Machine Learning for Predictive and Prescriptive Analytics of Operational Data in Smart Manufacturing

Katerina Lepenioti¹, Minas Pertselakis², Alexandros Bousdekis^{1(✉)},
Andreas Louca², Fenareti Lampathaki², Dimitris Apostolou^{1,3},
Gregoris Mentzas¹, and Stathis Anastasiou⁴

¹ Information Management Unit (IMU), Institute of Communication and Computer Systems (ICCS), National Technical University of Athens (NTUA), Athens, Greece

{klepenioti, albous, gmentzas}@mail.ntua.gr

² Suite5 Data Intelligence Solutions Limited, Limassol, Cyprus
{minas, andreas, fenareti}@suite5.eu

³ Department of Informatics, University of Piraeus, Piraeus, Greece
dapost@unipi.gr

⁴ M. J. Maillis S.A., Inofyta, Greece
stathis.anastasiou@maillis.com

Abstract. Perceiving information and extracting insights from data is one of the major challenges in smart manufacturing. Real-time data analytics face several challenges in real-life scenarios, while there is a huge treasure of legacy, enterprise and operational data remaining untouched. The current paper exploits the recent advancements of (deep) machine learning for performing predictive and prescriptive analytics on the basis of enterprise and operational data aiming at supporting the operator on the shopfloor. To do this, it implements algorithms, such as Recurrent Neural Networks for predictive analytics, and Multi-Objective Reinforcement Learning for prescriptive analytics. The proposed approach is demonstrated in a predictive maintenance scenario in steel industry.

Keywords: Recurrent Neural Network · Multi-Objective Reinforcement Learning · Deep learning · Industry 4.0 · Predictive maintenance

1 Introduction

Perceiving information and extracting business insights and knowledge from data is one of the major challenges in smart manufacturing [1]. In this sense, advanced data analytics is a crucial enabler of Industry 4.0 [2]. More specifically, among the major challenges for smart manufacturing are: (deep) machine learning, prescriptive analytics in industrial plants, and analytics-based decision support in manufacturing operations [3]. The wide adoption of IoT devices, sensors and actuators in manufacturing environments has fostered an increasing research interest on real-time data analytics. However, these approaches face several challenges in real-life scenarios: (i) They require a large amount of sensor data that already have experienced events (e.g. failures of -ideally- all possible causes); (ii) They require an enormous computational capacity

that cannot be supported by existing computational infrastructure of factories; (iii) In most cases, the sensor data involve only a few components of a production line, or a small number of parameters related to each component (e.g. temperature, pressure, vibration), making impossible to capture the whole picture of the factory shop floor and the possible correlations among all the machines; (iv) The cold-start problem is rarely investigated. On the other hand, there is a huge treasure of legacy, enterprise and operational systems data remaining untouched. Manufacturers are sitting on a goldmine of unexplored historical, legacy and operational data from their Manufacturing Execution Systems (MES), Enterprise Resource Planning systems (ERP), etc. and they cannot afford to miss out on its unexplored potential. However, only 20–30% of the value from such available data-at-rest is currently accrued [4].

Legacy data contain information regarding the whole factory cycle and store events from all machines, whether they have sensors installed or not (e.g. products per day, interruption times of production line, maintenance logs, causalities, etc.) [5]. Therefore, legacy data analytics have the credentials to move beyond KPIs calculations of business reports (e.g. OEE, uptime, etc.), towards providing an all-around view of manufacturing operations on the shopfloor in a proactive manner. In this direction, the recent advancements of machine learning can have a substantial contribution in performing predictive and prescriptive analytics on the basis of enterprise and operational data aiming at supporting the operator on the shopfloor and at extracting meaningful insights. Combining predictive and prescriptive analytics is essential for smarter decisions in manufacturing [2]. In addition mobile computing (with the use of mobile devices, such as smartphones and tablets) can significantly enable timely, comfortable, non-intrusive and reliable interaction with the operator on the shopfloor [6], e.g. for generating alerts, guiding their work, etc. through dedicated mobile apps.

