Please cite as follows:
Ruberti, A., Polenghi, A., & Macchi, M. (2024). Maintenance plan adaptation based on health ratings of servitised machines through a fleet-wide machine clustering method. *Journal of Manufacturing Systems*, *77*, 368–383.
https://doi.org/10.1016/j.jmsy.2024.10.001

Journal of Manufacturing Systems 77 (2024) 368–383

Contents lists available at ScienceDirect

Journal of Manufacturing Systems

journal homepage: www.elsevier.com/locate/jmansys

Technical paper

Maintenance plan adaptation based on health ratings of servitised machines through a fleet-wide machine clustering method

Alessandro Ruberti [a,b,1], Adalberto Polenghi [a,*,2], Marco Macchi [a,3]

[a] Politecnico di Milano, Department of Management, Economics and Industrial Engineering, Via Lambruschini 4/b, 20156 Milan, Italy
[b] Applied s.r.l., via Speranza 35, San Lazzaro di Savena, 40068 Bologna, Italy

ARTICLE INFO

Keywords:
Servitisation
Maintenance services
Machine clustering
Machine fleets
Collaborative prognostics

ABSTRACT

The increased requests for value-added services to integrate product performance push manufacturing companies to extend their service offerings to meet customers' needs. In this context, maintenance planning can leverage new possibilities offered by digital technologies for data analytics services. The present research then proposes an approach for maintenance plan adaptation based on a data-driven method applied over a fleet of machines installed in different production sites. The method relies on collaborative prognostics to develop a clustering of machines' behaviour aimed at providing the health ratings of the machines and the subsequent maintenance plan adaptation due to the deviation from the expected behaviour. The method is adopted from the perspective of an Original Equipment Manufacturer, as part of a transformation path towards an advanced provision of digitalization for maintenance service offerings. The method is validated in the context of two lines at selected customer's premises. This demonstrates the viability and effectiveness of adapting the maintenance plans thanks to the data analytics in light of the current behaviour of the machines within the lines.

1. Introduction

The run for servitisation is long-lasting [1] and the increased attention towards value generation for customers, by improving partnerships and relationships, is emphasising the requests for product-related services [2]. In this context, the digital technologies currently available extend the potentiality offered to manufacturing companies to widen product-related services [3–5], enabling them also to pursue sustainability and circularity performance [6–8]. The bundle of digital technologies applied for servitisation includes, among others, Internet of Things, Big data analytics, cloud computing, mixed reality, additive manufacturing, simulation, and Artificial Intelligence [9]. Each technology could unleash several benefits reachable by the service provider and the customers/users such as reduced downtime [10], mitigated risks [11], reduced energy consumption [12] and reduced environmental impacts [13]. This makes companies eager to apply such technologies to their products, to achieve digitally enabled product-related services [14, 15].

In the industrial sector, the Product-Service System (PSS) integrates products and services to meet customer needs [16,17]. As such, Original Equipment Manufacturers (OEMs) are among those companies interested the most in pursuing the servitisation [18], building on the competitive advantage allowed by the knowledge of the system/machine they produce [19]. This is translated in support through the system/machine lifecycle, with particular emphasis on the phases of operations and maintenance [20], with the latter one being highly valued for enhancing system availability and reducing costs [21]. Consequently, leading global OEMs often manage maintenance directly at customer sites [22], enhancing their offerings with additional services.

As part of the provided services, OEMs are incorporating condition-based and predictive maintenance to react to potential failures, with related economic and financial benefits [23,24]. They use fleet-wide data from connected machines in order to adopt advanced management strategies through collaborative maintenance. The benefits of collaborative maintenance may be seen at the operational (field) level, as, for example, fault thresholds could be better defined thanks to the data federation resulting from the fleet [25], and tactical level

* Corresponding author.
E-mail addresses: alessandro.ruberti@polimi.it (A. Ruberti), adalberto.polenghi@polimi.it (A. Polenghi), marco.macchi@polimi.it (M. Macchi).
[1] 0000–0003-1787–7290
[2] 0000–0002-3112–1775
[3] 0000–0003-3078–6051

https://doi.org/10.1016/j.jmsy.2024.10.001
Received 19 April 2024; Received in revised form 2 October 2024; Accepted 4 October 2024
Available online 9 October 2024

supporting integrated maintenance planning [26] based on fleet data; for example, improved RUL (Remaining Useful Life) estimation together with maintenance strategies can be reached [27]. Building on a collaborative maintenance scheme, it is possible to envision an integrated approach by OEMs, leveraging condition-based and predictive maintenance solutions from different production sites and maintenance planning that exploits fleet-wide data.

From an architectural perspective, collaborative prognostics offers a framework usable to develop condition-based services related to industrial assets through digital technologies. This may be seen as a result of the application of "collaborative agents concepts into the fields of prognostics and health management" [28]. This framework deploys real-time algorithms for each machine and utilizes fleet-level algorithms that benefit globally from larger datasets, enhancing prediction accuracy [29]. The data exchange among collaborative machines reveals common characteristics like failures and usage patterns, facilitating cost-effective predictions based on data from similar operating conditions [30].

OEMs may take advantage of the collaborative maintenance scheme. Nevertheless, different challenges have to be faced by OEMs to implement effective maintenance service offerings of this kind. On the technical side, the accessibility of the data should not be given for granted: there are various situations in which data collection and storage are prevented due to privacy, cybersecurity restrictions, or not-ready IT infrastructure [31]. On the managerial side, the requirements posed on the OEMs, are putting at the stack the traditional way of maintenance management: customer companies, with geographically dispersed systems and machines, ask for improved and reactive maintenance plans depending on the machines' degradation control [32]; this leads, on the one hand, to the adaptation of maintenance plans to the real-time behaviours of the servitised systems or machines and, on the other hand, to the need for cost-effective scheduling of the skilled maintenance technicians by OEMs to guarantee service levels worldwide, accounting for reliability and cost performance [33].

Given this context, the research work aims at supporting an OEM in maintenance plan adaptation for monitored production systems given limited access to machine data. The underlying assumption is that the OEM already released an anomaly detection monitoring service and related historical data are available; this part is out of the scope of this work, but it is a pre-requisite as the algorithm/s release useful information in terms of anomalies, alerts or alarms. A supervisory role is added as part of the proposal of the present work. It leverages a collaborative approach between machines in order to identify similar machines' behaviour as the primary goal. In particular, the maintenance plan adaptation method has, as the core part, a data-driven clustering of machines' behaviour in a common baseline, allowing the identification of deviation from the expected behaviour: the deviation is tracked as a health rating of the machines under observation; eventually, the deviation, as a change in the health rating, triggers the maintenance plan adaptation from on-shelf plan options already defined by the OEM. Overall, this is implemented in two architectural levels, in coherence with the collaborative maintenance approach. At the operational level, condition monitoring is locally put in place at each machine, aimed at the anomaly detection algorithms as a pre-requisite of this research work. At the tactical level, namely the novelty this work promotes, a machine clustering method is adopted on a global scale, aimed at making a similarity evaluation of the machines, according to their conduction, in the fleet-wide scope; finally, the similarity evaluation is used to decide the delivery of the maintenance plan when needed, based on the health rating. This is the novelty this research work claims, which finally consists of a supervisory control loop that, stemming from the information from anomaly detection algorithms and through the health rating, allows to automatically adapt maintenance plans.

The document is structured as such: Section 2 analyses the maintenance servitisation, with a particular insight on the current advances of condition-based maintenance services offered by OEMs; Section 3 describes the proposed method, nurtured by the collaborative prognostics approach; Section 4 deals with the application of the proposed method to two production lines realising medical devices; Section 5 elaborates over the results of the research, to bring about their managerial implications and to foster further research outlooks; Section 6 draws conclusions and envisions future works to advance the research and industrial practice.

2. State of the art of OEM maintenance service offerings and condition-based maintenance services

2.1. Overview of current advances in OEM maintenance service offerings

The maintenance service offering by OEMs was initially promoted in those sectors that are characterised by high Maintenance, Repair and Overhaul (MRO) spending [34,35], high availability requirements [36] or in which a one-to-one supplier-customer relationship is mandatory [32]. Therefore, the first sectors in which OEMs started to offer maintenance as product-related services are the military and aviation. Afterwards, extensions to other sectors have been experienced. Several studies have been performed in order to shed light on different perspectives on the maintenance service offerings, looking for increased profitability by adding services to OEMs' products, increasing partnership strength, and relying on the OEMs' knowledge of the industrial machines and systems they traditionally offer to the market [19]. Overall, four main trends may be identified. Firstly, contract definition is highly explored given also possible extensions from traditional to performance-based agreements and extended through-life service offerings [37,38]. Secondly, given the collaborative scope of the work between OEMs and customers, researchers focus also on the exploration of collaborative models considering the wide presence of IMPs (Independent Maintenance Providers) with already established services and market share [39]. Thirdly, maintenance planning challenge is faced to facilitate better-tuned plans and adaptation to the assets' age and production requirements [40–42]. Finally, cost-effectiveness driven by maintenance actions aggregation nurtured by statistical models as well as CBM (Condition-based Maintenance) and PdM (Predictive Maintenance) solutions has flourished especially in the last years where the connectivity of assets could be secured [43–45]. It is this last trend that is further explored in the following subsection 2.2, given the scope of this research work.

2.2. Insight on the current advances in condition-based maintenance in OEM service offerings

The deployment of CBM and PdM solutions is determined by almost consolidated methodologies that leverage predefined steps to reach expected goals, especially grounding on PHM as the reference process model [46]. Despite the steps may change according to each researcher's perspective or application details, a common underlying path is present that entails [47–49]: data acquisition, data manipulation, state detection, health assessment and prognostic assessment. Indeed, it is worth noticing that state detection is often declined as anomaly detection or novelty detection [50], while the health assessment does include both the health rating, also designated as health state evaluation, and the diagnosis [51,52]. As documented [53], PHM and related decisions result from a mix of technical and experience-related aspects; for example, the former includes tackling the data connectivity and the model development, while the latter deals with the threshold definition, health rating and maintenance actions judgements [54].

Despite the wide knowledge about CBM and PdM solutions development, their uptake by OEMs to innovate their business models is still limited, and some barriers, or inhibitors, can be spotted [55]. The first family of inhibitors refers to failure data that may become obsolete due to the redesign of components and systems or may be due to highly reliable machines that hence do not provide robust and significant datasets. The second family of inhibitors refers to over- or under-maintenance that implies, on one side, an abundance of preventive maintenance and, on the other side, an abundance of corrective maintenance; in turn, these imply that natural degradation processes are hidden behind the human interventions, preventing the full exploitation of algorithms' capabilities, and, from a business perspective, customers may be not convinced to change maintenance policies as they have different risk aversion propensions with respect to the OEM. These inhibitors work jointly with other two relevant barriers that prevent full exploitation of CBM and PdM as valuable services offered by OEMs: data privacy leading to "data islands" on the side of the OEM and the customer [56], and business model establishment so as to be convenient for both parties [31].

The scientific literature then appears still limited, or not ready, in terms of proving the viability and effectiveness of such service offerings in industrial settings implying customer-OEM relationships in an advanced maintenance strategy that exploits CBM and PdM solutions in an integrated manner. Recent examples include CBM evaluation for an optimal MRO strategy under a performance-based contracting, and concerning availability and cost optimization [57]. Mathematical optimisation is also proposed to balance reliability and CBM in performance-based contracting [58]. One of the latest studies also analyses performance-based contracting considering CBM and spare parts management for better contract definition [59]. However, CBM in these works is intended as inspection-based, that is: a component/machine is monitored at those times in which the OEM planned an activity in the customer's premises without any real-time data analysis. The integration that could be gained from digital capabilities aimed at real-time data analysis for condition monitoring, is not yet evidently shown in real industrial settings.

