown in Table 1.

Joint PxM and production scheduling is concerned with sequencing production orders, jobs, or operations, and maintenance interventions on machines. Fitouri, Fnaiech, Varnier, Fnaiech, and Zerhouni (2016) propose a heuristic, and Ladj, Tayeb, Varnier, Dridi, and Selmane (2017; 2017), and Zhai, Riess, and Reinhart (2019) propose a mixed integer linear program and genetic algorithms that optimize the scheduling of orders on a single machine. As they use optimization, all works assume that degradation values for each order are fixed and known. In contrast, Zarte et al. (2017) describe a laboratory simulation test bed and do not make assumptions about order-specific degradation values. However, breakdowns are simulated based on the mean time between failure, and other PPC steps are not addressed.

Publication	**PxM**	**Scheduling**	**Lot sizing**	**Control**	**Traditional maintenance**	**Simulation**	**Real industrial data**
(Ladj et al. 2017a; Ladj et al. 2017b; Zhai, Riess, and Reinhart 2019; Fitouri et al. 2016)	✓	✓					
(Zarte, Wunder, and Pechmann 2017)	✓	✓				✓	
(Dehghan, Nourelfath, and Hajji 2023; Zheng et al. 2021)	✓		✓				
(Yang, Zhao, and Han 2022)	✓		✓				✓
(Broek et al. 2021; Broek et al. 2020; Wesendrup and Hellingrath 2023)	✓			✓	✓	✓	
Our study	✓	✓	✓	✓	✓	✓	✓

Table 1. Comparison of related works

Joint PxM and lot sizing determines economic production quantities that facilitate opportunistic maintenance windows between lots. Zheng et al. (2021) use a semi-Markov decision process to calculate optimal lot sizes and maintenance windows but do not make comparisons to other PPC or traditional maintenance strategies. Dehghan, Nourelfath, and Hajji (2023) optimize the lot sizes using mathematical programming applied to an artificial turbofan dataset. In contrast, Yang, Zhao, and Han (2022) use industrial data from a steel factory to find an economic production quantity.

Joint PxM and production control adjusts the production levels or rates based on the degradation of the machines to postpone breakdowns. Broek, Teunter, Jonge, Veldman, and van Foreest (2021; 2020) control the machine degradation by dynamically adjusting the production rate using a simulation with different parameter sets. Wesendrup and Hellingrath (2023) use simulation-based reinforcement learning to set production levels based on machine degradation. All three works compare their models to scheduled maintenance strategies, but do not make use of industrial data.

While the identified works highlight the benefits of joint PxM and PPC, they also reveal some research gaps. First, no work compares multiple strategies, and only a few are benchmarked with 'traditional' maintenance (e.g., reactive, scheduled), leaving researchers and managers clueless about the relative strength of each strategy. Secondly, most works use optimization and, e.g., assume perfect knowledge of future degradation for each production order. This is unrealistic, as degradation is time-varying and “different operational conditions result in different system loads and thus different degradation effects” (Zhai et al., 2021, p. 2). Lastly, only a few studies validate their approaches with actual industrial data.

In contrast, our work proposes a simulation that compares joint prescriptive maintenance and production scheduling, lot sizing, and control and showcases whether they outperform other maintenance strategies using industrial data. Our research involves a comparative analysis of various maintenance strategies. Comparative studies are crucial in understanding the relative effectiveness of different approaches. By evaluating joint PxM alongside other strategies, our work provides valuable insights into which method offers the most efficient solution, a perspective the existing literature lacks. Furthermore, our work uses discrete-event simulation, which is particularly useful when dealing with complex and dynamic systems and allows us to examine different maintenance and production strategies under time-varying degradation effects and shocks, helping to identify resilient strategies. Lastly, many simulations rely on hypothetical or simulated data, which might not accurately represent the complexities of real-world scenarios. Our study uses actual industrial data to provide practical insights and solutions directly applicable to industrial settings.

3. SIMULATION PROBLEM

The scope of our simulation problem comprises a set O of n non-preemptive orders $o \in O = \{o_1, o_2, \dots, o_n\}$ that each produce several products. Each order is available at the beginning, takes a fixed time t_S to set up and a variable, order-dependent time t_i to process. Orders are processed on a single-machine system with health h that is measurable between production runs using condition monitoring. Each order exerts a product-specific, known load l_i over its processing time t_i that reduces h by an unknown degradation factor. The load l and processing time t, and consequently the degradation, are similar for similar products (i.e., products that consist of similar raw materials). If h drops below 0, the machine breaks down, and a reactive repair with a fixed repair time t_R must be carried out.

Machines can be preventively maintained in between orders to prevent lengthy breakdowns using condition-based maintenance, which requires a fixed maintenance time $t_P \gg t_R$. Additionally, by analyzing the load l and processing time t of upcoming orders and whether they would lead to breakdowns, production plans and controls can be adjusted using one of three strategies:

1) *Joint PxM and production scheduling* allows scheduling orders based on their load l_n and processing time t_n. This allows heavier orders to be processed at the beginning and lighter orders at the end of the machine's life. Additionally, it allows to check whether lighter orders can still be produced when heavier orders would lead to failure.

2) *Joint PxM and lot sizing* allows splitting an order into two lots. Here, l stays the same, but the processing time t is split, effectively “dividing” the caused deterioration. In between both lots, a condition-based maintenance intervention can be performed.

3) *Joint PxM and production control* allows for a reduction in the production rate, which lowers l and decreases the caused degradation. However, this causes delays in the order’s processing time t.

The main objective of our simulation problem is to minimize the total makespan of all orders O.

4. SIMULATION MODEL

The final concept of the discrete-event simulation model is summarized visually in Figure 1 and explained in the upcoming sections. It comprises a health indicator construction and failure prediction that provides inputs for the three joint PxM and PPC strategies.

4.1. Health and Load-based Failure Prediction

To solve the described decision problem, first, we need to approximate the machine's health and find a way to predict failures. We assume that there exists historical time series data about the order o (i.e., id), and sensor values of the machine's condition (e.g., vibration, temperature) and load parameters l (e.g., conveyor speed, heater temperature) and its breakdowns and maintenance repairs. From there, one can construct a health indicator proposed by Medjaher, Zerhouni, and Baklouti (2013). Hereby, each run-to-failure sensor time series is correlated to a healthy machine reference time series. Then, correlations are sampled, standardized to a failure threshold, smoothed, and a non-linear regression is fitted. The resulting standardized correlation coefficient h (from 0 to 1) of the two time series equals the machine's health. The further the sensor values deviate from the reference time series, the lower the correlation (i.e., $h \rightarrow 0$) and hence, the machine's health.

This health indicator h can be calculated for each historical time point and, consequently, the health at the beginning and end of an order o, depicted in Figure 2.

This analysis allows us to get historical data about the health h_i before producing an order o_i, the order’s load parameters l_i and processing time t_i, and whether the production of specific orders led to failure. We assume that the planned load l_i and processing time t_i of a new order o_i is available. With this information, we trained a machine learning classifier that is used as input in our simulation study. The fault predictor $f(h_i, l_i, t_i)$ is able to predict whether a new order is causing a failure based on the machine’s current health indicator h_i, the planned load l_i and the time t_i the load is