The current paper proposes an approach for predictive and prescriptive analytics on the basis of enterprise and operational data for smart manufacturing. To do this, it develops algorithms based on Recurrent Neural Networks (RNN) for predictive analytics, and Multi-Objective Reinforcement Learning (MORL) for prescriptive analytics. The rest of the paper is organized as follows: Sect. 2 presents the background, the challenges and prominent methods for predictive and prescriptive analytics of enterprise and operational data for smart manufacturing. Section 3 describes the proposed approach, while Sect. 4 shows a walkthrough scenario of the proposed approach in the steel industry. Section 5 presents the experimental results, while Sect. 6 concludes the paper and outlines the plans for future research.

2 Background, Existing Challenges and Novel Methods

2.1 Predictive Analytics for Smart Manufacturing

Background. Intelligent and automated data analysis which aims to discover useful insights from data has become a best practice for modern factories. It is supported today by many software tools and data warehouses, and it is known by the name “descriptive analytics”. A step further, however, is to use the same data to feed models that can make predictions with similar or better accuracy than a human expert. In the

framework of smart manufacturing, prognostics related to machines' health status is a critical research domain that often leverages machine learning methods and data mining tools. In most of the cases, this is related to the analysis of streaming sensor data mainly for health monitoring [7–9], but also for failure prediction [10–12] as part of a predictive maintenance strategy. However, in all of these approaches, the prediction is produced only minutes or even seconds before the actual failure, which, is not often a realistic and practical solution for a real industrial case. The factory managers need to have this information hours or days before the event, so that there is enough time for them to act proactively and prevent it. One way to achieve this is to perform data mining on maintenance and operational data that capture the daily life-cycle of the shop floor in order to make more high-level predictions [13–15].

Existing Challenges. The most notable challenges related to predictive analytics for smart manufacturing include: (a) Predictions always involve a degree of uncertainty, especially when the data available are not sufficient quantity-wise or quality-wise; (b) Inconsistent, incomplete or missing data with low dimensionality often result into overfitting or underfitting that can lead to misleading conclusions; (c) Properly preparing and manipulating the data in order to conclude to the most appropriate set of features to be used as input to the model is the most time-consuming, yet critical to the accuracy of the algorithms, activity; (d) Lack of a common “language” between data scientists and domain experts hinders the extraction of appropriate hypothesis from the beginning and the correct interpretation and explainability of results.

Novel Methods. *Time series forecasting* involves prediction models that analyze time series data and usually infer future data trends. A time series is a sequence of data points indexed in time order. Unlike regression predictive modeling, time series also adds the complexity of a sequence dependence among the input variables. *Recurrent Neural Networks (RNN)* are considered to be powerful neural networks designed to handle sequence dependence. *Long Short-Term Memory Network (LSTM)* is a type of RNN that is typically used in deep learning for its ability to learn long-term dependencies and handle multiple input and output variables.

2.2 Prescriptive Analytics for Smart Manufacturing

Background. Prescriptive analytics aims at answering the questions “What should I do?” and “Why should I do it?”. It is able to bring business value through adaptive, time-dependent and optimal decisions on the basis of predictions about future events [16]. During the last years, there is an increasing interest on prescriptive analytics for smart manufacturing [17], and is considered to be the next evolutionary step towards increasing data analytics maturity for optimized decision making, ahead of time.

Existing Challenges. The most important challenges of prescriptive analytics include [2, 17, 18]: (i) Addressing the uncertainty introduced by the predictions, the incomplete and noisy data and the subjectivity in human judgement; (ii) Combining the “learned knowledge” of machine learning and data mining methods with the “engineered knowledge” elicited from domain experts; (iii) Developing generic prescriptive analytics methods and algorithms utilizing artificial intelligence and machine learning

instead of problem-specific optimization models; (iv) Incorporating adaptation mechanisms capable of processing data and human feedback to continuously improve decision making process over time and to generate non-intrusive prescriptions; (v) Recommending optimal plans out of a list of alternative (sets of) actions.