2.3. Concluding remarks

Despite the relevance of CBM and PdM services as core solutions envisioned in new business models by OEMs, the scientific literature is still not supporting the evidence of a practical deployment that is capable of facing real-industry challenges in a viable and effective approach. On the whole, the challenges mostly refer to three spheres:

- the fulfilment of the PHM process, inhibited by data availability and quality in relation to the failures, and the constraints from the data privacy, which may prevent a full transfer of condition monitoring data by the running machines or systems out of the factory;
- the presence of human-based decisions in the PHM process, which may lead to subjectivity, especially in the definition of threshold, health rating and related actions; this makes service delivery challenging, especially when scaling worldwide;
- the arrangement of an organizational setting and business model where benefits from the CBM/PdM services can be exploited in terms of value-adding and cost-effectiveness for both the OEM and customer companies.

This research work is particularly interested in proposing and demonstrating the viability and effectiveness of the PHM process and the subsequent decision-making in maintenance plan adaptation so that it is performed in an objective way, drawing on past experience, to finally scale up so as to uniform the approach toward an enhanced service offering. This integrated maintenance approach leverages fleet-wide data across the operational and tactical levels of maintenance management tasks.

3. Proposed method for data-driven maintenance plan adaptation

The method leverages the potentialities offered by CBM to extend maintenance service offerings by OEM with a focus on maintenance plan adaptation along the machine lifecycle. The method considers the relevant assumptions hereafter listed:

- The existence of an IT infrastructure able to collect data from the machines and make them available to the OEM itself, where algorithms and digital services can be deployed.
- Condition monitoring and algorithm/s for anomaly detection are already in place on servitised machines (i.e., CBM is already enabled through condition monitoring, running in each servitised machine) and allow the generation of historical data related to machine operations.
- Exclusively sharing machine status, i.e. machine states, is performed to provide the basic information so as to facilitate the identification of similar behaviour/s of machines within the fleet of servitised machines (i.e., while the data logs tracking the machine status are shared across the fleet of machines, the condition monitoring signals, such as vibration, temperature, and others, are not; this is respecting some data privacy requirements on the detailed knowledge of the machine behaviour demanded by the customer).
- A set of pre-defined maintenance plans is already in place, as it is developed by the OEM to cope with the customers' requirements and contingent situations, defined according to past experience and knowledge. These pre-existing maintenance plans are not going to be changed in their content, sequence and structure, but will be automatically engaged upon verifying specific health rating results.

From a functional perspective, the goal of this work is to explore how to use the information coming from the anomaly detection algorithm/s to choose the proper maintenance plan adapting to the ongoing health rating of the machines in a way that is no longer subjective.

Fig. 1 reports where the novelty of this work stands, that is in the proposal of a data-driven method to guarantee objectivity and stability to the maintenance plan adaptation process.

Indeed, during the machine conduction, an anomaly detection algorithm is in charge of detecting deviations from expected behaviour. These anomalies become the knowledge base on which maintenance operators, technicians and managers may decide to confirm or change the maintenance plan based on experience and skills. The proposal this work aims to defend is to substitute this mechanism, which may imply different results as it is subjective, with a data-driven method that will guarantee an objective maintenance plan. Nonetheless, the solution is configured not to be prescriptive, but rather informative to the maintenance team.

In Fig. 2, the proposed data-driven method (yellow box in Fig. 1) is exploded in all of its constituent algorithms and steps (blue boxes), pointing out if the proposed steps should happen on customer premises or, potentially, on the OEM computational platform. The novel part of the method refers to the algorithms and steps running on the OEM computational platform (blue boxes) and the influence they have on the machines' conduction at customer's premises (grey boxes).

The main constituents of the method are hereafter summarized:

a. The creation of the Transition Matrix representing the machine behaviour is triggered by the arising of an anomaly.
b. The application of the Non-negative Matrix Factorization (NMF) helps in discerning between baseline production conduction and the specific behaviour of the machine under study.
c. The comprehension of the machine behaviour is then performed by the application of the Hidden Markov Model (HMM), which – in return – generates a posterior matrix per each set of data collected.
d. The calculation of the similarity index (on the HMM posterior matrix) enables similarity evaluation; in particular, the step performs the comparison of the posterior matrix with the existing benchmark (of past sets of data) of the machine behaviour so as to habilitate health rating.

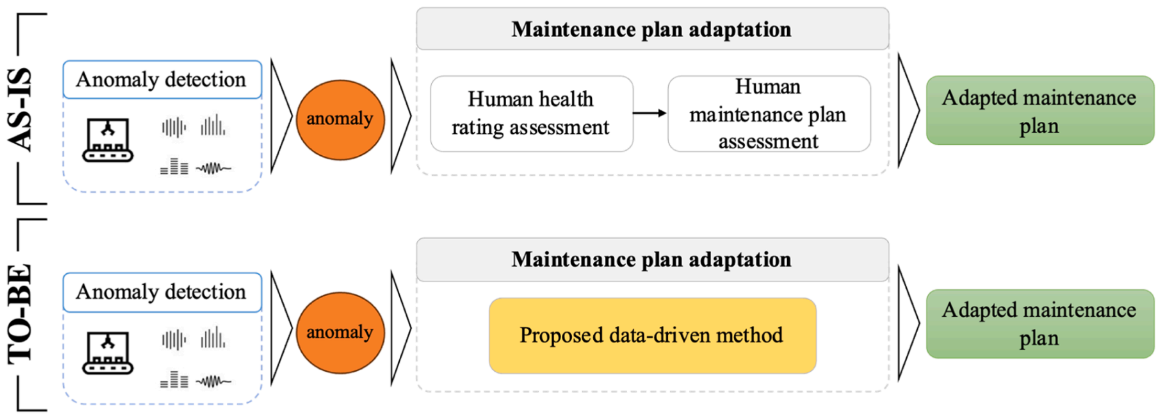

Fig. 1. Novelty of the work: data-driven method (yellow box) that improves maintenance plan adaptation.

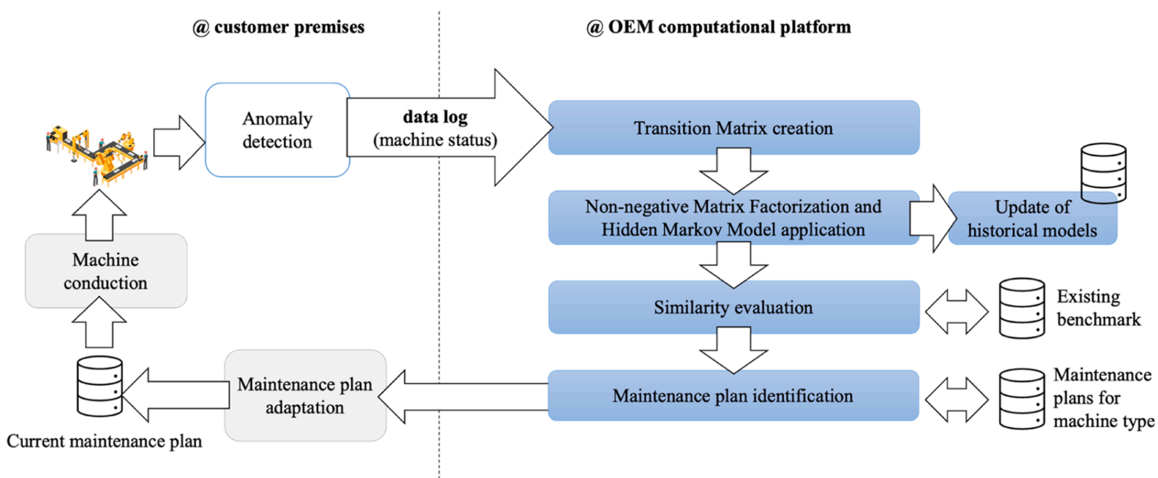

Fig. 2. Proposed method for data-driven maintenance plan adaptation.

e. Similarity with the existing benchmark and the subsequent health rating finally allows to compare the current health state to the designed, known and "accepted" state in order to select the maintenance plan that best fits the latest observed conditions in the machine under study; the customer is then informed whether the maintenance plan adaptation has to be updated on the specific servitised machine for which anomaly has been initially detected.

It is worth remarking that the maintenance plan is not only due to the system design and features – as determined by OEM knowledge – but also delivered in the light of the way the system (and each machine within the system) is operated – which leads to data-driven planning decisions.

For the sake of completeness, the entire data-driven method is presented, including also anomaly detection, limiting it to the functional utility as the triggering step of the method, despite the anomaly detection algorithm, its definition, configuration and development are outside the scope of this work. Each step is then described with the related modelling decisions, inputs, and outputs. Then, before applying the method to a real case (Section 4), the pseudo-code is proposed so as to give a flavor of its implementation view.

3.1. Anomaly detection

Anomaly detection is in charge of detecting possible anomalies in the machine's behaviour, arising during the running of the production operations. The anomaly detection is configured based on the following general assumptions:

- Several signals can be used to enable the detection capability, depending especially on the type of machine under analysis. Vibrations collected via accelerometers are one of the most frequent and widespread collected signals due to their relevance in detecting anomalies in rotary components.
- For what concern data analytics, there are many different techniques and algorithms that can be used for anomaly detection, including statistical methods, machine learning algorithms, and time series analysis [60]. The selection is based on the gathered signals, the availability of labels and the computational power at the stack.

In the proposed method for data-driven maintenance plant adaptation, anomaly detection triggers and starts the key steps of the method with the purpose of verifying the validity of the current maintenance plan and, when not validated, replan the maintenance activities based on current machine behaviors and existing benchmarks provided by the past machine conditions and states.

It is also worth remarking that the proposed method works under the assumption of limited access to data, which corresponds to transferring only machine status as data logs from the customers to the OEM. In other words, the detailed information arising from the condition monitoring at every single servitised machine is used for the CBM onsite, but not transferred to the OEM, which receives only data logs tracking the machine status.

3.2. Transition Matrix creation

The data logs, consisting of a sequence of machine states, is transformed into the so-called Transition Matrix. The machine states of the production asset are defined according to the European Standard 415–11 "Safety of packaging machines – Part 11: Determination of efficiency and availability" published in 2021 [61] and are collected by a monitoring tool. With these machine states, a Transition Matrix, also known as probability matrix or Markov matrix, is generated; the Transition Matrix is a mathematical transformation used to model a system that undergoes transitions from one state to another over time. It is commonly used in fields such as economics, finance, physics, and engineering to analyze the behaviour of complex systems.

In a Transition Matrix, both the rows and the columns represent the machine states of the system under study and each value cell in the matrix represents the probability of transitioning from one state to another; for example in Fig. 3, the numeric value inside the matrix cell at the crossing between row "2" ("state #2″) and column "3" ("state #3″) is the probability in percentage that, from the present state "state #2″, the machine will go into state "state #3″. The machine states may include, for example, disconnection and errors. Therefore, the entries in the matrix are typically non-negative, and the sum of the entries in each row must be equal to 1, representing the fact that the machine must make a transition to a new state at a given time step.

The Transition Matrix can be used to compute the probabilities of different sequences of transitions over time, as well as to analyze the long-term behaviour of the system. For example, the dominant eigenvalue of the matrix can be used to compute the long-term behaviour of the system, while the eigenvectors can be used to identify the most likely paths or sequences of transitions. In the specific case of maintenance plan adaptation, the Transition Matrix enables the analysis of the machine condition as well as possible deviations in terms of state change.

3.3. Non-negative matrix factorization application

The first transformation of the Transition Matrix is through the NMF, later followed by the HMM. The adoption of the NMF comes from the need to reduce the data complexity from the multiple machines and aims at isolating the relevant information to properly enable a collaborative prognostics approach. Therefore, the goal is to extract information about machine conduction.

An NMF is a technique for factorizing a non-negative matrix into two or more non-negative matrices. It is widely used for feature extraction, pattern recognition, and data compression [62]. NMF is a form of unsupervised learning, meaning that it does not require labelled data or the knowledge of pre-existing conditions to perform the factorization.

In this work, the goal of NMF is to find two non-negative matrices, W and H, such that their product approximates the original non-negative matrix X as closely as possible, where the matrix X represents the original Transition Matrix. The general representation of the method is represented as in Fig. 4.

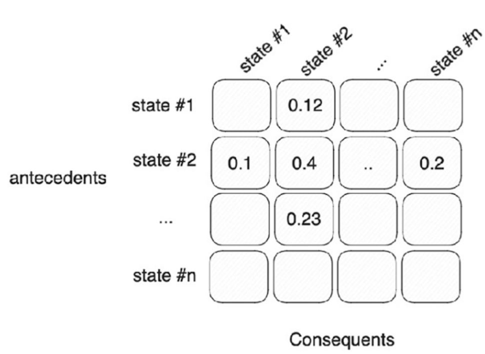

Fig. 3. General representation of a transition matrix.