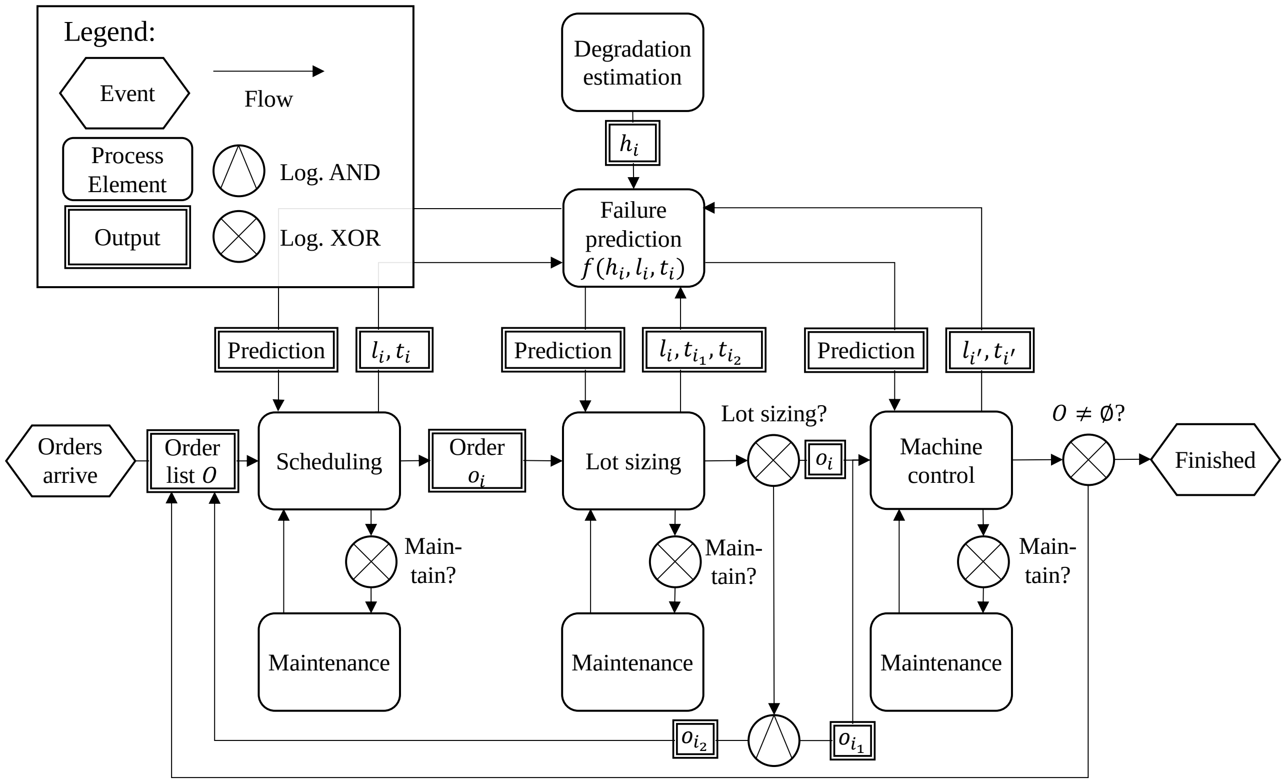

Figure 1. Conceptual model of the simulation.

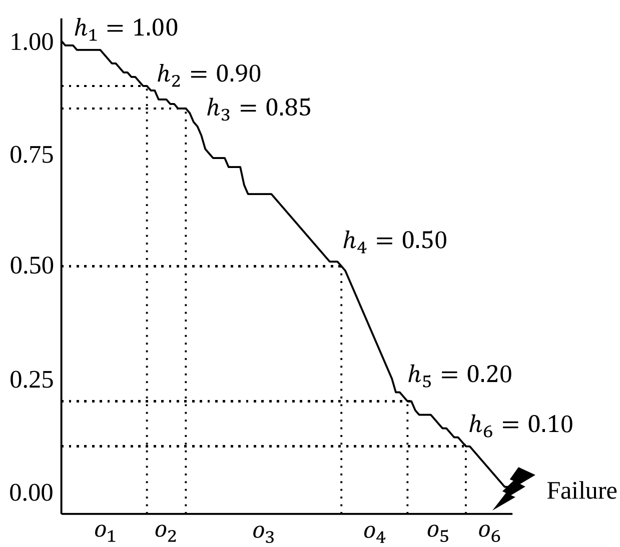

Figure 2. Exemplary run to failure

exerted on the machine. For our fault prediction, most classifiers are usable, but it is recommend that models are chosen that can handle many input variables and are supervised. For an exhaustive list of fault predictors, the reader is referred to Divya, Marath, and Santosh Kumar (2023).

In the following, we show how the attained information and failure predictor is used for the three joint PxM and PPC strategies.

4.2. Joint Prescriptive Maintenance and Scheduling

Scheduling is the first task of our PxM and PPC simulation model. It receives n orders and must decide the sequence in which the orders should be produced on the machine, as well as schedule maintenance interventions. Here, condition-based scheduling allows the assessment of each order's 'intensity' by looking at load and processing time used to sort the orders depending on their caused degradation.

In our model, we use the common best-fit decreasing heuristic suggested for condition-based scheduling by Ladj, Varnier, et al. (2017). We chose the best-fit decreasing heuristic due to its simplicity and ability to outperform other algorithms "in important ways" (Coffman, Garey, & Johnson, 1984, p. 38). With this algorithm, orders are sorted based on the product of l and t, and the order with the highest value is chosen. Then, the current health and l and t are input into the failure predictor and the next orders are processed as long as they do not lead to failure. Should the next order risk a breakdown, it is skipped and the next, 'less intense' order is chosen until a fitting candidate is found. If no fitting order is found, the machine's health is too low, and a preventive condition-based maintenance intervention is performed with time t_P.

The advantage of condition-based scheduling is that orders are kept in their original state and are not split (which incurs additional setup time t_S) or processed with a slower speed (which extends t_i). The disadvantage is that it is less flexible and depends on the availability of adequate orders. If only orders with high processing time t and load l are left, more of the remaining useful life of the machine could be wasted.

If PxM is not applied for scheduling, i.e., if we use other strategies, orders are processed first-in-first-out.

4.3. Joint Prescriptive Maintenance and Lot Sizing

After the orders are scheduled, they are forwarded to lot sizing. In case an order o_n will not cause a breakdown, it is fully processed. In the other case, the order quantity is split into two (o_{i1} and o_{i2}) by iteratively testing smaller split coefficients. Based on a decreasing lot split coefficient $c \in\ ]0,1[$, a new lot size is calculated. We assume that the new lot size divides the processing time proportionally (i.e., $t_{i1} = t_i \cdot c$ and $t_{i2} = t_i \cdot (1 - c)$), while the split does not have a significant effect on the load. Hereby, bigger lot splits are created first, and the new processing time of the first lot t_{i1} is put into the fault predictor.

As soon as the predictor is confident that the smaller lot o_{n1} is not leading to failure, it is produced, and the second lot o_{n2} is returned to the front of the order queue. Then, as the end of life of the machine has been reached, a condition-based maintenance intervention is performed.

The advantage of condition-based lot sizing is its flexibility, as orders can be split to exploit the remaining useful life of orders optimally. The disadvantage is that additional setup time t_S is incurred for each new lot, and exploiting more remaining useful life of the machine increases the risk of an unexpected failure.

If PxM is not applied for lot sizing, orders are not split.

4.4. Joint Prescriptive Maintenance and Production Control

After the lot sizes to be produced are defined, short-term production control ensures that the operating regimes are adequately adjusted. This might include setting machine control parameters such as cutting, conveyor speeds, or working temperatures. In PHM theory, it is assumed that different operating regimes influence the degradation of a machine (Wang, Yu, Siegel, & Lee, 2008).

In our work, we assume that the load can be reduced at the beginning of processing an order o_i, leading to a new load $l_{i'}$ which in turn extends the processing time $= t_{i'}$ but also decreases degradation overall. By how much t_i is extended is known if the desired $l_{i'}$ is determined. For instance, halving the trajectory speed of a single-conveyor machine doubles the processing time t. Both $t_{i'}$ and $l_{i'}$ can then be input into the fault predictor again to approximate whether they would lead to failure.

In our model, we process orders with their standard control parameters if they do not cause a failure. When this is not the case, we reduce l_i by small increments, compute $t_{i'}$ and produce the order with the maximum $l_{i'}$ that does not lead to failure. After processing the order with a reduced load, a preventive condition-based maintenance intervention is performed with processing time t_P.

The advantage of joint PxM and production control is its flexibility to adjust l_i. In turn, this might lead to increases in t_i that might exceed t_P, which is a big disadvantage. In this case, performing a preventive maintenance intervention might make more sense without adjusting the production control. Additionally, maxing out the use of machine life also increases the risk of an unexpected failure.

If PxM is not applied for production control, orders are always produced with their standard operating regime.

4.5. Traditional Maintenance Strategies

To not only compare the joint PxM and PPC strategies inter se, we will also compare them to traditional maintenance strategies (i.e., reactive, scheduled), which are explained next.

Reactive maintenance. In the case of reactive maintenance, no condition data is available. The best-fit-decreasing heuristic is not applied during scheduling, and orders are processed first-in-first-out. Further, original lot sizes are kept, and the machine control is not adjusted. As neither h nor failures can be predicted, the machine is run to failure and then subsequently repaired with a repair time of t_R.