Novel Methods. *Reinforcement Learning (RL)* is considered to be a third machine learning paradigm, alongside supervised learning and unsupervised learning [19]. RL shows an increasing trend in research literature as a tool for optimal policies in manufacturing problems (e.g. [20, 21]). In RL, the problem is represented by an environment consisting of states and actions and learning agents with a defined goal state. The agents aim to reach the goal state while maximizing the rewards by selecting actions and moving to different states. In *Interactive RL*, there is the additional capability of incorporating evaluative feedback by a human observer so that the RL agent learns from both human feedback and environmental reward [22]. Another extension is *Multi-Objective RL (MORL)*, which is a sequential decision making problem with multiple objectives. MORL requires a learning agent to obtain action policies that can optimize multiple objectives at the same time [23].

3 The Proposed Approach

The proposed approach consists of a predictive analytics component (Sect. 3.1) and a prescriptive analytics component (Sect. 3.2) that process enterprise and operational data from manufacturing legacy systems, as depicted in Fig. 1. The communication is conducted through an event broker for the event predictions and the actions prescriptions, while other parameters (i.e. objective values and alternative actions) become available through RESTful APIs. The results are communicated to business users and shopfloor operators through intuitive interfaces addressed to both computers and mobile devices.

Fig. 1. The architecture of the proposed approach.

3.1 Recurrent Neural Network for Predictive Analytics

The proposed predictive analytics approach aims to: (i) exploit hidden correlations inside the data that derive from the day-to-day shop floor operations, (ii) create and adjust a predictive model able to identify future machinery failures, and (iii) make estimations regarding the *timing* of the failure, i.e. when a failure of the machinery may occur, given the history of operations on the factory. This type of data usually contains daily characteristics that derive from the production line operations and are typically collected as part of a world-wide best practice for monitoring, evaluation and improvement of the effectiveness of the production process. The basic measurement of this process is an industry standard known as Overall Equipment Effectiveness (OEE) and is computed as: $OEE(\%) = Availability(\%) \times Performance(\%) \times Quality(\%)$. Availability is the ratio of actual operational time versus the planned operational time, Performance is the ratio of actual throughput of products versus the maximum potential throughput, and the Quality is the ratio of the not-rejected items produced vs the total production. The OEE factor can be computed for the whole production line as an indication of the factory's effectiveness or per machine or a group of machines. The proposed methodology takes advantage of these commonly extracted indicators and processes them in two steps: in predictive model building (learning) and predictive model deployment.

Predictive Model Building. The predictive analytics model incorporates LSTM and exploits its unique ability to “remember” a sequence of patterns and its relative insensitivity to possible time gaps in the time series. As in most neural network algorithms, LSTM networks are able to seamlessly model non-linear problems with multiple input variables through the iterative training of their parameters (weights). Since the predictive analytics model deals with time-series, the LSTM model is trained using supervised learning on a set of training sequences assigned to a known output value. Therefore, an analyst feeds the model with a set of daily features for a given machine (e.g. the factors that produce the OEE) and use as outcome the number of days until the next failure. This number is known since historical data can hold this information. Nevertheless, when the model is finally built and put in operation, it will use new input data and will have to estimate the new outcome.

Predictive Model Deployment. When the LSTM model is fed with new data it can produce an estimation of when the next failure will occur (i.e. number of days or hours) and what is the expected interruption duration in the following days. Although this estimation may not be 100% accurate, it could help factory managers to program maintenance actions proactively in a flexible and dynamic manner, compared to an often rigid and outdated schedule that is currently the common practice. This estimation feeds into prescriptive analytics aiming at automating the whole decision-making process and provide optimal plans.

3.2 Multi-Objective Reinforcement Learning for Prescriptive Analytics

The proposed prescriptive analytics approach is able to: (i) recommend (prescribe) both perfect and imperfect actions (e.g. maintenance actions with various degrees of

restoration); (ii) model the decision making process under uncertainty instead of the physical manufacturing process, thus making it applicable to various industries and production processes; and, (iii) incorporate the preference of the domain expert into the decision making process (e.g. according to their skills, experience, etc.), in the form of feedback over the generated prescriptions. To do these, it incorporates Multi-Objective Reinforcement Learning (MORL). Unlike most of the Multi-objective Optimization approaches which result in the Pareto front set of optimal solutions [24], the proposed approach provides a single optimal solution (prescription), thus generating more concrete insights to the user. The proposed prescriptive analytics algorithm consists of three steps: prescriptive model building, prescriptive model solving, and prescriptive model adapting, which are described in detail below.