The matrix W is a matrix of basis vectors or *base matrix*, where each column represents a basis vector that captures a particular feature or pattern in the data. The matrix H is a matrix of coefficients, where each row represents the coefficients that determine how much each basis vector contributes to the observations in the Transition Matrix X.

The factorization is performed iteratively, using an optimization algorithm on the NMF objective functions based on the Frobenius norm to update the values of W and H until a good approximation of the original matrix is achieved. The solution of the optimization algorithm can depend on the specific problem and the characteristics of the data; in general, it can be solved either with gradient descent, alternating least squares (ALS) or multiplicative updates of the matrices [63,64].

One of the main benefits of the NMF is that it can produce sparse representations of the data, meaning that only a small number of basis vectors are required to represent the data accurately. This can be useful in reducing the dimensionality of the data and removing noise or irrelevant features. NMF can be also used for feature extraction, where the basis vectors are used as a reduced set of features for further analysis.

Considering a procedural viewpoint of NMF application in the proposed method, a few steps are required:

a. The NMF technique is developed upon a batch of historical data composed of past Transition Matrices, collected over a long period of time: in this way, it is possible to calculate the common base matrix shared among all the (collected and calculated) Transition Matrices.
b. The application of the NMF upon historical data then grants the ability to identify a unique base matrix (identical per each machine type) thanks to Eq. 1:

$$(H_1, H_2, \ldots, H_n) = NMF\{W * (X_1, X_2, \ldots, X_n)\} \tag{1}$$

where:
X_i = Transition Matrix in input of the event *i*
H_i = Coefficient Matrix in output of the event *i*
W = common Base Matrix in output

c. As the base matrix W is known, when a new event is recorded (i.e., a new Transition Matrix X_{curr_date} is created), a function (f^{NMF}) based on the NMF and on the given base matrix W generates the new specific Coefficient Matrix related to the new event (Eq. 2):

$$f^{NMF}(X_{curr\_date}, W) = H_{curr\_date} \tag{2}$$

It is worth remarking that, in order to apply the NMF, the batch of historical data is essential and should be made available as a prerequisite: the batch is created upon the event *i* traced along the history of the machine conduction, being it an anomaly detection. Considering the multitude of machines in a fleet (i.e., machines of the same type), the NMF and, in particular, the base matrix W can then be obtained during the training phase (i.e., step b), taking advantage of the machines featuring similar patterns in terms of anomalies along their history. Any new anomaly arising during the runtime (step c) should be evaluated in its deviation from the expected behaviour, resulting from the training phase.

On the whole, the NMF step is key for what follows as the common base matrix W holds the information about the machine type, its core features and its intended or predefined usage, while the coefficient matrix (H_i) holds only the specific information regarding the event and, specifically, the conduction that led it to that.

3.4. Hidden Markov Model application

After the NMF, the Hidden Markov Model (HMM) is a statistical tool used to model sequences of observed events linked to unknown or hidden states. Commonly used in speech and handwriting recognition, bioinformatics, and financial time series analysis [65]. HMMs trace

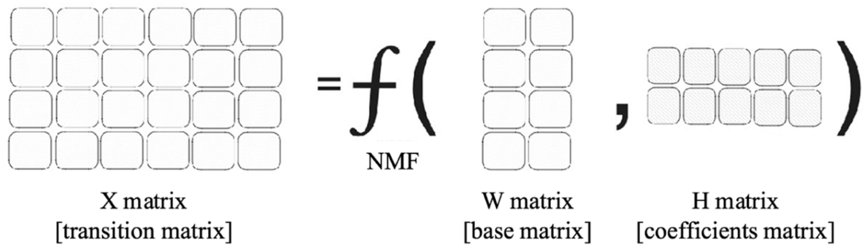

Fig. 4. General representation of the Non-negative Matrix Factorization.

events back to invisible generating states. These hidden states are modelled as a Markov process, with the system's current state only influenced by its previous state. Each observation is a probabilistic function of its hidden state, aiming to estimate the sequence of hidden states that most accurately explains the observed data.

The development of an HMM algorithm centres on exploring these hidden states and their dynamics, recognizing them as the true drivers behind observable events, regardless of whether anomalies are detected or if the system operates normally. The sequences of hidden states effectively 'emit' the visible states corresponding to observable events, allowing probabilistic insights into the causes behind these occurrences. This marks a shift from examining the effects to directly analyzing the causes.

Mathematically speaking, the HMM consists of three components, as represented in Fig. 5: the state transition probabilities or posterior probability matrix, the emission probabilities matrix (usually indicated with B), and the initial state distribution matrix. The posterior probability matrix describes the probability of transitioning from one hidden state to another.

The emission probabilities matrix describes the probability of observing a particular observation given the hidden states. The initial state distribution describes the probability of starting in each of the hidden states. Therefore, whenever a common emission probabilities matrix is identified, each observed event is connected to the previous event in the observed sequence via the initial states which, combined with the common matrix B, generates the posterior probability matrix. The HMM algorithm is then used to estimate the sequence of hidden states that best explains the observed data.

The Baum-Welch algorithm [66] is the algorithm used in this specific analysis and is a variant of the Expectation-Maximization (EM) algorithm used to estimate the parameters of the model from a set of observed data. Using the batch of original input data, the algorithm iteratively estimates the parameters of the model to maximize the likelihood of the observed data [67].

Considering a procedural viewpoint of HMM application in the proposed method, a few steps are required:

a. A common emission probabilities matrix B is identified for all the machines of the same type by using historical data.
b. As a posterior probability matrix P depends both upon the initial states and the common emission probabilities matrix B, a model is trained using known sequences of observed events to optimize the weights as in Eq. 3 where H matrices represent the Coefficient Matrices from NMF:

$$(P_1, P_2, \ldots, P_n) = \mathrm{HMM}\{B^*(H_1, H_2, \ldots, H_n)\} \tag{3}$$

where:

P_i = Posterior Probability Matrix in input of the event *i*

H_i = Coefficient Matrix in input of the event *i* (i.e., output of the previous NMF corresponding to event *i*), as the initial state in the HMM model

B = common emission probabilities matrix

c. As the emission probabilities matrix B is known when a new event is recorded (i.e., a new coefficient matrix H_{curr_date} is created and becomes the new initial state) a function (f^{HMM}) – based on the HMM and on the given base matrix B – generates the new specific posterior matrix [P_{curr_date}] related to the new event [H_{curr_date}] (Eq. 4):

$$f^{HMM}(H_{curr\_date}, B) = P_{curr\_date} \tag{4}$$

The HMM model and, in particular, the common emission probabilities matrix B can then be obtained during the training phase (i.e., step a and b) using historical data by taking advantage of machines of the same type in the fleet, and with similar patterns of anomalies. Any new anomaly arising during the runtime (step c) should be evaluated in its deviation from the expected behaviour embedded inside the HMM model.

More specifically, the HMM step of the method is key for what follows as the common emission matrix B holds the information about the relations for that specific machine type between (hidden) states, while the posterior probability matrix (P_i) holds the specific information regarding the causes of what will be observable manifestations in the machine. In terms of maintenance issues, this means that modelling the hidden states operating on the system/machine and the transition from one machine state to another translates into the ability to know when the machine is moving away from a specific health state before a manifestation occurs: from state of healthy production to an anomalous state. In particular, whenever the cause of an observable manifestation becomes evident through the posterior probability matrix which drifts and changes before an event becomes observable in the machine, this leads to a practical recognition of the health state and subsequent impact on the maintenance activities (and therefore plan) adopted as countermeasure of the cause of the observed event. The posterior matrix is then the information adopted to move further on to the next step of the method.

3.5. Similarity evaluation and maintenance plan identification

The similarity evaluation takes place to match the newly obtained posterior probability matrix (i.e., output of the previous HMM step) with past, existing ones from historical records: by assessing the similarity of a new posterior probability matrix with past matrices, the health rating is then habilitated, and it becomes possible to draw certain conclusions on the healthiness of the machine, finally taking into consideration the possibility of a maintenance plan adaptation.

More specifically, the posterior probability matrices of past events constitute the 'memory', traced in the historical data of the machine in the light of matrices W and B as information of the correspondent machine type. As such, it contains information on the conduction of the machine and its states (production, stop, failure, product quality deviation, etc.) together with the health state of the machine itself.

This determines the ability to detect when health rating changes and, therefore, to potentially adapt maintenance plans so to return to the ideal health status of the machine by observing the conduction and the

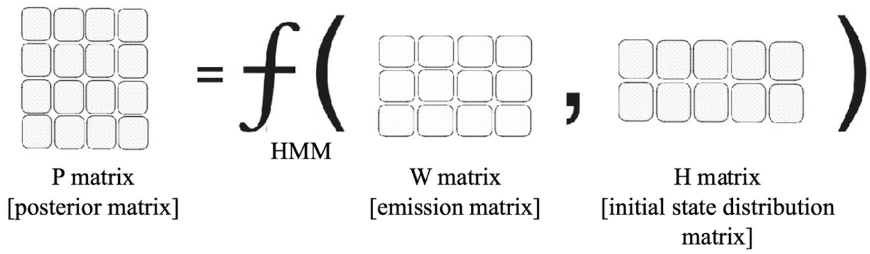

Fig. 5. General representation of the Hidden Markov Model.

sequence of the related machine states: this is inherently due to the fact that the cause of an observed event, at the basis of the health rating, can be related to the recommended maintenance activities as countermeasures. This depends on the OEM engineering knowledge. Indeed, a set of maintenance plans is already in place as countermeasures of the different causes of observable events; the plans are then picked up based on the evaluated machine behaviour, which comes with the posterior probability matrices.

To make the similarity evaluation, the following requirements apply with respect to the historical data:

- The historical database should be built by collecting a proper number of production periods (e.g., weeks) coming from different machines in complex production systems over a period of operations.
- The historical database should be then analysed and clustered and the posterior matrices of past events should be split into a number of groups based on their health label. In this work, 3 groups are considered: matrices of past events that underwent minimal or regular and standard maintenance interventions [group #1, this corresponds to causes of observed events that require standard maintenance], matrices of past events that underwent significant but not critical maintenance interventions [group #2: this maintenance requires the substitution of parts subjected to wear and fatigue, and is required by more impacting causes of observed events] and finally matrices related to critical failures and urgent maintenance interventions [group #3, this can mean substitution of hardware components as major equipment/functional unit of the machine, which is the worst case when the causes of observed events are the most impacting ones].

Once the historical database is available, for each machine type and for each of the identified historical cases of the individual machines of that type, there is the ordered sequence (in time) of the posterior probability matrices generated by the algorithm (from the HMM step) and that constitutes the history of the machine states: when a new posterior probability matrix is recorded, this one is compared by similarity with each cluster and their latest (in time) posterior matrix.

The similarity is calculated using the Jaccard method, which is one of the general-purpose similarity coefficients used in the industry [68]. Jaccard similarity is calculated by dividing the size of the intersection of the two sets by the size of the union of the two sets. In other words, the Jaccard similarity coefficient is the number of items that the two sets have in common divided by the total number of distinct items in both sets. For example, consider two sets: set A contains {1, 2, 3} and set B contains {2, 3, 4}. The intersection of the two sets is {2, 3}, and the union of the two sets is {1, 2, 3, 4}. Therefore, the Jaccard similarity coefficient of A and B is 2/4 or 0.5.

Jaccard similarity is commonly used in data mining, machine learning, and information retrieval applications [69]. It is particularly useful when dealing with categorical data, such as text documents, where the sets represent the presence or absence of certain terms or features.

One advantage of Jaccard similarity is that it is simple to compute and can be calculated quickly for large datasets. It also provides a measure of similarity that is independent of the size of the sets being compared. However, Jaccard similarity has some limitations. Firstly, it does not take into account the frequency or importance of the items in the sets, and it can be biased towards sets with a larger number of items, even if the sets have low overall similarity. In addition, Jaccard similarity may not be appropriate for datasets with a high degree of sparsity or where the sets are of different sizes.