Scheduled maintenance. The reactive maintenance strategy allows us to calculate the mean time between failures that can be a reference for setting an appropriate scheduled maintenance interval. Our simulation study uses an optimal interval that minimizes the makespan objective. Like reactive maintenance, PPC decisions are not changed within this strategy. In contrast, however, preventive maintenance interventions with a processing time of t_P are performed regularly. Hence, scheduled maintenance prevents some of the time-consuming reactive repairs but might also lead to over-maintenance as the actual condition of the machine can not be regarded.

5. SIMULATION STUDY

Our simulation study was carried out at a plastic packaging manufacturer that currently uses a balanced mixture of scheduled and reactive maintenance strategies. The investigated production process is an extrusion process transforming different raw materials (e.g., plastic granules, dye) into different types of plastic films. The focal machine is an extrusion machine comprising extruders, a blowhead, an air ring, haul-off, edge trim, and a winder, as shown in Figure 3. Here, the latter is the critical component comprising multiple mechanical failure modes such as deformations or bearing failures.

To perform our simulation study, we have used the VDI 3633 standard that distinguishes between the phases of simulation preparation, execution, and evaluation.

5.1. Simulation Preparation

We acquired roughly one year of failure, maintenance, and sensor data (e.g., condition monitoring and load) to build the simulation model. The latter was available with a frequency of 0.1 Hertz, leading to almost 2 million observations, and comprised 329 features for each observation that could be attributed to the extruder, blowhead, winder, or that were metadata. The features comprised information about the load, temperature, force, energy consumption, raw material, and finished product. Each observation can be attributed to one of 256 unique orders with a processing time between 1 and 17 hours. The setup time between two production orders is fixed at one hour ($t_S = 1$ hour). Over 35 failure and maintenance events were available, each attributed to each observation of the preceding order (e.g., we assumed that the order was 'responsible' for failure).

Using t-SNE (Maaten & Hinton, 2008) for dimensionality reduction, we could condense the load of each order into one dimension. In contrast to other techniques (e.g., principal component or linear discriminant analysis), t-SNE can capture non-linear relationships and preserve local structures, and showed good performances for fault prediction (Liang, Zhang, & Wang, 2023). Thus, it allows to reveal patterns that are not apparent in high dimensional data and, together with the constructed health indicators for each observation, shows the influence of processing time and load on degradation. Figure 4 depicts the relationship between processing time (x-axis) and load (y-axis) on degradation (bubble size and color). The marginal histograms show the statistical distribution of different loads and processing times. The histograms show that lower processing times are more common, while the loads are uniformly distributed with a gap between low to moderate and high load orders. The bubble sizes and colors show that higher processing times increase degradation even for low-load orders, while high-load orders can be processed for smaller durations without causing much degradation. Logically, high-load orders with long processing times cause the most degradation.

Further, we were able to generate a realistic machine and degradation model of the winder using the failure, maintenance, and sensor data and the Python package *progPy* (Teubert, Jarvis, Corbetta, Kulkarni, & Daigle, 2023). Therefore, we tuned the 'loading' by inputting the processing times of our dataset and adjusting the 'states' parameters so that the variances and mean-time-between-failures match our real-world data using sensor values. We used the observed sensor values and expert knowledge from the manufacturer to inform our parameter estimation. The resulting model exponentially degrades with random shocks (e.g., through bearing spalling). The maintenance data and experts also indicated that reactive repairs of the winder take

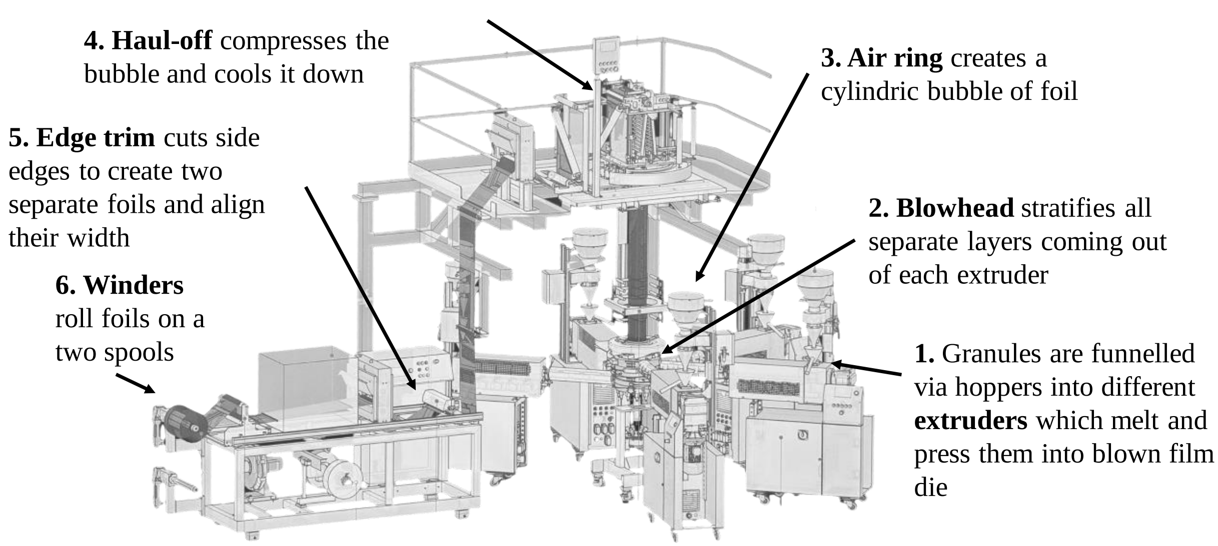

Figure 3. Overview of extrusion machine.

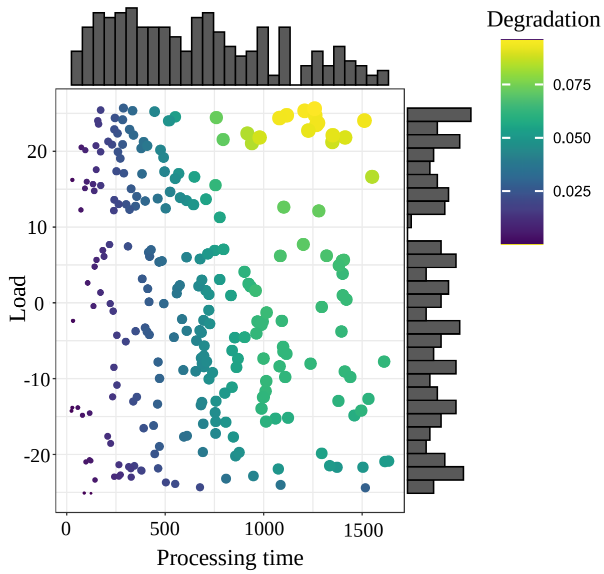

Figure 4. Influence of load and processing time on degradation.

roughly two days to fix ($t_R = 48$ hours), while a preventive maintenance intervention takes roughly half a working day ($t_P = 4$ hours).

5.2. Simulation Execution

For the PxM and PPC simulation experiments, the Java-based simulation software AnyLogic® was used with an interface connected to the failure predictor programmed in Python. All training, testing, and simulation runs were carried out on a Windows computer with 8 GB RAM and a 3.00 GHz Intel® g Core™ i5-9500 CPU with six cores.

Failure prediction. We have used 160,000 random observations to train the fault predictor presented in section 4.1. The predictor variables comprise an order's dimensionality-reduced sensor values and load information. The outcome variable contains information about whether a breakdown occurred during or after processing the current order or not. As with many other condition monitoring datasets, this implies that the dataset is strongly imbalanced with many 'healthy' data points but only a few failures. Hence, we have used the Synthetic Minority Over-sampling Technique (SMOTE), presented by Chawla, Bowyer, Hall, and Kegelmeyer (2011), to balance the training set and generate new failure instances. SMOTE preserves the failure class characteristics and non-linear relationships, does not limit the choice of other classification models, and showed excellent performance over other techniques for condition-based maintenance (Sridhar & Sanagavarapu, 2021).