Prescriptive Model Building. The prescriptive analytics model representing the decision making process is defined by a tuple (S, A, T, R) , where S is the state space, A is the action space, T is the transition function $T : S \times A \times S \rightarrow \mathbb{R}$ and R is the vector reward function $R : S \times A \times S \rightarrow \mathbb{R}^n$ where the n -dimensions are associated with the objectives to be optimized O_n . The proposed prescriptive analytics model has a single starting state S_N , from which the agent starts the episode, and a state S_B that the agent tries to avoid. Each episode of the training process of the RL agent will end, when the agent returns to the normal state S_N or when it reaches S_B . Figure 2 depicts an example including 3 alternative (perfect and/or imperfect maintenance) actions (or sets of actions) $S_{A_i}, i = 1, 2, 3$, each one of which is assigned to a reward vector. The prescriptive analytics model is built dynamically. In this sense, the latest updates on the number of the action states S_{A_i} and the estimations of the objectives' values for each state S_k are retrieved through APIs from the predictive analytics. Each action may be implemented either before the breakdown (in order to eliminate or mitigate its impact) or after the breakdown (if this occurs before the implementation of mitigating actions). After the implementation of each action, the equipment returns to its normal state S_N . Solid lines represent the transitions α_i that have non-zero reward with respect to the optimization objectives and move the agent from one state to another.

Fig. 2. An example of the prescriptive analytics model with 3 alternative (sets of) actions.

Prescriptive Model Deployment. On the basis of event triggers for predicted abnormal situations (e.g. about the time of the next breakdown) received through a

message broker, the model moves from the normal state S_N to the dangerous state S_D . For each objective, the reward functions are defined according to whether the objective is to be maximized or minimized. On this basis, the optimal policy $\pi_{O_i}(s, a)$, for each objective O_i is calculated with the use of the actor-critic algorithm, which is a policy gradient algorithm aiming at searching directly in (some subset of) the policy space starting with a mapping from a finite-dimensional (parameter) space to the space of policies [23]. Assuming independent objectives, the multi-objective optimal policy is derived from: $\pi_{opt}(s, a) = \prod_{i \in I} \pi_{O_i}(s, a)$. The time constraints of the optimal policy (prescription) are defined by the prediction event trigger. The prescription is exposed to the operator on the shop-floor (e.g. through a mobile device) providing them the capability to accept or reject it. If accepted, the prescribed action is added to the actions plan.

Prescriptive Model Adaptation. The prescriptive analytics model is able to adapt according to feedback by the expert over the generated prescriptions. This approach learns from the operator whether the prescribed actions converge with their experience or skills and incorporates their preference to the prescriptive analytics model. In this way, it provides non-disruptive decision augmentation and thus, achieves an optimized human-machine interaction, while, at the same time, optimizing manufacturing KPIs. To do this, it implements the policy shaping algorithm [25], a Bayesian approach that attempts to maximize the information gained from human feedback by utilizing it as direct labels on the policy. For each prescription, optional human feedback is received as a signal of approval or rejection, numerically mapped to the reward signals and interpreted into a step function. The feedback is converted into a policy $\pi_{feedback}(s, a)$, the distribution of which relies on the consistency, expressing the user's knowledge regarding the optimality of the actions, and the likelihood of receiving feedback. Assuming that the feedback policy is independent from the optimal multi-objective policy, the synthetic optimal policy for the optimization objectives and the human feedback is calculated as: $\pi_{opt}(s, a) = \pi_{opt}(s, a) * \pi_{feedback}(s, a)$.

4 A Predictive Maintenance Scenario in the Steel Industry

The case examined is the cold rolling production line of M. J. Maillis S.A. Cold rolling is a process of reduction of the cross-sectional area through the deformation caused by a pair of rotating in opposite directions metal rolls in order to produce rolling products with the closest possible thickness tolerances and an excellent surface finish. In the milling station, there is one pair of back up rolls and one pair of work rolls. The deformation takes place through force of the rolls supported by adjustable strip tension in both coilers and de-coilers. Over the life of the roll some wear will occur due to normal processing, and some wear will occur due to extraneous conditions. During replacement, the rolls are removed for grinding, during which some roll diameter is lost, and then are stored in the warehouse for future use. After several regrinding, the diameter of the roll becomes so small that is no longer operational.