Within the algorithm, the value of Jaccard similarity is used for the attribution of similarities and eventual labelling for health rating: the new (latest in time) posterior probability matrix will be then associated with the cluster with the highest similarity. Each cluster of similarity is characterized by a specific level of health rating: the first level (with the highest health rating) corresponds to an Overall Equipment Effectiveness[4] above 85 % for all the labelled periods under consideration; the second level has an OEE between 70 % and 85 %; while the level (#3) has OEE below the 70 %. Upon this cluster attribution, the health rating is updated in order to reflect the latest conditions of the machine; subsequently, the maintenance plan may also be adapted or not. More specifically, the following cases may happen, as synthesized in Table 1.

Summarizing, all the above-mentioned levels (#1, #2, #3) are essentially different levels of health ratings based on posterior probability matrices comparison, as represented in Fig. 6.

The maintenance plans to be adopted by the user of the machines are correspondingly adapted, when it is the case (only levels #2, and #3).

Table 1
Attribution of similarities, health rating and maintenance plan adaptation.

Health level	Description
Health level #1	If the latest recorded matrix of the individual machine under control is similar (as Jaccard Similarity) to the cluster of posterior probability matrices of past events that have a good health rating, then no further, additional maintenance is needed; the machine stays in health level #1 or "healthy" state.
Health level #2	If the latest recorded matrix is similar (as Jaccard Similarity) to the cluster of posterior probability matrices of past events that have worse (compared to the level #1) health rating, then additional maintenance might be necessary such as increase in frequency of inspections, mandatory clean-ups at the end of each production shift, daily system deep analysis, log check and alarms analysis; the machine is considered in health level #2 or "monitored" state.
Health level #3	If the latest recorded matrix is similar (as Jaccard Similarity) to the cluster of posterior probability matrices of past events that have much worse (compared to the level #1) health rating, then an important set of maintenance interventions might be necessary such as complete stop of the machine, complete clean-up, deep maintenance service and substitution of key components; finally, the machine is considered in health level #3 or "critical" state.

[4] As well known, the OEE formula is the following: OEE $=$ Availability $\times$ Performance $\times$ Quality

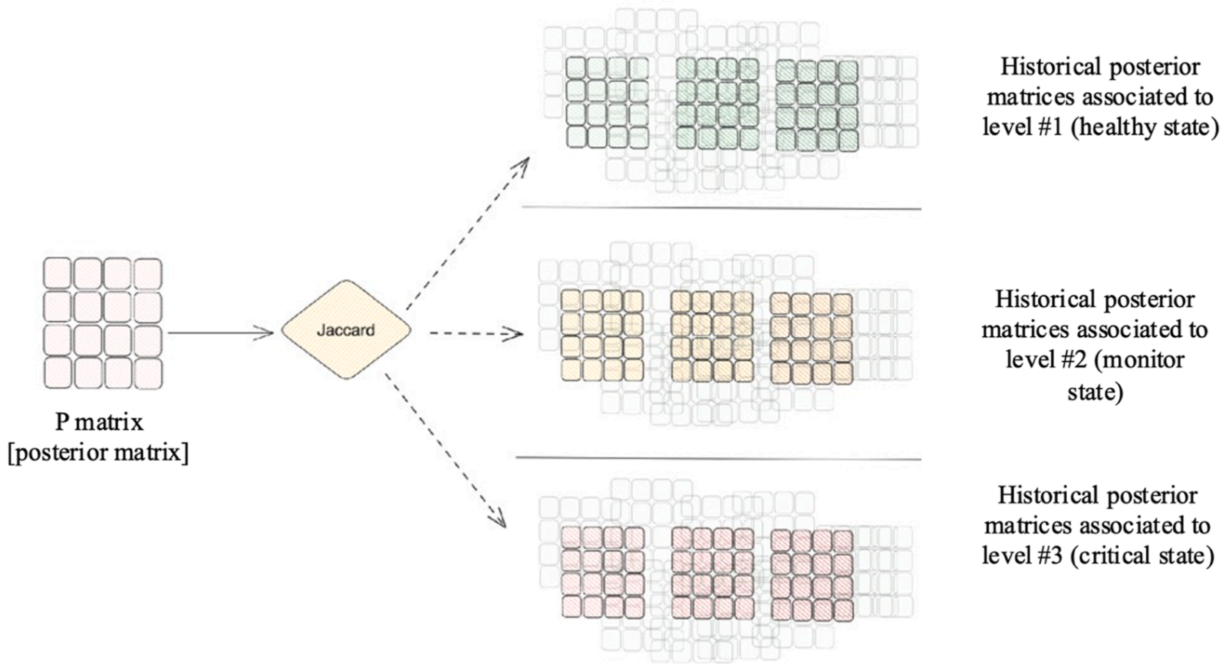

Fig. 6. General representation of the Jaccard similarity for the state, and related maintenance plan, identification.

3.6. Implementation view of the proposed method

The following Fig. 7 reports the pseudo-code of the method proposed in this Section 3 and summarised in the above Fig. 2. This provides a general idea of the implementation view, as well as a summary perspective, of how the underlying algorithm is organised.

4. Application to an OEM company with maintenance service offerings

The data-driven maintenance plan adaptation method is now assessed via a real case of an OEM providing maintenance services to production lines run on customers' premises. This builds on the willingness to better define the maintenance plan over time adapting to the health rating, considering the customers' need to keep the production lines the healthiest by having a prompt response from a maintenance

```
START Algorithm

DECLARE anomaly_event  # this variable is a Boolean: 0 means "there is no anomaly", while 1 means "there is an anomaly"
DECLARE motor^i  # this variable is the identification of one of the brushless motors
DECLARE ALL_MOTORS  # these are all the motors under monitoring
DECLARE ARRAY motor_variables^i  # this array contains the relevant variables extracted from an  operating brushless motor such as current, velo
                                   acceleration, position

WHILE anomaly_event == 0:
    anomaly_event = anomaly_detection(motor^i , motor_variables^i )

IF anomaly_event == 1:
    window-of-time = time( hour-of-the-day , hour-of-the-day + 7 days)
    dataset^window-of-time = collect_data_log(duration = window-of-time , variable = 'machine status')
    X_matrix ^window-of-time  = transition_matrix( dataset^window-of-time )
    Coefficient_Matrix ^window-of-time  = NMF(Common-Base-Matrix , X_matrix ^window-of-time)
    Posterior_Matrix ^window-of-time  = HMM(Common-Emission-Matrix , Coefficient_Matrix ^window-of-time)
    health_status = Jaccard_Similarity(Posterior_Matrix^window-of-time) # admissible values for health_status is "healthy", "monitor" and "critical"

IF health_status == "healthy":
    PRINT("no change in the maintenance plan")
ELSE IF health_status == "monitor":
    ENABLE  New-Maintenance-Policy-LEVEL-1 # system deep analysis, log check, alarms analysis, more frequent inspection
    PRINT("maintenance policy changed")
ELSE IF health_status == "critical":
    ENABLE  New-Maintenance-Policy-LEVEL-2 # plan for a complete stop of the machine and deep maintenance service
    PRINT("maintenance plan improved")

END Algorithm
```

Fig. 7. Pseudo-code of the method for data-driven maintenance plan adaptation.

standpoint.

To this end, the method has been applied to 8 production lines installed in two geographically dispersed plants of the same customer: all the production lines are quite similar and composed of modules. The production lines are furthermore making a sequence of automated, sequential operations leading to the realisation of the final product, which is a class IIa medical device for cosmetic or therapeutic purposes. A class IIa medical device is a non-invasive device used for channelling or storing blood, body liquids, cells or tissues (https://health.ec.europa.eu/). As such, the product must guarantee the highest possible quality, reflecting the need to have always control over the health states of the machines along the production line, so as to guarantee no operations are uncompliant with respect to the predefined requirements.

The essential components of the machines are brushless motors, which generate the movements needed to operate all mechanical parts. A brushless motor, also known as an electronically commutated motor, is a synchronous motor powered by direct current (DC). It operates with an electronic controller that switches DC currents to the motor windings, creating magnetic fields that the rotor follows. This controller fine-tunes the DC current's phase and amplitude to manage the motor's speed and torque.

For organizational efficiency, production lines are divided into sections, each performing specific tasks within the production sequence. Each section includes five brushless motors managed by a dedicated Programmable Logic Controller (PLC). This PLC monitors the motors' position, velocity, acceleration, and current, and also tracks important data for each section, such as states and alarms. Notably, while all brushless motors are identical across all production lines and sections, the load varies depending on the specific tasks each motor performs.

Since the initial installation, the production lines have been inspected once a week by both the OEM and the customer to verify the conditions of the entire lines and also to assess the machine's health state.

The available dataset has been generated over 241 production weeks of monitoring of 14 sections of 2 production lines over a period of 3 months. The historical database is analysed and clustered and the posterior probability matrices of past events are split into 3 clusters - each with a specific level - based on their health label: matrices of past events with good health ratings [level #1], matrices of past events with higher health ratings where the motors underwent significant but not critical maintenance interventions as it is in the "monitor" state [level #2] and finally matrices related to low health ratings which showed critical failures and urgent maintenance interventions [level #3].

The assessment of the belonging group is performed manually by the operators, based on their experience and on a set of multiple pieces of information, namely, the log of the machine states and related alarms of a given week, and the register of maintenance interventions performed by the operators for regular setups, standard maintenance actions or exceptional interventions for the same week of analysis.

Based on this information, the practice is to classify any monitored window of production period (typically a week) in the three possible clusters of health states (expressed according to the current OEM-customer agreement and documented in accordance with what is generally presented in the previous Section 3.5) that are:

1.11. [level #1] good health ratings which means a week with no anomalies or with anomalies that, when detected, proved to be not harmful to the motors and, consequently, to the quality of the output: this level is called "healthy".
1.12. [level #2] worse than level #1 health rating, which means a week with detected anomalies causing minor impacts on the production with low effort but regular interventions: this level is called "monitor" state.
1.13. [level #3] Critical/low health rating, which refers to a week with recorded anomalies with potential impacts on the machine's health and therefore on the product quality and, hence, that must be dealt with as soon as possible: this level is called "critical".

This is a practice not uncommon in OEM-customer relationships as proven by Wang et al., 2020 [58].

According to the levels above, when a new production period is analyzed, it is determined the belonging cluster so that the maintenance is planned accordingly. For example, if in the previous week, the maintenance team (or the future data-driven method, see next subsections) determined that a given section of the production line was in level #1 and in the current week in level #2, a new maintenance plan will be issued accordingly to improve the control of the machine's health. In case, for n-periods in a row, the belonging group/cluster stays the same, the defined maintenance plan is kept unchanged.

It has to be noted that, contrary to the initial study phase, at regime the evaluation and the run of the full method will be done discretely, only when the anomaly detection algorithm is triggered. In synthesis, the proposed method has been applied so to develop a data-driven algorithm able to automatically inform about the need of maintenance plan adaptation, just when the anomaly is detected.

In subsection 4.1 the proposed method is applied and intermediate results are shown; then, in subsection 4.2 the assessment of the deployed solution is reported to support the goodness of the reached results.

4.1. Application of the proposed method

Due to the significant number of data and alerts that may occur in complex production lines, a surveillance algorithm is continuously in charge of triggering the evaluation of the ongoing production period (typically a week) in order to assess the need to evaluate any change in the current maintenance plan. The surveillance algorithm consists of an unsupervised anomaly detection model.

Even though the development of the anomaly detection algorithm is outside the scope of this work, the proposed model and reached results are reported for the sake of completeness.

The selected unsupervised algorithm is an isolation forest algorithm, which is often used as a machine learning algorithm for anomaly detection [70]. It is a tree-based algorithm that works by randomly partitioning the data points into subsets, or "isolation trees," until the anomalies are isolated.

The algorithm starts by selecting a random feature and a random value within the range of that feature. The data is then partitioned into two subsets based on whether each point is above or below the selected value. This process is repeated recursively, with each new subset being split into two based on a randomly selected feature and value. This creates a binary tree structure where each leaf node represents an isolated subset of data points.