We used the Python packages *scikit-learn* and *xgboost* to implement our failure prediction, providing state-of-the-art machine learning algorithms. Using these packages, we have analyzed nine different combinations of hyperparameters for three machine learning algorithms (XGBoost, random forest, and Gradient Boosting) using a random search and five-fold cross-validation ($9 * 3 * 5 = 135$ fits). We found that a random forest with a maximum depth of 20 and 150 trees performed best.

Joint PxM & PPC strategies. The final computational implementation of our conceptual model is shown in Figure 5. We generate O, following the statistical distribution of orders shown in the histograms of Figure 4, during 'orderSource' which is forwarded to a queue. From then, orders are processed through the different stages of PPC, as presented in section 4. The different PxM or 'traditional' maintenance strategies can be activated using the parameter 'maintenanceStrategy'. Different process elements and functions realize the described behavior. When plans and controls have been finalized, an order is produced in the service block 'production', which is connected to the extrusion machine.

To model the machine, we used a resource pool entity in AnyLogic® ('extruder'), which is connected to a Python implementation of the machine model based on *progPy* (Teubert et al., 2023).

To simulate the planning of our model, namely the three joint PxM and PPC strategies (scheduling, lot sizing, control), and to compare them with 'traditional' (reactive, scheduled) maintenance strategies, we ran five experiments (by setting 'maintenanceStrategy') for four different time horizons for 100 replications each (= 2,000 runs) to ensure the robustness of the evaluation. Different time horizons from one month to a whole year were simulated by setting the size of O ('numOrders') to 75 (~1 month), 200 (~3 months), 400 (~6 months), and 800 (~1 year).

5.3. Simulation Evaluation

The fault predictor's total training time, including hyperparameter tuning, was 6.5 hours. The total computation time for all 2,000 simulation runs was 14 hours.

Failure prediction. We tested the forest on 40,000 random observations not included in the training dataset, which led to an excellent accuracy of 0.94. As an overall objective, we wanted to predict every costly breakdown, even if that implies that some healthy machine states are misclassified as failures. Therefore, we have used recall as the loss criterion to make the model sensitive to finding failures. This led to a good overall recall of 0.91 but a poor precision of 0.42. This means that more than every second failure was classified wrongly. Fortunately, these misclassifications often happen at the end of the life of the machine, so in the worst case, it means that the machine is maintained even though it can maintain a few further orders without failure. Consequently, the F1 score, which is the harmonic mean of precision and recall, is only mediocre, with a score of 0.57.

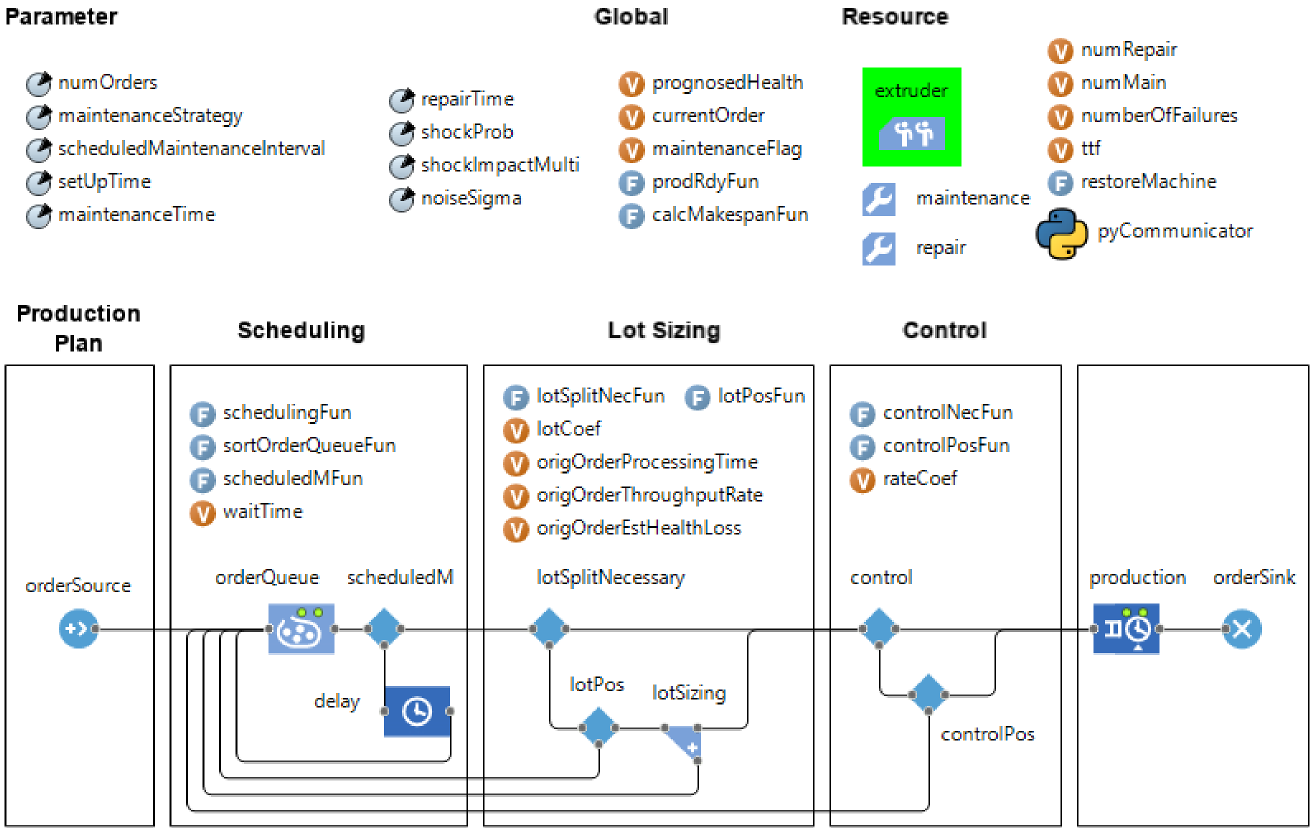

Figure 5. Computer simulation model.

The predictions can be visualized in a confusion matrix (Figure 6) to understand our model's behavior better. The matrix shows that most failure-free observations are classified correctly (36274 true negatives). Additionally, 1462 observations have been correctly classified as failures (true positives). However, many failure-free observations still have been classified as failures (2058 false positives), while luckily, only a few failures have not been identified (148 false negatives). Saxena et al. (2008) state that false positives are suboptimal as they may lead to unnecessary maintenance interventions, while false negatives are much more critical, as they would lead to unexpected failures.

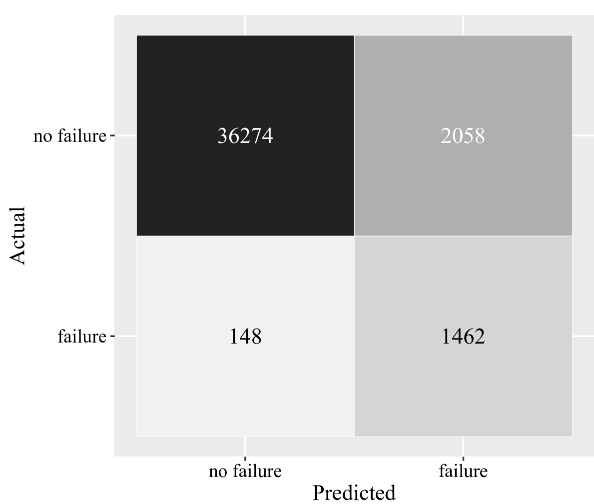

Figure 6. Confusion matrix of failure prediction

While in the context, the results look promising at first glance, Saxena et al. (2008) also suggest analyzing the receiver operator characteristic curve to visualize the trade-off between true and false positives for different prediction thresholds, depicted in Figure 7.

It can be seen that the curve looks very good, which is confirmed by an excellent area under the curve score of 0.93, which signifies a good trade-off has been achieved with our model.