The LSTM model of predictive analytics was created using the Keras library with Tensorflow as backend and the MORL using Brown-UMBC Reinforcement Learning

and Planning (BURLAP) library, while the event communication between them is performed with a Kafka broker. In the M. J. Maillis S.A case, the system predicts the time of the next breakdown and the RUL of the available rolls. For the latter, the operator can select one of the repaired rollers, having been subject to grinding, or a new one. Therefore, the alternative actions are created dynamically according to the available repaired rollers existing in the warehouse. Each one has a different RUL, according to its previous operation, and a different cost (retrieved from enterprise systems) due to its depreciation. Each roller has an ID and is assigned to its characteristics/objectives of MORL (i.e. cost to be minimized and RUL to be maximized) in order to facilitate its traceability. The available rolls along with the aforementioned objectives values are retrieved on the basis of a predicted breakdown event trigger.

Table 1. The alternative actions, their costs and their RULs.

Actions	Cost (Euro)	RUL (days)
Replace with new roller	8,000	12
Replace with repaired roller ID1	5,284	5
Replace with repaired roller ID2	3,510	4
Replace with repaired roller ID3	4,831	6

Fig. 3. Illustration of the scenario under examination.

The alternative actions for the current scenario along with their costs and RULs are shown in Table 1. The action “Replace with new roller” represents a perfect maintenance action, while the rest ones represent imperfect maintenance actions. Figure 3 depicts an example of the process in which the prescription “Replace with repaired roller ID3” is generated on the basis of a breakdown prediction and previously received feedback and instantly communicated to the operators through a dedicated mobile app. The operators are also expected to provide feedback so that their knowledge and preferences are incorporated in the system and the models are adapted accordingly.

5 Experimental Results

5.1 Evaluation of the Predictions About the Timing of Interruptions

The legacy datasets used are related to the OEE of M. J. Maillis S.A. factory for the years 2017, 2018 and 2019. The datasets included a total of 21 features and some of the most useful were the real operational time, the time of interruptions and the duration of the breakdown events. A multivariate time series model was designed, as several input features were used in order to predict one output feature. To this direction, the first time series analysis performed in order to predict when the next interruption will occur (*'When do we expect the next interruption?'*). A new feature named 'Days from next breakdown event' was created, and the model was trained. The input features selected for this prediction were: Availability, Performance, Minutes of Breakdown, Minutes of Interruptions, Real Gross Production and the date. After preprocessing the data, the appropriate sequences were extracted that would allow for several timesteps to be selected and tested for the analysis. In our case, timestep represented the last "n" days that the model will use. The LSTM model was then created by testing several layers, optimizers, neurons, batch sizes and epochs until the best performing model was designed. The final result ended being a sequential model with a first LSTM layer of 32 neurons, a second LSTM layer of 16 neurons, a dropout layer with rate 0.1, and finally a dense layer. The model was then trained using data from 2017, 2018 and six months of 2019; with an rmsprop optimizer, a batch size of 1, a timestep of 1, an epochs value of 300, and an early stopping that reached the best performance around 100 epochs. Predictions deal with the last six months of 2019, and the result can be seen in Fig. 4. The blue line represents the actual values and the orange line represents the predicted values. The RMSE came to be 1.26, meaning that there is an average of 1.26 days uncertainty in every result predicted.

Fig. 4. LSTM result - prediction of when the next interruption will occur. (Color figure online)

5.2 Evaluation of the Predictions About the Duration of Interruptions

The second analysis aimed to predict the expected interruption duration for the following day (*‘Which is the expected interruption duration for the following day?’*). The input features used in this LSTM model were: Availability, Performance, Minutes of breakdown, Real Gross Production, Number of breakdowns, and month (date). Again, several LSTM parameters and layers were tested and the final model resulted to be a sequential model with a first LSTM layer of 24 neurons and an activation function ‘relu’, a second layer of 12 neurons with a ‘relu’ activation function, a dropout layer of 0.1 rate, and finally a dense layer. The model was trained using data from 2017 and 2018; using a batch size of 20, 100 epochs, a timestep of 3 and an rmsprop optimizer. Predictions were performed in 2019 data and results are depicted in Fig. 5. The blue line represents the actual value whereas the orange line represents the predicted value. The overall RMSE is 107.57, meaning that there is an average of 107.57 min uncertainty in each prediction.