Once the isolation trees have been constructed, the anomalies are identified as the data points that are isolated in small subsets. The rationale behind this is that anomalies are likely to be isolated more quickly than normal data points, as they tend to have different values or characteristics that make them stand out.

One of the advantages of the unsupervised isolation forest algorithm is that it does not require labelled data for training. This means that it can be used for anomaly detection in datasets where anomalies are rare or difficult to identify [71].

By modelling the unsupervised isolation forest for the brushless motors under analysis, it is possible to look for anomalies whenever they arise. Indeed, Fig. 4 reports an example over four days of monitoring of the vibration signals. The red dots are the anomalies classified as potentially dangerous for the machine..

To avoid and prevent misclassification of events and minimize noise, several steps are taken: an algorithm monitors mid-to-long-term drifts within the time series to detect shifts; the data is then de-noised by smoothing the original dataset and scaled and normalized using the Min-Max algorithm; finally, the outcomes of the selected anomaly detection algorithm are compared with other models like autoencoders and Local Outlier Factor to enhance robustness by compensating any limitation of each single models.

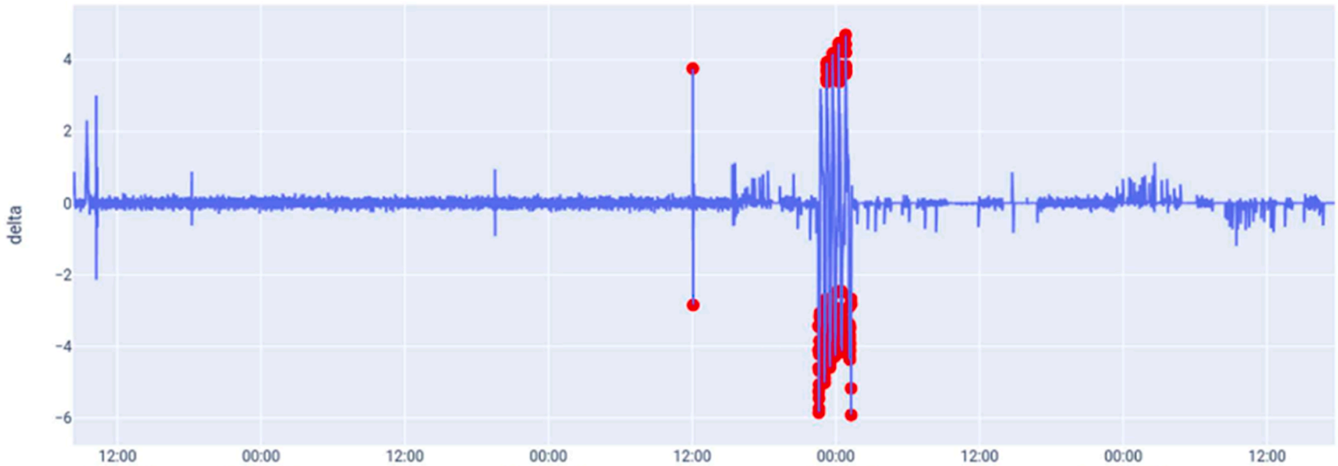

Fig. 8. Application of unsupervised isolation forest for vibrations of brushless motors.

As the algorithm picks up an anomaly, it triggers the core of the proposed method, that is the calculation of the Transition Matrix on the previous production period, application of NMF as the main transformation, subsequent modelling through HMM and the similarity evaluation.

When the surveillance algorithm is triggered, the Transition Matrix based on the machine states is evaluated (see Fig. 9); in this specific case, the brushless motors may experience ten states.

The Transition Matrix is then reduced by an NMF operation which removes the common base of the starting matrix and achieves the reduction of the complexity as well as highlights the specific situation of the analysed section of the production line. Afterwards, the HMM is applied and a new posterior matrix is generated: the displayed result of the HMM is reported in Fig. 10 as a fingerprint; indeed, posterior probability matrices can be also designated as the fingerprints of the events.

As explained in Section 3, the posterior probability matrix (Pi) does represent the specific chain of events leading to the anomaly, after

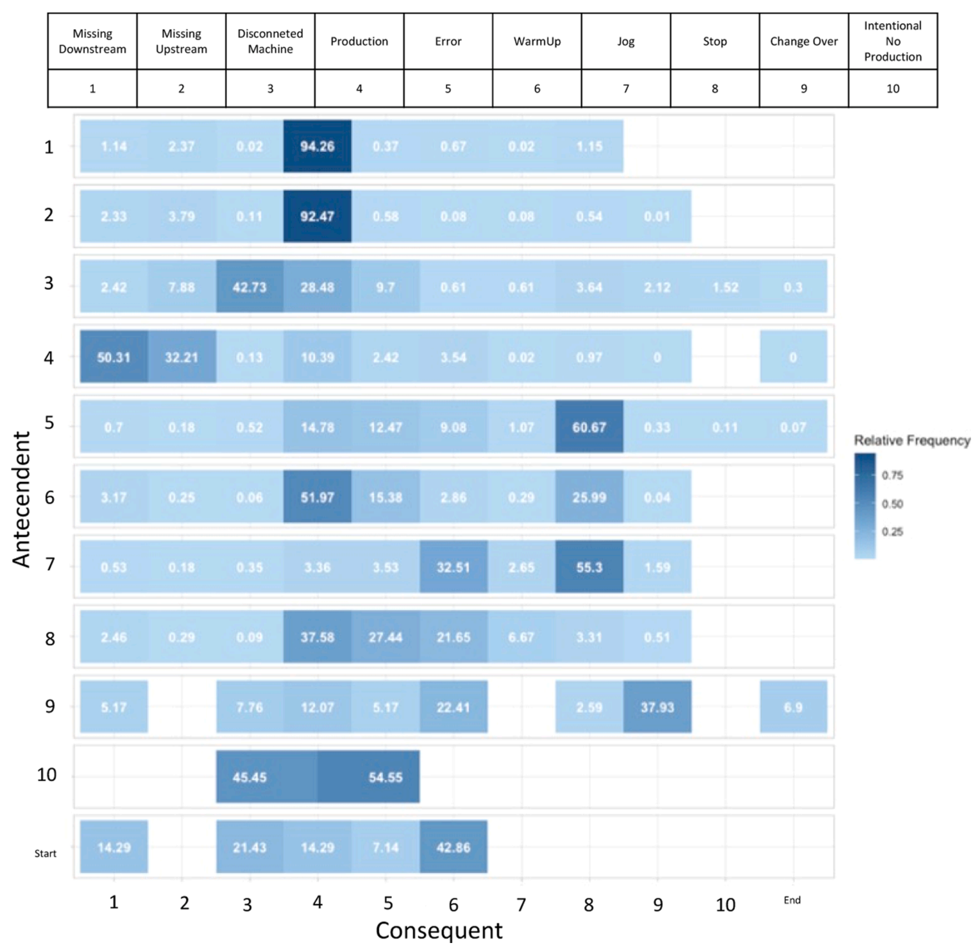

Fig. 9. Transition matrix with 10 machine states for a brushless motor.

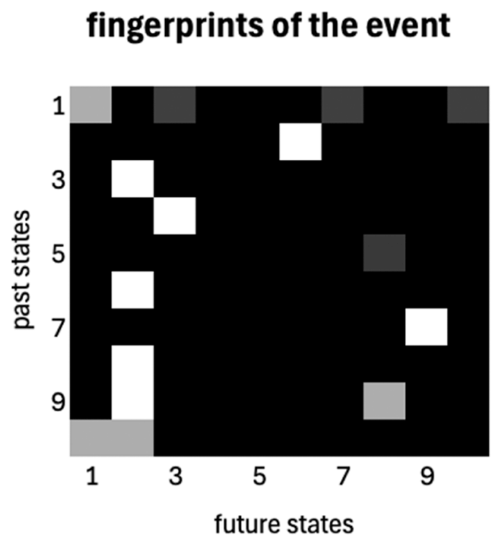

Fig. 10. Posterior probability matrix (P_i) realised after the detection of an anomaly.

removing the baseline or common path for all sections in the production line, which do represent the nominal conduction of the machine.

Once a new posterior probability matrix is generated, the similarity evaluation happens, which compares the previous posterior probability matrices with the newly generated one. Thanks to the previous matching of each posterior probability matrix with the three possible clusters to classify the production period, it is possible to label it as “healthy”, “monitor”, or “critical”, according to its health ratings. The shift from one case to another impacts the maintenance plan for the following week, including production schedule modifications. It has to be noted that, in order to grant the maximum robustness of the solution, the production lines are not continuously operational but run in batches that last 2–3 days. Each batch is followed by a thorough check-up and setup for the next batch. This routine ensures that each production cycle starts under uniform conditions, effectively minimizing the impact of any short-term underlying variations while isolating more significant, longer-term effects. This also leads to a positive effect on the data analytics as it allows to absorb potential weaknesses of the HMM step (namely potential lack of independence of each observation).

The overall picture of the functioning of the data-driven method is presented in Fig. 11, which draws graphically what the pseudo-code in Fig. 7 does continuously.

4.2. Results of the application

A 3-month long period of observations has been mined, so to validate the proposed method: from 2 production lines, 241 records (i.e., each record is a production week) have been collected. The 241 observations have been evaluated using the 2 technological solutions: on the one hand the human-based, experience-based assessment of the maintenance operators and asset managers, on the other the data-driven algorithm described in Section 3.

At the end of the data mining phase, it was available a dataset in which – for each observation – both the human experts and the algorithm reported the assessed health belonging cluster (level #1, #2 or #3): the results obtained by the proposed method in terms of classification in the three groups were eventually compared with the classification made by the operator. .

Over the application period, the results can be expressed through the confusion matrix reported in Fig. 7.

The matrix shows the comparison between *observed* assessment (verified, logged records by humans – y-axis) and *suggested* assessment (generated by the proposed data-driven method – x-axis): for example, in the 168 times that the state of the motors was considered healthy by human observation and verification, 159 times also the algorithm reported the state ‘healthy’ whereas, in 7 occasions, the algorithm stated a to be ‘monitored’ state and in 2 cases reported a critical health state.

Overall, the method’s accuracy is 91 % which was considered satisfactory as the target model’s performance. To calculate accuracy, all values in the main diagonal of the confusion matrix are summed up and then divided by the total number of instances, thus applying the well-known definition of the accuracy indicator. The accuracy provides an overall measure of the model’s performance by considering both true positives and true negatives. Table 2 summarizes the performance in terms of an entire set of well-known indicators (including related formulas), computed from the results reported in the confusion matrix.

It is worth now remarking on the change with respect to an extant solution. The previous method of health rating was based on the operators’ experience and on the assessment involving all the stakeholders of the production line: the output of those multiple interactions among various parties can be considered as the “ground truth” thanks to the deep experience in maintenance activities on every single line. The new, proposed method is based on a data-driven health rating via the proposed method, thus using the human-experience-based assessment as a reference. With an accuracy of 91 %, the method proves that the maintenance assessment can be automatized and be performed continuously using the human experience not as the decision maker but as the supervisor of the automated method and its results.

Despite the cost of the IT infrastructure, the entire algorithm/s development and the creation of a knowledge base, the new, proposed solution manages to be faster and cheaper in comparison to the existing, human-based monitoring: for the existing solution, the higher cost was determined by the human capital and the time dedicated to maintaining the knowledge base so as to keep the assessments as objective and fact-based as possible. Over a longer period, the proposed method should also be compared with the real impacts on the production line/s, this is further considered in the managerial implications.

5. Managerial implications and reasonings over OEM adoption of maintenance service

5.1. Managerial implications of the proposed method

The method has been developed in coherence with the objective of an OEM to offer the proper maintenance services to the customers based on the conduction of their production lines and machines. This relates to a major business driver leading to the prominent need for customized maintenance service offerings based on the current requirements of the extant plants. Accordingly, the data-driven method has the end purpose of activating synergies between the knowledge and information from the OEM and the customer.

Considering the OEM viewpoint, it is remarkable that lines and machines, where this research is contextualized, are built on a core ‘structure’ and some add-ons that can be configured according to customer requirements. This originates from the modularization of the design of the lines and is the common ground that motivates the OEM high-end capability to exploit engineering knowledge about each machine type, to translate into well-known maintenance requirements and, subsequently, maintenance plans to be adopted for a running line/machine.