Still, there are some caveats, as few false negatives and many false positives are predicted. To understand the effects of these false predictions, we can look at a run-to-failure to see how the algorithm would have influenced decision-making. Figure 8 shows an example of particularly bad model behavior during one run-to-failure. The x-axis depicts the elapsed time; the y-axis and black solid line the true health of the machine. Each vertical line denotes the time point between two orders, where a fault prediction is made. Black dotted lines denote that the fault prediction correctly estimated that the next order will not lead to a failure (true

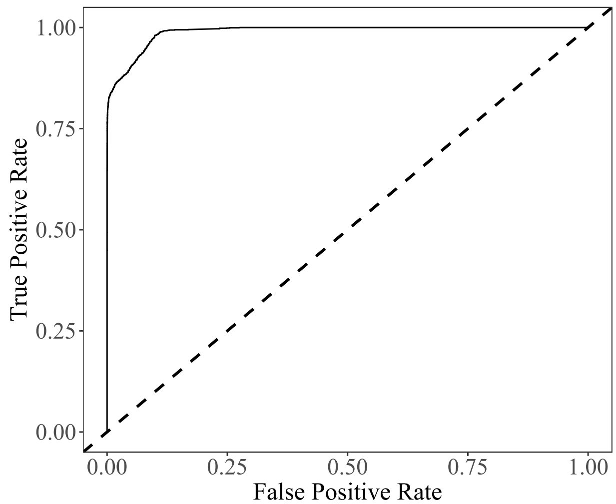

Figure 7. Receiver operating characteristic curve.

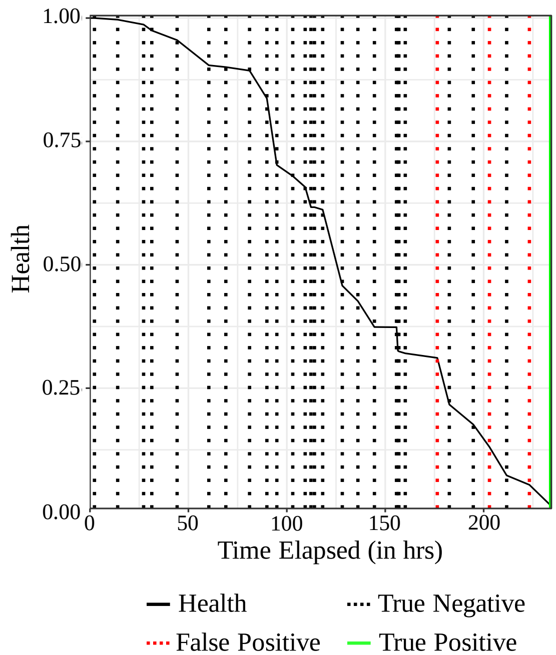

Figure 8. Misclassification of failures

negatives). Red dotted lines denote failure predictions, even though the machine could produce the order without failure (false positives). The green solid line depicts correct failure predictions (true positives). Here, the predictor suggests maintaining early (the first red vertical line) even though the machine can still run six more orders. Also, an alert would have been triggered twice before an actual failure was correctly identified.

We could derive reasons for false positives and negatives from analyzing multiple predictions. The former happens when the health is low (< 0.1) while either the load (<20) or processing time (> 800 minutes) is high. It also happens when the health is only moderately low ($0.3 > h > 0.1$) while load and processing time are high. False negatives can be attributed to observations when the machine health indicator is > 0.1, while winder clamping forces or rotational speeds are very high, and processing times are moderate but not extremely high ($800 < t < 1100$ minutes).

Condition-based Production Planning and Control. The fault predictor has been used as an input for decision-making, leading to joint PxM and PPC. Table 2 shows the average makespan and number of breakdowns and maintenance interventions of each strategy for the four different sizes of O over 100 replications each.

The results show that for $\#O = 75$ (~1 month), PxM-enabled PPC strategies already lead to one to multiple days of savings in the makespan versus scheduled or reactive maintenance strategies. Looking at only one month, *joint PxM and lot sizing* performs best, but the other strategies are very close. Why *reactive maintenance* fails to perform well becomes clear when looking at the number of breakdowns and repairs per run. The winder must be repaired 2.81 times on average in just a month. In contrast, *scheduled maintenance* reduces the average number of breakdowns per month to 0.52, but at the cost of 3.28 maintenance interventions. In contrast to PxM, *scheduled maintenance* has higher average numbers of breakdowns and maintenance interventions, highlighting the inflexibility of the predetermined maintenance interval.

In contrast, the PxM strategies allow prescribing optimal actions based on the real condition of the machine and, thus, reduce the number of preventive maintenance interventions while maintaining a low level of failures. Surprisingly, *joint PxM and production control* is the ‘safest’ strategy, as it reduces failures to a minimum of 0.25 per 75 orders, while it matches the number of maintenance interventions of the best PxM strategy (i.e., *joint PxM and lot sizing*). However, reducing the load of the winder also decreases production speed so that the ‘safest’ strategy does not outperform the other PxM strategies.

Increasing the planning horizon to three months ($\#O = 200$), half a year ($\#O = 400$), and one year ($\#O = 800$) leads to similar results, although there are some performance fluctuations. For instance, when looking at a quarter of a year, *joint PxM and production scheduling* performs best, while *PxM and lot sizing* is the worst PxM strategy. Also, for half a year, joint *PxM and lot sizing* is the safest strategy, with an average of only 1.85 breakdowns.

Over most periods, *joint PxM and lot sizing* consistently performs best, closely followed by the other two PxM strategies. Moreover, *joint PxM and control* is also consistently the safest strategy (i.e., leads to the fewest breakdowns), while *PxM and scheduling* is the most efficient (i.e. minimizes the number of maintenance interventions). Hence, if the objective was concerned with safety or maintenance

No. of orders / strategy	**75**	**200**	**400**	**800**
Makespan in days				
PxM scheduling	32.45 (±2.17)	**86.85 (±3.53)**	174.38 (±4.55)	348.02 (±6.56)
PxM lot sizing	**32.43 (±2.14)**	87.45 (±3.70)	**173.22 (±4.35)**	**347.14 (±6.77)**
PxM control	32.58 (±1.93)	87.12 (±3.49)	174.78 (±4.51)	347.92 (±6.92)
Reactive	36.84 (±2.35)	100.09 (±3.97)	199.44 (±5.18)	400.93 (±7.37)
Scheduled	33.62 (±2.22)	90.42 (±3.34)	181.55 (±6.19)	363.87 (±6.84)
Average number of breakdowns and repairs				
PxM scheduling	0.42 (±0.67)	1.07 (±1.03)	2.54 (±1.41)	4.87 (±2.22)
PxM lot sizing	0.44 (±0.62)	1.13 (±1.16)	**1.85 (±1.23)**	4.1 (±2.01)
PxM control	**0.25 (±0.48)**	**1.02 (±0.83)**	2.07 (±1.44)	**3.88 (±1.87)**
Reactive	2.81 (±0.53)	8.19 (±0.84)	16.49 (±0.94)	33.86 (±1.37)
Scheduled	0.52 (±0.70)	1.72 (±1.31)	3.72 (±1.72)	7.32 (±2.50)
Average number of maintenance interventions				
PxM scheduling	**2.59 (±0.77)**	**7.84 (±1.27)**	**15.97 (±1.67)**	**32.45 (±2.54)**
PxM lot sizing	2.75 (±0.78)	8.09 (±1.36)	16.91 (±1.67)	33.84 (±2.11)
PxM control	2.74 (±0.68)	7.96 (±0.95)	16.36 (±1.80)	33.43 (±2.27)
Reactive	-	-	-	-
Scheduled	3.28 (±0.68)	9.08 (±1.15)	18.43 (±1.49)	37.57 (±2.38)

Table 2. Simulation results (± standard deviation)

cost, the other PxM strategies might perform better and should not be disregarded. However, this does not apply to *reactive* and *scheduled maintenance* that consistently performed worse than PxM strategies. Compared to the ‘traditional’ strategies, time savings of over 15% or 50 days can be achieved over one year ($\#O = 800$) with joint PxM and PPC strategies.

While these figures draw a conclusive picture of the advantages of PxM for PPC, it is interesting to see how this is achieved precisely. Therefore, we visualized an excerpt of one week of an exemplary simulation run, close before and after a first failure or maintenance, for each of the five strategies in Figure 9. The excerpt starts with a relatively healthy machine with an indicator of 0.6. Using a *reactive maintenance* strategy, the machine fails after producing the twelfth order and is subsequently repaired during two working days. After repair, the machine's state is as good as new and further orders can be produced. *Scheduled maintenance* performs interventions at regular intervals. In the shown example, instead of producing twelve orders, only seven orders are produced before the scheduled maintenance interval is reached.