Fig. 5. LSTM result - prediction of the expected interruption duration for the following day. (Color figure online)

5.3 Evaluation of Feedback Impact on Prescriptions

For this experiment, the actor-critic algorithm, which calculates the associated optimal policy sequentially within 10000 episodes, consists of a Boltzmann actor and a TD-Lambda critic with learning rate = 0.3, lambda = 0.4 and gamma = 0.99. The generated policies are then integrated into a single policy taking into account the consistency (C = 0.7) and likelihood (L = 0.8) values. Table 2 presents five “snapshots” of positive and negative feedback along with the resulting shaped prescriptions and their respective policies. Each “snapshot” is compared to the previous one.

Table 2. Evaluation of feedback impact on prescriptive analytics.

#	Initial prescription	Initial policy	Positive feedback	Negative feedback	Shaped prescription	Shaped policy
1	ID2	0.998900	5	7	ID2	0.155000
2	ID2	0.155000	5	12	ID2	0.000410
3	ID2	0.000410	5	18	ID3	0.000075
4	ID3	0.000075	7	14	ID1	0.000022
5	ID1	0.000022	0	4	New roller	0.000004

6 Conclusion and Future Work

In this paper, we proposed an approach for predictive and prescriptive analytics aiming at exploiting the huge treasure of legacy enterprise and operational data and to overcome some challenges of real-time data analytics. The potential of the proposed approach is high, especially in traditional industries that have not benefit from the advancements of Industry 4.0 and that have just started investigating the potential of data analytics and machine learning for the optimization of their production processes. The traditional manufacturing sectors (e.g. textile, furniture, packaging, steel processing) have usually older factories with limited capacity on investing in modern production technologies. Since the neural networks are inherently adaptive, the proposed approach could be applied to similar production lines (e.g. at a newly established factory of the same type) overcoming the cold-start problem, due to which other techniques usually fail. It also exploits both the “voice of data” and the “voice of experts”. Regarding future work, we plan to evaluate our proposed approach in additional use cases, with different requirements, as well as to investigate approaches and algorithms for fusion of the outcomes derived from real-time data analytics and operational data analytics that represent different levels of information.

Acknowledgments. This work is funded by the European Commission project H2020 UPTIME “Unified Predictive Maintenance System” (768634).

References

1. Gröger, C.: Building an industry 4.0 analytics platform. *Datenbank Spektrum* **18**, 5–14 (2018). <https://doi.org/10.1007/s13222-018-0273-1>
2. Menezes, B., Kelly, J., Leal, A., Le Roux, G.: Predictive, prescriptive and detective analytics for smart manufacturing in the information age. *IFAC-PapersOnLine* **52**, 568–573 (2019)
3. Big Data Challenges in Smart Manufacturing: A discussion paper for BDVA and EFFRA Research & Innovation roadmap alignment BDVA. <http://www.bdva.eu/node/1002>
4. The age of analytics: competing in a data-driven world. <https://www.mckinsey.com/business-functions/mckinsey-analytics/our-insights/the-age-of-analytics-competing-in-a-data-driven-world>
5. Pertselakis, M., Lampathaki, F., Petralli, P.: Predictive maintenance in a digital factory shop-floor: data mining on historical and operational data coming from manufacturers’ information systems. In: Proper, H., Stirna, J. (eds.) *Advanced Information Systems*