Looking at the customer viewpoint, different factors in the runtime are influential for the observed behaviour, so the performance, of each line and its machines, and are monitored during the conduction of the line/machines. As postulated in the proposed method, the conduction of the machines should be made accessible by the OEM: a surveillance algorithm, aimed at detecting anomalies, leads to a first level of degradation control. However, the anomaly detection determines just the prerequisite to start mobilizing knowledge and information between OEM and customers. The collaborative maintenance approach is further

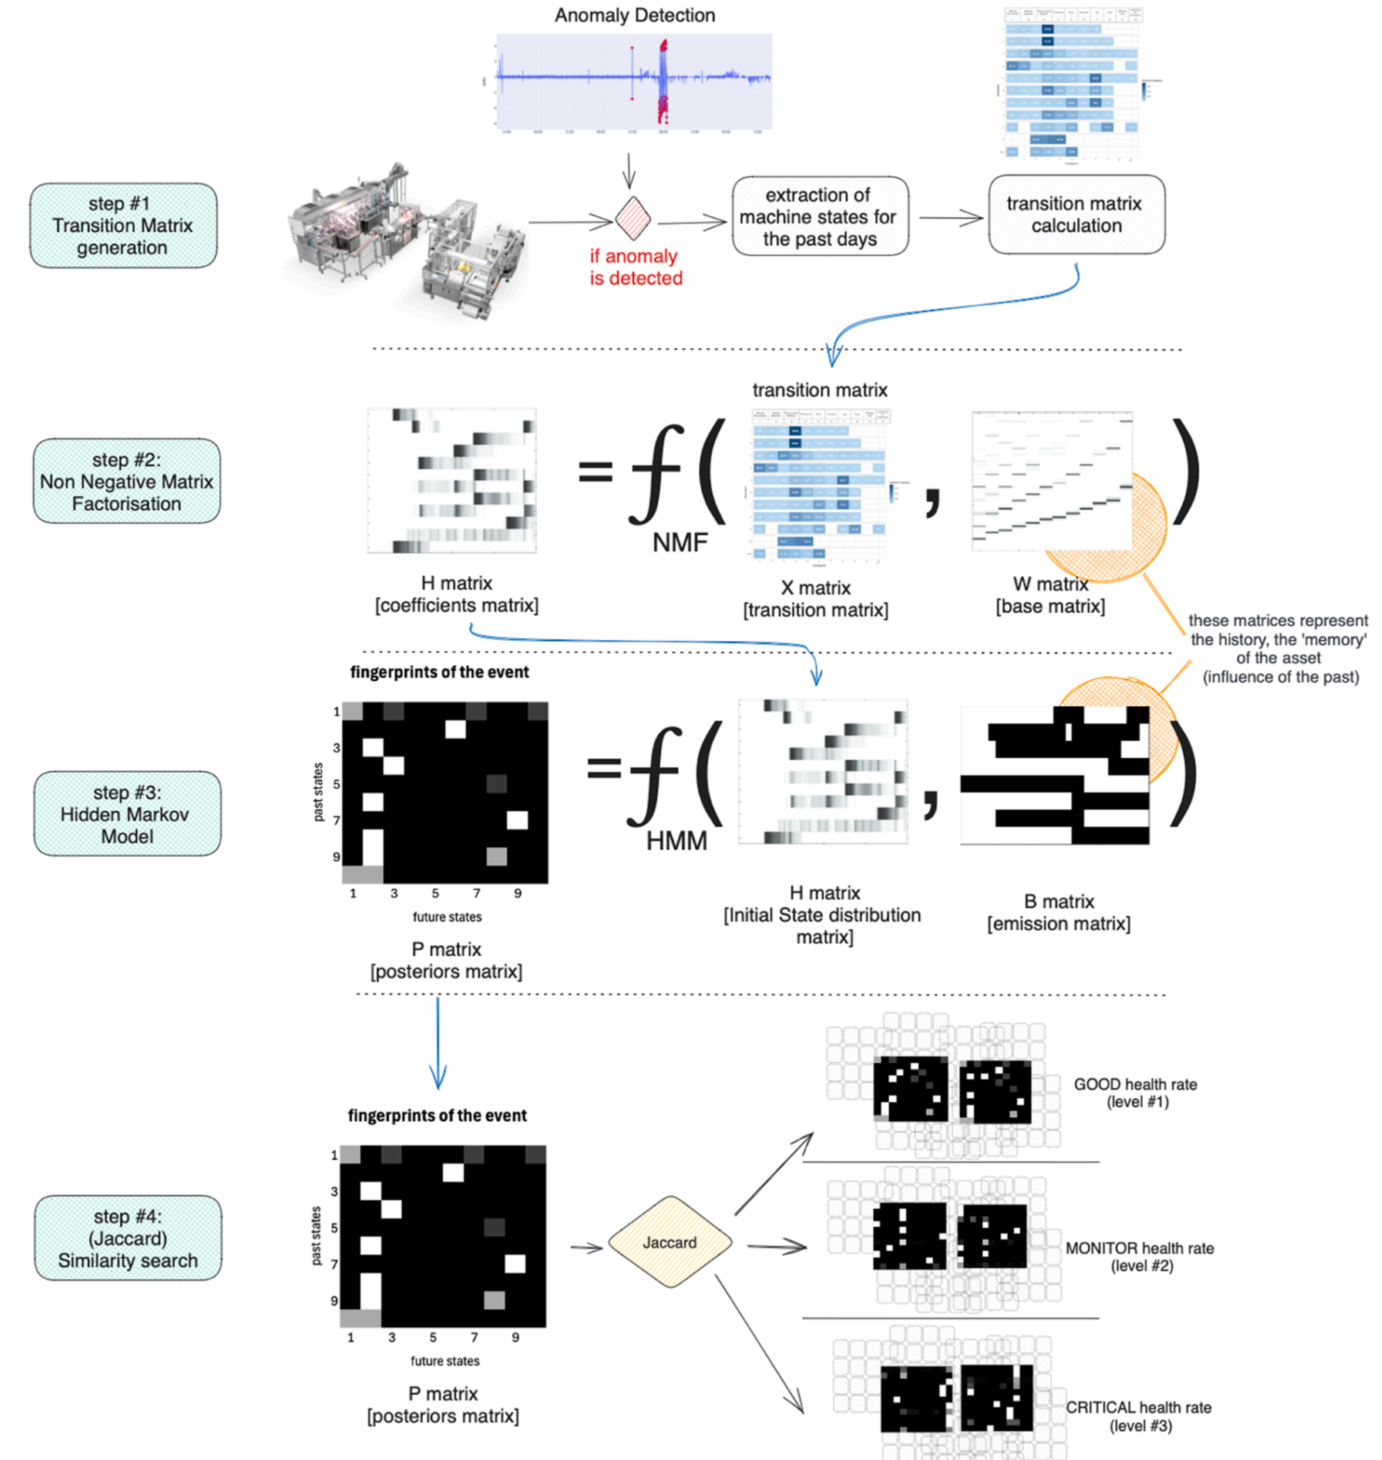

Fig. 11. Instantiation of the data-driven method for maintenance plan adaptation to the industrial case.

developed thanks to the data-driven method proposed in the paper.

In particular, the maintenance strategy, designed to proactively monitor and maintain based on the health ratings over time, can be provided leveraging a quite traditional way, building on regular inspections and testing, but also extending to the condition monitoring habilitated by an IT infrastructure and surveillance algorithms, as it was the case of this research. This finally leads to real-time analytics to identify potential issues before they become serious problems as severe impacts from failures. Even if this already brings benefits, it is not deemed sufficient in the light of the complexity of the production lines: CBM programs could not be so cost-effective, as the maintenance working practice is driven majorly by local performance deviations and alarms without getting the overall picture over the behaviour of the machines running within the production line; this may induce a cognitive overload of maintenance and OEM technical operators, to evaluate the required follow-ups from monitoring tasks from each machine, with the hidden risk also to cause over-maintenance and line stoppages, due to the monitored conditions based on local deviations and alarms. This is one of the major problems at the origin of the research, leading to a service offering at a more tactical level of maintenance plan adaptation

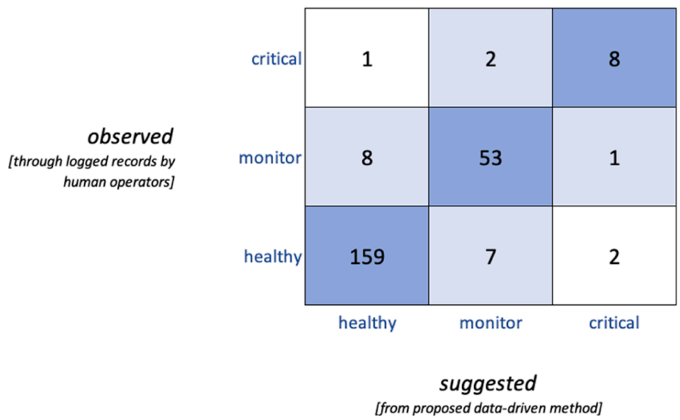

Fig. 12. Results from the testing over three months.

based on machine status data. Therein, the method exploits the traced history of single machines in a production line, providing an automated approach to support a novel supervisory role to recommend the maintenance plan adaptation. On the one hand, the supervisory role works in synergy with the tasks at the operational level, located in each servitised machine, through the anomaly detection algorithm/s; at this level, technical analysis is always possible, to get more insights from the condition monitoring. On the other hand, the supervisory role demonstrates, according to the validation results, to achieve a performance comparable to the experience-based judgements of the maintenance operators and asset managers; beyond that, the method is an automatic workflow, and therefore is scalable, thus enabling to take advantage, in a cost-effective practice, of more lines and machines supervised within the fleet of servitised plants, to fully grasp the benefits of the collaborative prognostics at large scale. The overall observation, collected from the operators in the lines where the method is now running, is that, with the method in place, the cognitive overload would be reduced, leading to idle times to develop more technical insights at the operational level when it is the case. Moreover, as feedback from the asset managers and OEM technical operators, assuming the quality of the data-driven method as evident in the validation, the decision-making becomes standardized, defining clear rules and processes for maintenance plan adaptation, which can be then automated. Eventually, such supervision at the tactical level may support a balanced approach, avoiding the risk of an excess of stoppages due to technical measurements and a hyper-protective approach adverse to any risk. As a cascade effect, this may lead to cost-effective management of the MRO materials to be replaced and related working capital.

The cost-effective scalability of the solution can be better understood in its feasibility considering the integrability of the proposed workflow within the IT infrastructure and the related monitoring activities. Hereafter, the OEM viewpoint is primarily taken. Different reasons are to be considered as a foundation to scale the proposed solution in a large scope of work (more customers, more lines, more machines).

First of all, for the OEM business model, adopting a workflow embedding a machine clustering aimed at similarity evaluation is a relevant objective of the proposed method, and can be justified within a larger perspective of the product-related service offerings including performance and health supervision of industrial assets in general (line, machines, components). In particular, it is worth noticing that, unlike most collaborative prognostics applications, using physical data from pressure and temperature sensors, vibrations and currents from accelerometers, and speeds, the proposed method for data-driven machine clustering is designed to be used on a dataset comprised of only machine status data. This type of data is coherent for the calculation of well-known performance indicators in production lines, such as the OEE. Therefore, introducing the workflow up to the similarity-based clustering in pre-existing software services for the monitoring of machine performances, involves no additional storage costs, no modifications to configured connectivity and automation (i.e., connected PLCs), and no costly and time-consuming modifications to data pipelines. In addition to this, the machine status data is more easily standardized than data from physical variables, as the different machines are often equipped with distinct sensors. The sensor data and their heterogeneity are, according to this proposal, managed at the edge computing level, thus limiting herein the complexity of the design in the overall architecture. Last but not least, it is worth considering that the accumulation of data is typically constrained by data storage capacity and related costs; therefore, it is remarkable that the volume of machine status data is certainly significantly smaller than that of physical variables data, which enables relevant savings to this regard. Nonetheless, the presence of historical data is fundamental for the full model accuracy. Without the historical base, it would be complicated to create critical components of the NMF and the HMM: having a baseline grants the success of the model. In case of missing historical data, it could be possible to use testing data from the machine testing or commissioning but this case is out of the scope of this research.