Using *joint PxM and production scheduling*, the machine is maintained even earlier, but this is because orders with higher loads are produced first (best-fit first scheduling heuristic). Thus, the initial health indicator is also much lower at the beginning of the excerpt. Nevertheless, due to PxM, the machine's useful life is exploited much better than the previous ‘traditional’ strategies. *Joint PxM and lot sizing* showcases the advantage of dynamic maintenance times in contrast to fixed *scheduled maintenance* intervals. Instead of maintaining after seven orders, our failure prediction provides confidence for production planning to produce three further orders. Then, instead of maintaining directly, *joint PxM and lot sizing* splits o_{56} into two lots and maintains in-between. *Joint PxM and production control* maintains at a similar point in time, but the strategy works slightly differently. Instead of splitting up the order into lots, the machine produces order o_{56} with a lower load and reduces degradation while extending its processing time. This leaves a slightly higher health safety margin but extends processing time slightly compared to its *joint PxM and lot sizing* pendant.

5.4. Sensitivity Analysis

The considerable makespan reductions through our joint PxM and PPC strategies can be attributed to the excellent performance of our underlying failure prediction model. However, such performances are not always the case, and many uncertainties, particularly occurring in sensor measurements, can impede the usefulness of condition-based maintenance (Sankararaman, 2015). Hence, we performed a sensitivity analysis to showcase how robust our model is and whether it is also useful for more uncertain production contexts.

To achieve that, we induce measurement noise that influences the health indicator, the central component in our failure classification model. In the analysis, Gaussian noise is added to the health h with a standard deviation of 0.01 to 0.1 (in 0.01 steps). In the worst case, this means that the health indicator h is, on average, 10% higher than the true health of the machine. We added noise to the same 40,000 observations we used for evaluation earlier and tested the

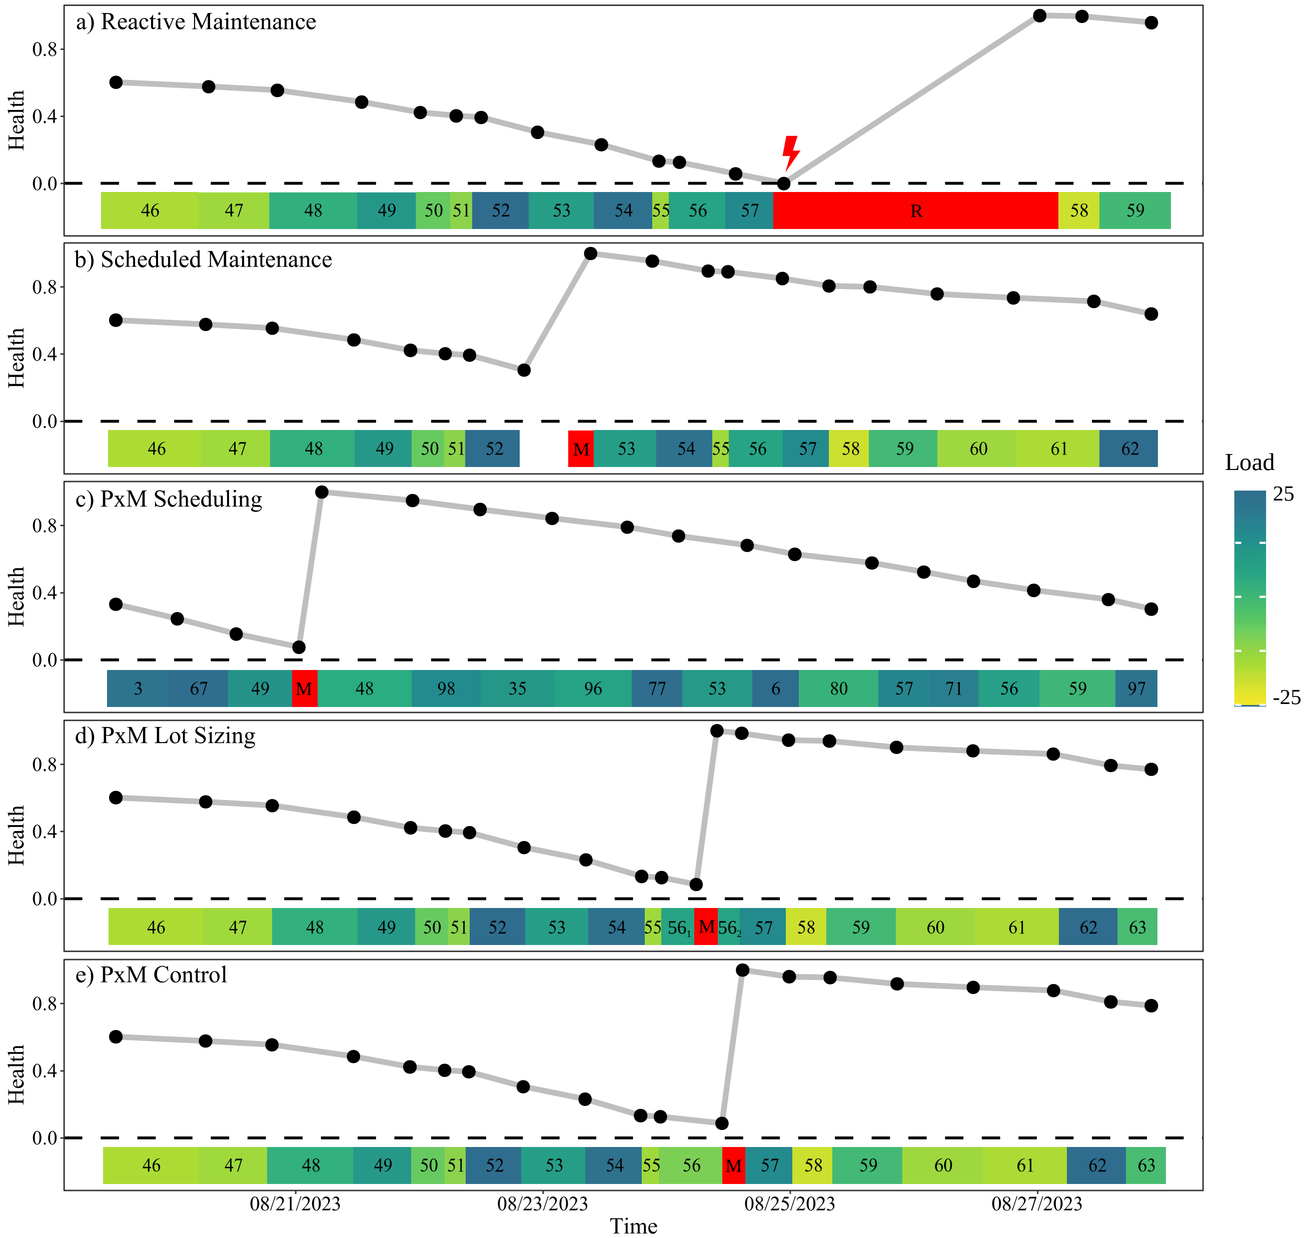

Figure 9. Direct comparison of different strategies

noisy data with our failure prediction algorithm, as shown in Table 3.

With increasing noise, the accuracy stays stable and even increases with high noise levels. The higher the noise, the more observations are classified as non-failure. Because of the imbalance of the classes, this leads to a net benefit in accuracy. However, our main classification objective, recall, suffers from the increased noise, and now, critical false negatives increase. While our recall without noise was 0.91, which means more than nine out of ten breakdowns can be predicted correctly, with a noise level of 0.1, only two of three breakdowns can be correctly identified ($\text{recall} = 0.68$). In contrast, this leads to inversely proportional changes in the precision, while consequently, the F1 score, the harmonic mean of precision and recall, is stable. Additionally, the area under the receiver operator characteristic curve decreases, which indicates worse discrimination between failure and non-failures.

The failure prediction scores indicate that noise might also substantially impact our joint PxM and PPC strategies. To analyze their impact, we simulated the three PxM strategies for approximately one year ($\#O = 800$) over the ten noise levels with 100 replications each ($= 3{,}000$ runs), which took two days and 17 hours to compute. The results are shown in Table 4.