- Engineering Workshops. LNBIP, vol. 349, pp. 120–131. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-20948-3_11
6. Yang, H., Kumara, S., Bukkapatnam, S., Tsung, F.: The internet of things for smart manufacturing: a review. *IIE Trans.* **51**, 1190–1216 (2019)
 7. Lee, J., Lapira, E., Bagheri, B., Kao, H.: Recent advances and trends in predictive manufacturing systems in big data environment. *Manuf. Lett.* **1**, 38–41 (2013)
 8. Kobbacy, K., Fawzi, B., Percy, D., Ascher, H.: A full history proportional hazards model for preventive maintenance scheduling. *Qual. Reliab. Eng. Int.* **13**, 187–198 (1997)
 9. Lin, C., Tseng, H.: A neural network application for reliability modelling and condition-based predictive maintenance. *Int. J. Adv. Manuf. Technol.* **25**, 174–179 (2004)
 10. Choudhary, A., Harding, J., Tiwari, M.: Data mining in manufacturing: a review based on the kind of knowledge. *J. Intell. Manuf.* **20**, 501–521 (2008)
 11. Harding, J., Shahbaz, M., Kusiak, A.S.: Data mining in manufacturing: a review. *J. Manuf. Sci. Eng.* **128**, 969–976 (2005)
 12. Bey-Temsamani, A., Engels, M., Motten, A., Vandenplas, S., Ompusunggu, A.P.: A practical approach to combine data mining and prognostics for improved predictive maintenance. In: *The Data Mining Case Studies Workshop (DMCS), 15th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD 2009)*, Paris, pp. 37–44 (2009)
 13. Bastos, P., Lopes, I., Pires, L.C.M.: Application of data mining in a maintenance system for failure prediction. In: *Safety, Reliability and Risk Analysis: Beyond the Horizon: 22nd European Safety and Reliability*, vol. 1, pp. 933–940 (2014)
 14. Romanowski, C.J., Nagi, R.: Analyzing maintenance data using data mining methods. In: Braha, D. (ed.) *Data Mining for Design and Manufacturing*. MACO, vol. 3, pp. 235–254. Springer, Boston (2001). https://doi.org/10.1007/978-1-4757-4911-3_10
 15. Susto, G., Schirru, A., Pampuri, S., McLoone, S., Beghi, A.: Machine learning for predictive maintenance: a multiple classifier approach. *IEEE Trans. Ind. Inf.* **11**, 812–820 (2015)
 16. Šikšnys, L., Pedersen, T.B.: Prescriptive analytics. In: Liu, L., Özsu, M. (eds.) *Encyclopedia of Database Systems*. Springer, New York (2016). <https://doi.org/10.1007/978-1-4899-7993-3>
 17. Vater, J., Harscheidt, L., Knoll, A.: Smart manufacturing with prescriptive analytics. In: *2019 8th International Conference on Industrial Technology and Management (ICITM)*, pp. 224–228. IEEE (2019)
 18. Lepenioti, K., Bousdekis, A., Apostolou, D., Mentzas, G.: Prescriptive analytics: literature review and research challenges. *Int. J. Inf. Manag.* **50**, 57–70 (2020)
 19. Sutton, R., Barto, A.G.: *Reinforcement Learning: An Introduction*. MIT Press, Cambridge (2018)
 20. Dornheim, J., Link, N., Gumbsch, P.: Model-free adaptive optimal control of episodic fixed-horizon manufacturing processes using reinforcement learning. *Int. J. Control Autom. Syst.* (2019). <https://doi.org/10.1007/s12555-019-0120-7>
 21. Rocchetta, R., Bellani, L., Compare, M., Zio, E., Patelli, E.: A reinforcement learning framework for optimal operation and maintenance of power grids. *Appl. Energy* **241**, 291–301 (2019)
 22. Li, G., Gomez, R., Nakamura, K., He, B.: Human-centered reinforcement learning: a survey. *IEEE Trans. Hum. Mach. Syst.* **49**, 337–349 (2019)
 23. Liu, C., Xu, X., Hu, D.: Multiobjective reinforcement learning: a comprehensive overview. *IEEE Trans. Syst. Man Cybern. Syst.* **45**, 385–398 (2015)
 24. Tozer, B., Mazzuchi, T., Sarkani, S.: Many-objective stochastic path finding using reinforcement learning. *Expert Syst. Appl.* **72**, 371–382 (2017)
 25. Griffith, S., Subramanian, K., Scholz, J., Isbell, C., Thomaz, A.L.: Policy shaping: integrating human feedback with reinforcement learning. In: *NIPS* (2013)