Combining the scalability herein discussed with the model performance expected in its automatic workflow from the results achieved in this work (as in Section 4), it can be concluded that the data-driven method for maintenance plan adaptation reveals both viable and with a potential to demonstrate its cost-effectiveness over the lifecycle of the machines.

5.2. Reasonings over maintenance services by OEM

In the current industrial market, OEMs are increasingly looking at maintenance service offerings to improve their profitability. Digital technologies are current levers to enable digitally enhanced services. Reaching this objective is not straightforward, which requires, to the authors understanding, to develop future research in order to address valuable scientific problems.

As evidence of this work, the collaborative prognostics approach seems to provide a promising innovative scheme to enrich the as-designed knowledge on machines and systems an OEM has, by adding information thanks to data federation from servitised systems and machines [72]. Nevertheless, lots of challenges are still in place in the undergoing trend of digital servitisation. Keeping those that appear more relevant for this research and its possible follow-ups, the following can be considered remarkable:

- Advanced technologies may extend the portfolio of offered services as well as they can empower those already present. However, a technology-centred service offering does not match customers' needs and is not enough. The effect on the maintenance process must be cleared out in the first place [73,74]; moreover, the business model establishment will be essential to effectively lead to the full exploitation of advanced valuable services [31] such as CBM and PdM run

Table 2
Model performance.

Total number of samples	Accuracy formula: (TP +TN)/ (TP+TN+FP+FN)	Precisionformula: TP / (TP+FP)	Recallformula: TP/ (TP+FN)	F1 scoreformula: 2 *recall*precision/ (recall + precision)
241	91.3 %	94.6 %	92.2 %	93.4 %

Legend: TP = true positive; TN = true negative; FP = false positive; FN = false negative within a collaborative maintenance approach of the OEM with the customers.

- Information is a critical resource for the effective implementation of maintenance servitisation, and the OEMs can leverage upon wide product knowledge to better fit their offering [72], integrating it with proper data pooling from multiple machines/systems of various customers [75].
- Amongst the constraints and requirements, the data privacy and cybersecurity reasons could impact the effectiveness of the proposed maintenance service offerings, especially when including CBM and PdM, as datasets from operating machines may be restricted in their access, being partially or not available at all [56].

Considering these challenges, an OEM must find the proper organizational and technological settings to pursue profitability while ensuring cost-effectiveness through the offered services, to achieve higher gains from the two viewpoints, OEM and customer, with respect to the "do-it-yourselves" on the customer side [76]. The solution, leveraging the collaborative prognostics to identify similar behaviour in servitised machines, appears promising to the end purpose of a cost-effective maintenance servitisation.

6. Conclusions

This research work has proposed a data-driven maintenance plan adaptation method for servitised machines in complex production systems. This is conceived by relying on similarity evaluation through a fleet-wide machine clustering that relates the extant behaviour of individual machines, as observed in the runtime (e.g. a production week), to a benchmark built from past datasets collected from machines of the same type (e.g., over different past weeks) so as to build a health rating system, and correspondingly enable to recommend the required maintenance plan.

In the proposed concept, the method applies a collaborative prognostics approach that leverages events recorded in data logs of machine status to spot similar behaviour. This is a novel way with respect to what is currently put in evidence in collaborative prognostics, typically based on condition monitoring data (vibrations, pressures, temperatures, currents, etc.) only. This novelty is coherent with the concurrent need to deal with automatic maintenance planning, which requires a supervisory role at a tactical level. Maintenance plan adaptation is also a key issue within CBM research in general and is one of the focus traced in the research streams within maintenance servitisation, also in light of the enabler of digitalization for maintenance services. Thus, this work is meant to contribute both to the collaborative prognostics and the digital servitisation streams.

Overall, it is worth remarking that the major novelty provided by the present research relates to the automatic maintenance planning capability based on health rating developed by proper data pooling from multiple machines/systems in different production systems. As this capability improves the operational level of CBM with adaptation at tactical planning, it fosters, through automated workflow, the adaptiveness to the ever-evolving behaviour of the machines present in a complex production system. Over the long term, this may open the possibility of observing changing maintenance requirements impacted by ageing as well as newly introduced production requirements.

Limits of the work majorly regard the scope of experiments. The number of servitised lines/machines where the method is run will be extended, and this will allow a wider range of empirical proofs. Moreover, the utilization over a longer time of the method will enable to confirm not only the model performance but also to measure the resulting contribution to the Key Performance Indicators (KPI) targeted by the customers in their lines, such as, e.g., the OEE or the level of working capital due to the managed MRO. This will finally enable to get proof also of the impacts at the business level.

Other limits regard the verification of the optimality of maintenance plans, to fit the current conditions. Nowadays, the plans, in their content in terms of maintenance tasks, are taken as given, according to the long-term experiences and knowledge of the OEM on the systems and lines is building. Nevertheless, it would be required to investigate if the optimality is close or far away when considering the current conditions of every single machine within a production system. This is also a research opportunity and should require extending the data-driven approach in order to further exploit the shreds of evidence that may result from other machine learning as well as optimization methods that may better fit in some parts of the proposed data-driven method.

Eventually, it is important to remember that the proposed solution considers the anomaly detection algorithm/s and related IT architecture as given for granted. Therefore, an additional effort has to be spent for a global solution on which the OEM could rely for improved maintenance service offerings, combining different options of services offered both at the operational and tactical level, namely as digital services local to each machine, and other digital services globally available for the fleet.

In future research, the collaborative prognostics architecture could be extended. In the current understanding of the authors, two main objectives can be suggested: i) to develop prognostics capability at the machine level, by exploiting at the edge (so local to the machines) the features and machine learning models that could be transferred from machines in other domain (i.e. other customer's sites); this should be done while respecting the data privacy constraint and also the more technical need of dependability of the architecture in terms of managed information; this motivates the need to develop local algorithms at each machine, fed by some transfer learning outcome from federated machines in the fleet; ii) to develop a supervisory approach to make optimal health management; this may extend the current achievements by means of an optimization methodology aimed to support in searching for the optimum of key characteristics in multi-objective functions, representative of both cost minimization and risk mitigation objectives. Last but not least, a study on the business impact of the adoption of collaborative prognostics in the provision of maintenance services will be an interesting extension of the current research experiences. This is clearly aimed at developing a different kind of study. The main interest would be to get evidence of the benefits, barriers and opportunities that may result from collaborative models in maintenance servitisation, with a special emphasis on the exploitation of the collaborative prognostics framework.

CRediT authorship contribution statement

Adalberto Polenghi: Writing – review & editing, Writing – original draft, Methodology, Conceptualization. **Marco Macchi:** Writing – review & editing, Writing – original draft, Validation, Methodology, Conceptualization. **Alessandro Ruberti:** Writing – review & editing, Writing – original draft, Validation, Investigation, Formal analysis, Data curation, Conceptualization.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

[1] Schmenner RW. Manufacturing, service, and their integration: some history and theory. Int J Oper Prod Manag 2009;29:431–43. https://doi.org/10.1108/01443570910953577.

[2] Lightfoot H, Baines T, Smart P. The servitization of manufacturing. Int J Oper Prod Manag 2013;33:1408–34. https://doi.org/10.1108/IJOPM-07-2010-0196.

[3] Parida V, Sjödin DR, Wincent J, Kohtamäki M. Mastering the transition to product-service provision: insights into business models, learning activities, and capabilities. Res-Technol Manag 2014;57:44–52. https://doi.org/10.5437/08956308×5703227.