Noise	Accuracy	Recall	Precision	F1 score	ROC AUC
0.00	0.94	0.91	0.42	0.57	0.93
0.01	0.95	0.88	0.43	0.58	0.92
0.02	0.95	0.86	0.45	0.59	0.91
0.03	0.95	0.84	0.46	0.60	0.90
0.04	0.96	0.82	0.48	0.60	0.89
0.05	0.96	0.80	0.48	0.60	0.88
0.06	0.96	0.76	0.49	0.60	0.87
0.07	0.96	0.75	0.50	0.60	0.86
0.08	0.96	0.72	0.52	0.60	0.85
0.09	0.96	0.70	0.52	0.60	0.84
0.10	0.96	0.68	0.54	0.60	0.83

Table 3. Failure prediction under noise

Noise	Scheduling	Lot Sizing	Control	Reactive	Scheduled
0.00	348.02 (±6.56)	**347.14 (±6.77)**	347.92 (±6.92)	400.93 (±7.37)	363.87 (±6.84)
0.01	352.83 (±7.32)	349.46 (±6.40)	**349.02** (±7.04)		
0.02	354.53 (±7.74)	**349.88** (±7.25)	350.25 (±6.80)		
0.03	357.19 (±8.03)	353.46 (±7.30)	**351.28** (±7.45)		
0.04	359.45 (±7.44)	353.56 (±6.70)	**352.83** (±6.81)		
0.05	362.03 (±7.82)	355.06 (±6.54)	**353.50** (±7.08)		
0.06	361.75 (±7.70)	356.67 (±7.12)	**355.19** (±6.76)		
0.07	362.97 (±8.23)	357.20 (±8.12)	**356.93** (±7.54)		
0.08	364.97 (±8.14)	358.55 (±7.57)	**357.99** (±9.12)		
0.09	367.38 (±8.70)	360.24 (±8.29)	**359.71** (±7.67)		
0.10	368.02 (±7.82)	362.35 (±7.54)	**361.94** (±7.16)		

Table 4. Simulated average makespan (± standard deviation) in days under noise

Overall, noise and the decreased failure prediction performance also significantly negatively affected the PxM results. *Joint PxM and scheduling* performance suffers the most. Each percent of noise causes delays in the makespan of multiple days on average. Above a moderate noise level of 0.07, the joint PxM strategy is inferior to a *scheduled maintenance* approach. However, it is not outperformed by reactive maintenance, even at the highest noise levels.

Joint PxM and lot sizing, which outperformed other strategies over most of the 'regular' experiments, consistently exceeds the makespan of *joint PxM and production control* for comparable noise levels when a moderately low noise of 0.03 and higher is reached. Still, while it seems to be inferior to its PxM counterpart, it outperforms *scheduled* and *reactive maintenance* strategies. However, the residual difference at 10% noise is only marginal, and at least *scheduled maintenance* is expected to prove superior when more noise is increased.

Joint PxM and production control is the most noise-resilient strategy. While it also suffers from noise, the makespan increases are minor for each percent noise increment, and it quickly surpasses the performance of other strategies. This is in line with it being the safest joint PxM and PPC strategy, as described previously. Similarly to *joint PxM and lot sizing*, it is still superior to 'traditional' maintenance strategies even at a noise level of 0.1.

6. CONCLUSION

In this paper, we have developed a simulation that compares joint PxM and PPC strategies and showcases whether they outperform other maintenance strategies using industrial data. We have created a conceptual simulation model using health- and load-based failure prediction to implement PxM for production scheduling, lot sizing, and production control. Within an extensive simulation study, we have instantiated our model using industrial data of an extrusion process and compared the performance of the three joint PxM and PPC with 'traditional' maintenance strategies. The failure prediction showed excellent results that translated well into our production planning and control simulation and led to multiple days of makespan savings compared to reactive and scheduled maintenance strategies. Further, a sensitivity analysis based on sensor measurement noise revealed that the PxM-based strategies are robust in uncertain environments.

However, our work also comes with some limitations. First, we used only simple heuristics to integrate PHM knowledge into PPC and did not use advanced optimizations such as mathematical programming, genetic algorithms, or reinforcement learning, which could improve results further. For more information about how this can be achieved for production scheduling, lot sizing, or production control, the reader is referred to the literature analyzed in Section 2. Further, we did not combine multiple PxM-based PPC strategies. For instance, one could first perform prescriptive maintenance and scheduling and use production control or lot sizing if no 'light' order is available. Thirdly, we constructed a health indicator and used it a) as the 'true' health to construct our simulation model and b) as direct input to our joint PxM and PPC strategies. However, this might lead to a machine simulation model with unrealistically good predictions that are hard to attain in practice. Still, the performed sensitivity analysis could indicate that our model might be robust if the health had to be predicted, too. However, our sensitivity analysis is only limited to sensor measurement noise, and uncertainties in processing times and load are not analyzed. Lastly, we only analyzed the makespan and did not regard monetary or other objectives.

This leads to some avenues for future research. First, future PxM models should compare to other advanced strategies and, if not possible, at least to 'traditional' maintenance. Additionally, it is interesting to investigate how joint PxM and scheduling, lot sizing, and production control strategies can be combined. Further, modeling uncertainties of the health indicator and extending sensitivity analyses to incorporate noise in processing time and load could validate the model more realistically. In the same vein, the proposed results should be validated in additional industrial settings, using real-world data and considering variations in machinery types, operational conditions, maintenance practices, or even industrial sectors. Lastly, how PxM solutions can be implemented in real production environments must be investigated, e.g., by investigating the human-in-the-loop.

Still, our work contributes to research by showcasing how complex production environments can be modeled and simulated and by benchmarking different condition-based PPC strategies, ultimately fostering the understanding of complex, dynamic industrial PHM systems. On the practical side, the simulation provides managerial insights on where to implement joint PxM and PPC and supports the justification of PHM business cases.

ACKNOWLEDGMENT

The authors would like to thank the entire team of the project "Prescriptive Maintenance for Production Planning and Control", as well as Anirudh Ravi and Zoi Nikolarakis for their valuable input.

REFERENCES

Bousdekis, A., Magoutas, B., Apostolou, D., & Mentzas, G. (2018). Review, analysis and synthesis of prognostic-based decision support methods for condition based maintenance. *Journal of Intelligent Manufacturing*, *29*(6), 1303–1316. https://doi.org/10.1007/s10845-015-1179-5

Broek, M. A. J. uit het, Teunter, R. H., Jonge, B. de, & Veldman, J. (2021). Joint condition-based maintenance and condition-based production optimization. *Reliability Engineering and System Safety*, *214*. https://doi.org/10.1016/j.ress.2021.107743

Broek, M. A. J. uit het, Teunter, R. H., Jonge, B. de, Veldman, J., & van Foreest, N. D. (2020). Condition-Based Production Planning: Adjusting Production Rates to Balance Output and Failure Risk. *Manufacturing & Service Operations Management*, *22*(4), 792–811. https://doi.org/10.1287/msom.2019.0773

Cadavid, J. P. U., Lamouri, S., Grabot, B., Pellerin, R., & Fortin, A. (2020). Machine learning applied in production planning and control: a state-of-the-art in the era of industry 4.0. *Journal of Intelligent Manufacturing*, *31*(6), 1531–1558. https://doi.org/10.1007/s10845-019-01531-7

Chawla, N. V., Bowyer, K. W., Hall, L. O., & Kegelmeyer, W. P. (2011). SMOTE: Synthetic Minority Over-sampling Technique. Advance online publication. https://doi.org/10.48550/arXiv.1106.1813

Coffman, E. G., Garey, M. R., & Johnson, D. S. (1984). Approximation Algorithms for Bin-Packing — An Updated Survey. In G. Ausiello, M. Lucertini, & P. Serafini (Eds.), *CISM International Centre for Mechanical Sciences. Algorithm Design for Computer System Design* (Vol. 284, pp. 49–106). Vienna: Springer Vienna. https://doi.org/10.1007/978-3-7091-4338-4_3

Dehghan, H. S., Nourelfath, M., & Hajji, A. (2023). A deep learning approach for integrated production planning and predictive maintenance. *International Journal of Production Research.* Advance online publication. https://doi.org/10.1080/00207543.2022.2162618