[4] Ardolino M, Rapaccini M, Saccani N, Gaiardelli P, Crespi G, Ruggeri C. The role of digital technologies for the service transformation of industrial companies. Int J Prod Res 2018;56:2116–32. https://doi.org/10.1080/00207543.2017.1324224.
[5] Pirola F, Boucher X, Wiesner S, Pezzotta G. Digital technologies in product-service systems: a literature review and a research agenda. Comput Ind 2020;123:103301. https://doi.org/10.1016/j.compind.2020.103301.
[6] Tukker A. Product services for a resource-efficient and circular economy – a review. J Clean Prod 2015;97:76–91. https://doi.org/10.1016/j.jclepro.2013.11.049.
[7] Ding K, Jiang P, Zheng M. Environmental and economic sustainability-aware resource service scheduling for industrial product service systems. J Intell Manuf 2017;28:1303–16. https://doi.org/10.1007/s10845-015-1051-7.
[8] Parida V, Wincent J. Why and how to compete through sustainability: a review and outline of trends influencing firm and network-level transformation. Int Entrep Manag J 2019;15:1–19. https://doi.org/10.1007/s11365-019-00558-9.
[9] Paschou T, Rapaccini M, Adrodegari F, Saccani N. Digital servitization in manufacturing: a systematic literature review and research agenda. Ind Mark Manag 2020;89:278–92. https://doi.org/10.1016/j.indmarman.2020.02.012.
[10] Grubic T. Remote monitoring technology and servitization: exploring the relationship. Comput Ind 2018;100:148–58. https://doi.org/10.1016/j.compind.2018.05.002.
[11] Ziaee Bigdeli A, Baines T, Schroeder A, Brown S, Musson E, Guang Shi V, et al. Measuring servitization progress and outcome: the case of 'advanced services. Prod Plan Control 2018;29:315–32. https://doi.org/10.1080/09537287.2018.1429029.
[12] Georgakopoulos D, Jayaraman PP. Internet of things: from internet scale sensing to smart services. Computing 2016;98:1041–58. https://doi.org/10.1007/s00607-016-0510-0.
[13] Höller J, Tsiatsis V, Mulligan C. Toward a machine intelligence layer for diverse industrial IoT use cases. IEEE Intell Syst 2017;32:64–71. https://doi.org/10.1109/MIS.2017.3121543.
[14] Yang X, Moore P, Chong SK. Intelligent products: from lifecycle data acquisition to enabling product-related services. Comput Ind 2009;60:184–94. https://doi.org/10.1016/j.compind.2008.12.009.
[15] Bustinza OF, Vendrell-Herrero F, Baines T. Service implementation in manufacturing: an organisational transformation perspective. Int J Prod Econ 2017;192:1–8. https://doi.org/10.1016/j.ijpe.2017.08.017.
[16] Mont OK. Clarifying the concept of product–service system. J Clean Prod 2002;10:237–45. https://doi.org/10.1016/S0959-6526(01)00039-7.
[17] Esmaeilian B, Behdad S, Wang B. The evolution and future of manufacturing: a review. J Manuf Syst 2016;39:79–100. https://doi.org/10.1016/j.jmsy.2016.03.001.
[18] Gao J, Yao Y, Zhu VCY, Sun L, Lin L. Service-oriented manufacturing: a new product pattern and manufacturing paradigm. J Intell Manuf 2011;22:435–46. https://doi.org/10.1007/s10845-009-0301-y.
[19] Li AQ, Kumar M, Claes B, Found P. The state-of-the-art of the theory on Product-Service Systems. Int J Prod Econ 2020;222:107491. https://doi.org/10.1016/j.ijpe.2019.09.012.
[20] Ren S, Shi L, Liu Y, Cai W, Zhang Y. A personalised operation and maintenance approach for complex products based on equipment portrait of product-service system. Robot Comput-Integr Manuf 2023;80:102485. https://doi.org/10.1016/j.rcim.2022.102485.
[21] Sala R, Bertoni M, Pirola F, Pezzotta G. Data-based decision-making in maintenance service delivery: the D3M framework. J Manuf Technol Manag 2021;32:122–41. https://doi.org/10.1108/JMTM-08-2020-0301.
[22] Chang F, Zhou G, Cheng W, Zhang C, Tian C. A service-oriented multi-player maintenance grouping strategy for complex multi-component system based on game theory. Adv Eng Inform 2019;42:100970. https://doi.org/10.1016/j.aei.2019.100970.
[23] Neely A. Exploring the financial consequences of the servitization of manufacturing. Oper Manag Res 2008;1:103–18. https://doi.org/10.1007/s12063-009-0015-5.
[24] Serradilla O, Zugasti E, Ramirez de Okariz J, Rodriguez J, Zurutuza U. Methodology for data-driven predictive maintenance models design, development and implementation on manufacturing guided by domain knowledge. Int J Comput Integr Manuf 2022;35:1310–34. https://doi.org/10.1080/0951192X.2022.2043562.
[25] Jia Pandhare V, Lee X. Collaborative prognostics for machine fleets using a novel federated baseline learner. Annu Conf PHM Soc 2021;13. https://doi.org/10.36001/phmconf.2021.v13i1.2989.
[26] Upasani K, Bakshi M, Pandhare V, Lad BK. Distributed maintenance planning in manufacturing industries. Comput Ind Eng 2017;108:1–14. https://doi.org/10.1016/j.cie.2017.03.027.
[27] Han X, Wang Z, Xie M, He Y, Li Y, Wang W. Remaining useful life prediction and predictive maintenance strategies for multi-state manufacturing systems considering functional dependence. Reliab Eng Syst Saf 2021;210:107560. https://doi.org/10.1016/j.ress.2021.107560.
[28] Palau AS, Dhada MH, Parlikad AK. Multi-agent system architectures for collaborative prognostics. J Intell Manuf 2019;30:2999–3013. https://doi.org/10.1007/s10845-019-01478-9.
[29] Li H, Parlikad AK. Social internet of industrial things for industrial and manufacturing assets. IFAC-Pap 2016;49:208–13. https://doi.org/10.1016/j.ifacol.2016.11.036.
[30] Balbi M, Cattaneo L, Nucera DD, Macchi M. On the relevance of clustering strategies for collaborative prognostics. IFAC-Pap 2021;54:37–42. https://doi.org/10.1016/j.ifacol.2021.08.004.
[31] González Chávez CA, Unamuno G, Despeisse M, Johansson B, Romero D, Stahre J. Analyzing the risks of digital servitization in the machine tool industry. Robot Comput-Integr Manuf 2023;82:102520. https://doi.org/10.1016/j.rcim.2022.102520.
[32] Wang W. A model for maintenance service contract design, negotiation and optimization. Eur J Oper Res 2010;201:239–46. https://doi.org/10.1016/j.ejor.2009.02.018.
[33] Si G, Xia T, Gebraeel N, Wang D, Pan E, Xi L. A reliability-and-cost-based framework to optimize maintenance planning and diverse-skilled technician routing for geographically distributed systems. Reliab Eng Syst Saf 2022;226:108652. https://doi.org/10.1016/j.ress.2022.108652.
[34] Lee SG, Ma Y-S, Thimm GL, Verstraeten J. Product lifecycle management in aviation maintenance, repair and overhaul. Comput Ind 2008;59:296–303. https://doi.org/10.1016/j.compind.2007.06.022.
[35] Von Grumbkow M. Successful asset management in the paper industry from an OEM point of view: Second to the owner, the OEM has the strongest interest in maintaining and improving the profitability of existing assets. Paper360 2013;8:16–21.
[36] Stip J, Van Houtum G-J. On a method to improve your service BOMs within spare parts management. Int J Prod Econ 2020;221:107466.
[37] Darghouth M, Aït-Kadi D, Chelbi A. Joint optimization of design, warranty and price for products sold with maintenance service contracts. Reliab Eng Syst Saf 2017;165:197–208.
[38] Hezarkhani B, Nagarajan M, Tong C. Toward servitization: optimal design of uptime-guarantee maintenance contracts. Prod Oper Manag 2022;31:3806–22.
[39] Zhu Y, Xia T, Chen Z, Pan E, Xi L. Optimal maintenance service strategy for OEM entering competitive MRO market under opposite patterns. Reliab Eng Syst Saf 2022;217:108060.
[40] Greenough RM, Grubic T. Modelling condition-based maintenance to deliver a service to machine tool users. Int J Adv Manuf Technol 2011;52:1117–32. https://doi.org/10.1007/s00170-010-2760-x.
[41] Chang F, Zhou G, Zhang C, Xiao Z, Wang C. A service-oriented dynamic multi-level maintenance grouping strategy based on prediction information of multi-component systems. J Manuf Syst 2019;53:49–61.
[42] Tsyokhla I, Griffo A, Wang J. Online condition monitoring for diagnosis and prognosis of insulation degradation of inverter-fed machines. IEEE Trans Ind Electron 2019;66:8126–35. https://doi.org/10.1109/TIE.2018.2885740.
[43] Poppe J, Boute RN, Lambrecht MR. A hybrid condition-based maintenance policy for continuously monitored components with two degradation thresholds. Eur J Oper Res 2018;268:515–32. https://doi.org/10.1016/j.ejor.2018.01.039.
[44] Si G, Xia T, Zhu Y, Du S, Xi L. Triple-level opportunistic maintenance policy for leasehold service network of multi-location production lines. Reliab Eng Syst Saf 2019;190:106519. https://doi.org/10.1016/j.ress.2019.106519.
[45] Xia T, Si G, Wang D, Pan E, Xi L. Progressive opportunistic maintenance policies for service-outsourcing network with prognostic updating and dynamical optimization. IEEE Trans Reliab 2022;71:1340–54. https://doi.org/10.1109/TR.2021.3074506.
[46] Teixeira HN, Lopes I, Braga AC. Condition-based maintenance implementation: a literature review. Procedia Manuf 2020;51:228–35. https://doi.org/10.1016/j.promfg.2020.10.033.
[47] Lei Y, Li N, Guo L, Li N, Yan T, Lin J. Machinery health prognostics: a systematic review from data acquisition to RUL prediction. Mech Syst Signal Process 2018;104:799–834. https://doi.org/10.1016/j.ymssp.2017.11.016.
[48] Xia T, Dong Y, Xiao L, Du S, Pan E, Xi L. Recent advances in prognostics and health management for advanced manufacturing paradigms. Reliab Eng Syst Saf 2018;178:255–68. https://doi.org/10.1016/j.ress.2018.06.021.
[49] Cattaneo L, Polenghi A, Macchi M. A framework to integrate novelty detection and remaining useful life prediction in Industry 4.0-based manufacturing systems. Int J Comput Integr Manuf 2021;35:388–408. https://doi.org/10.1080/0951192X.2021.1885062.
[50] Divya D, Marath B, Santosh Kumar MB. Review of fault detection techniques for predictive maintenance. J Qual Maint Eng 2022;29:420–41. https://doi.org/10.1108/JQME-10-2020-0107.
[51] Tahan M, Tsoutsanis E, Muhammad M, Abdul Karim ZA. Performance-based health monitoring, diagnostics and prognostics for condition-based maintenance of gas turbines: A review. Appl Energy 2017;198:122–44. https://doi.org/10.1016/j.apenergy.2017.04.048.
[52] Liu J, Wang Q, Wei B, Li Z. A fault early warning and health status rating method for ensuring safe operation of rotating equipment. 6th Int Conf Inf Sci Control Eng (ICISCE) 2019;2019:635–43. https://doi.org/10.1109/ICISCE48695.2019.00132.
[53] Li Y, Peng S, Li Y, Jiang W. A review of condition-based maintenance: its prognostic and operational aspects. Front Eng Manag 2020;7:323–34. https://doi.org/10.1007/s42524-020-0121-5.
[54] Chen C, Fu H, Zheng Y, Tao F, Liu Y. The advance of digital twin for predictive maintenance: The role and function of machine learning. J Manuf Syst 2023;71:581–94. https://doi.org/10.1016/j.jmsy.2023.10.010.
[55] Akkermans H, Basten R, Zhu Q, Van Wassenhove L. Transition paths for condition-based maintenance-driven smart services (n/a) J Oper Manag 2024. https://doi.org/10.1002/joom.1295.
[56] Zhang K, Xia T, Wang D, Chen G, Pan E, Xi L. Privacy-preserving and sensor-fused framework for prognostic & health management in leased manufacturing system. Mech Syst Signal Process 2023;184:109666. https://doi.org/10.1016/j.ymssp.2022.109666.
[57] Xiang Y, Zhu Z, Coit DW, Feng Q. Condition-based maintenance under performance-based contracting. Comput Ind Eng 2017;111:391–402. https://doi.org/10.1016/j.cie.2017.07.035.

[58] Wang Y, Liu Y, Chen J, Li X. Reliability and condition-based maintenance modeling for systems operating under performance-based contracting. Comput Ind Eng 2020; 142:106344. https://doi.org/10.1016/j.cie.2020.106344.
[59] Wang Y, Gao W, Li X, Liu Y. Joint optimization of performance-based contracting, condition-based maintenance and spare parts inventory for degrading production systems. Reliab Eng Syst Saf 2024;243:109845. https://doi.org/10.1016/j.ress.2023.109845.
[60] Chandola V., Deepthi C., Vipin K., Detecting anomalies in a time series database. Retrieved from the University Digital Conservancy 2009; https://hdl.handle.net/11299/215791.
[61] BS EN 415–11. Safety of packaging machines - Part 11: Determination of efficiency and availability. BSI Standards Publication 2021.
[62] Ghahramani Z. An introduction to hidden markov models and bayesian networks. Int J Pattern Recogn Artif Intell 2001;15:9–42. https://doi.org/10.1142/S0218001401000836.
[63] Lee D, Seung HS. In: Leen T, Dietterich T, Tresp V, editors. Algorithms for Non-negative Matrix Factorization. Advances in Neural Information Processing Systems, vol. 13. MIT Press; 2000.
[64] Gillis N, Glineur F. Accelerated Multiplicative Updates and Hierarchical ALS Algorithms for Nonnegative Matrix Factorization. Neural Comput 2012;24:1085–105. https://doi.org/10.1162/NECO_a_00256.
[65] Bishop CM. Pattern recognition and machine learning. 1st ed..,. New York: Springer; 2006.
[66] Rabiner LR. A tutorial on hidden Markov models and selected applications in speech recognition. Proc IEEE 1989;77:257–86. https://doi.org/10.1109/5.18626.
[67] Guo C, Liang Z. A predictive Markov decision process for optimizing inspection and maintenance strategies of partially observable multi-state systems. Reliab Eng Syst Saf 2022;226:108683. https://doi.org/10.1016/j.ress.2022.108683.
[68] Abdelhadi A, Alwan LC, Yue X. Managing storeroom operations using cluster-based preventative maintenance. J Qual Maint Eng 2015;21:154–70. https://doi.org/10.1108/JQME-10-2013-0066.
[69] Chao A, Chazdon RL, Colwell RK, Shen T-J. A new statistical approach for assessing similarity of species composition with incidence and abundance data. Ecol Lett 2005;8:148–59. https://doi.org/10.1111/j.1461-0248.2004.00707.x.
[70] Liu FT, Ting KM, Zhou Z-H. Isolation-based anomaly detection. 1-3:39 ACM Trans Knowl Discov Data 2012;6(3). https://doi.org/10.1145/2133360.2133363.
[71] Zhong S, Fu S, Lin L, Fu X, Cui Z, Wang R. A novel unsupervised anomaly detection for gas turbine using Isolation Forest. IEEE Int Conf Progn Health Manag (ICPHM) 2019;2019:1–6. https://doi.org/10.1109/ICPHM.2019.8819409.
[72] Lehtonen O, Ala-Risku T, Holmström J. Enhancing field-service delivery: the role of information. J Qual Maint Eng 2012;18:125–40. https://doi.org/10.1108/13552511211244175.
[73] Paluch S. Customer expectations of remote maintenance services in the medical equipment industry. J Serv Manag 2014;25:639–53. https://doi.org/10.1108/JOSM-07-2013-0195.
[74] Kamp B, Ochoa A, Diaz J. Smart servitization within the context of industrial user–supplier relationships: contingencies according to a machine tool manufacturer. Int J Inter Des Manuf 2017;11:651–63. https://doi.org/10.1007/s12008-016-0345-0.
[75] Dourado A, Viana FAC. Early life failures and services of industrial asset fleets. Reliab Eng Syst Saf 2021;205:107225. https://doi.org/10.1016/j.ress.2020.107225.
[76] Colen PJ, Lambrecht MR. Product service systems: exploring operational practices. Serv Ind J 2013;33:501–15. https://doi.org/10.1080/02642069.2011.614344.