Divya, D., Marath, B., & Santosh Kumar, M. B. (2023). Review of fault detection techniques for predictive maintenance. *Journal of Quality in Maintenance Engineering*, *29*(2), 420–441. https://doi.org/10.1108/JQME-10-2020-0107

Elbasheer, M., Longo, F., Mirabelli, G., Padovano, A., Solina, V., & Talarico, S. (2022). Integrated Prescriptive Maintenance and Production Planning: a Machine Learning Approach for the Development of an Autonomous Decision Support Agent. *IFAC-PapersOnLine*, *55*(10). https://doi.org/10.1016/j.ifacol.2022.10.102

Fitouri, C., Fnaiech, N., Varnier, C., Fnaiech, F., & Zerhouni, N. (2016). A Decision-Making Approach for Job Shop Scheduling with Job Depending Degradation and Predictive Maintenance. *IFAC-PapersOnLine*, *49*(12), 1490–1495. https://doi.org/10.1016/j.ifacol.2016.07.782

Goby, N., Brandt, T., & Neumann, D. (2023). Deep reinforcement learning with combinatorial actions spaces: An application to prescriptive maintenance. *Computers & Industrial Engineering*, *179*, 109165. https://doi.org/10.1016/j.cie.2023.109165

Guillén, A. J., Crespo, A., Macchi, M., & Gómez, J. (2016). On the role of Prognostics and Health Management in advanced maintenance systems. *Production Planning and Control*, *27*(12), 991–1004. https://doi.org/10.1080/09537287.2016.1171920

Jia, X., Huang, B., Feng, J., Cai, H., & Lee, J. (2018). A Review of PHM Data Competitions from 2008 to 2017: Methodologies and Analytics. In *2018 Annual Conference of the Prognostics and Health Management Society*.

Kiran, D. R. (2019). *Production Planning and Control: A Comprehensive Approach* (1st ed.). Oxford: Butterworth-Heinemann.

Ladj, A., Tayeb, F. B.-S., Varnier, C., Dridi, A. A., & Selmane, N. (2017). A Hybrid of Variable Neighbor Search and Fuzzy Logic for the permutation flowshop scheduling problem with predictive maintenance. *Procedia Computer Science*, *112*, 663–672. https://doi.org/10.1016/j.procs.2017.08.120

Ladj, A., Varnier, C., Tayeb, F. B.-S., & Zerhouni, N. (2017). Exact and heuristic algorithms for post prognostic decision in a single multifunctional machine. *International Journal of Prognostics and Health Management*, *8*(2).

Laguna, M., & Marklund, J. (2018). *Business Process Modeling, Simulation and Design*. Boca Raton, FL: Chapman and Hall / CRC. https://doi.org/10.1201/9781315162119

Liang, Z., Zhang, L., & Wang, X. (2023). A Novel Intelligent Method for Fault Diagnosis of Steam Turbines Based on T-SNE and XGBoost. *Algorithms*, *16*(2), 98. https://doi.org/10.3390/a16020098

Liu, Y.-Y., Chang, K.-H., & Chen, Y.-Y. (2023). Simultaneous predictive maintenance and inventory policy in a continuously monitoring system using simulation optimization. *Computers and Operations Research*, *153*. https://doi.org/10.1016/j.cor.2023.106146

Maaten, L. van der, & Hinton, G. (2008). Visualizing Data using t-SNE. *Journal of Machine Learning Research*, *9*(86), 2579–2605.

Medjaher, K., Zerhouni, N., & Baklouti, J. (2013). Data-driven prognostics based on health indicator construction: Application to PRONOSTIA's data. In *European Control Conference (ECC)*.

Meissner, R., Meyer, H., & Wicke, K. (2021). Concept and Economic Evaluation of Prescriptive Maintenance Strategies for an Automated Condition Monitoring System. *International Journal of Prognostics and Health Management*, *12*(3). https://doi.org/10.36001/ijphm.2021.v12i3.2911

Pinciroli, L., Baraldi, P., & Zio, E. (2023). Maintenance optimization in industry 4.0. *Reliability Engineering & System Safety*, *234*, 109204. https://doi.org/10.1016/j.ress.2023.109204

Sankararaman, S. (2015). Significance, interpretation, and quantification of uncertainty in prognostics and remaining useful life prediction. *Mechanical Systems and Signal Processing*, *52-53*(1), 228–247. https://doi.org/10.1016/j.ymssp.2014.05.029

Saxena, A., Celaya, J., Balaban, E., Goebel, K., Saha, B., Saha, S., & Schwabacher, M. (2008). Metrics for evaluating performance of prognostic techniques. In *2008 International Conference on Prognostics and Health Management, PHM 2008*.

Schmidt, M., & Schäfers, P. (2017). The Hanoverian Supply Chain Model: modelling the impact of production planning and control on a supply chain's logistic objectives. *Production Engineering*, *11*(4-5), 487–493. https://doi.org/10.1007/s11740-017-0740-9

Sridhar, S., & Sanagavarapu, S. (2021). Handling Data Imbalance in Predictive Maintenance for Machines using SMOTE-based Oversampling. In *2021 13th International Conference on Computational Intelligence and Communication Networks (CICN)* (pp. 44–49). IEEE. https://doi.org/10.1109/CICN51697.2021.9574668

Teubert, C., Jarvis, K., Corbetta, M., Kulkarni, C., & Daigle, M. (2023). ProgPy v1.5 [Computer software]. Zenodo: Zenodo.

Wang, T., Yu, J., Siegel, D., & Lee, J. (2008). A similarity-based prognostics approach for remaining useful life estimation of engineered systems. In *2008 International Conference on Prognostics and Health Management, PHM 2008* (pp. 1–6). https://doi.org/10.1109/PHM.2008.4711421

Wesendrup, K., & Hellingrath, B. (2023). Post-prognostics demand management, production, spare parts and maintenance planning for a single-machine system using Reinforcement Learning. *Computers & Industrial Engineering*, *179*, 109216. https://doi.org/10.1016/j.cie.2023.109216

Yang, J., Zhao, X., & Han, M. (2022). Joint optimization of imperfect condition-based maintenance and lot sizing via an availability-cost hybrid factor. *Engineering Reports*, *4*(2). https://doi.org/10.1002/eng2.12462

Zarte, M., Wunder, U., & Pechmann, A. (2017). Concept and first case study for a generic predictive maintenance simulation in AnyLogic™. In *IECON 2017 - 43rd Annual Conference of the IEEE Industrial Electronics Society*.

Zhai, S., Gehring, B., & Reinhart, G. (2021). Enabling predictive maintenance integrated production scheduling by operation-specific health prognostics with generative deep learning. *Journal of Manufacturing Systems.* Advance online publication. https://doi.org/10.1016/j.jmsy.2021.02.006

Zhai, S., Riess, A., & Reinhart, G. (2019). Formulation and solution for the predictive maintenance integrated job shop scheduling problem. In *2019 IEEE International Conference on Prognostics and Health Management, ICPHM 2019.*

Zheng, R., Zhou, Y., Gu, L., & Zhang, Z. (2021). Joint optimization of lot sizing and condition-based maintenance for a production system using the proportional hazards model. *Computers and Industrial Engineering, 154.* https://doi.org/10.1016/j.cie.2021.107157

BIOGRAPHIES

Kevin Wesendrup received his B.A. degree in Information Systems from the Westphalian University of Applied Sciences Gelsenkirchen, Germany, in 2017, and his M.S. in the same field from the University of Münster in 2019. During his studies, he worked as a BI consultant at TECE before joining the Chair of Information Systems and Supply Chain Management as a graduate research assistant in 2019. He is currently pursuing a Ph.D. at the same institution. His research interests are post-prognostics decision-making and prescriptive maintenance for production planning and control.

Bernd Hellingrath received his diploma in Computer Science and Mathematics from the Technical University Dortmund, Germany, in 1988. 2001 he finished his Ph.D. (Dr.-Ing) at the same university. After his Ph.D., he held leading positions at the Fraunhofer Institute for Material Flow and Logistics in Dortmund from 1991 to 2008. In 2006 he was appointed as professor at the Chair of Information Systems at the University of Paderborn, Germany. Since 2008 he has been appointed as professor at the Chair for Information Systems and Supply Chain Management and is director of the European Research Center for Information Systems at the University of Münster, Germany. His research interests include fleet-based prognostics and prognostics and health management in production.