

Enabling predictive analytics for smart manufacturing through an IIoT platform

Original

Enabling predictive analytics for smart manufacturing through an IIoT platform / Cerquitelli, T; Nikolakis, N; Bethaz, P; Panicucci, S; Ventura, F; Macii, E; Andolina, S; Marguglio, A; Alexopoulos, K; Petrali, P; Pagani, A; van Wilgen, P; Ippolito, M. - 53:(2020), pp. 179-184. (Intervento presentato al convegno IFAC-PapersOnLine) [10.1016/j.ifaco1.2020.11.029].

Availability:

This version is available at: 11583/2924299 since: 2021-09-16T14:34:33Z

Publisher:

ELSEVIER

Published

DOI:10.1016/j.ifaco1.2020.11.029

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

Enabling predictive analytics for smart manufacturing through an IIoT platform ^{*}

T. Cerquitelli ^{*}, N. Nikolakis ^{**}, P. Bethaz ^{*}, S. Panicucci ^{***},
 F. Ventura ^{*}, E. Macii ^{****}, S. Andolina [†], A. Marguglio [‡],
 K. Alexopoulos ^{**}, P. Petrali [§], A. Pagani [§], P. van Wilgen [‡],
 M. Ippolito ^{***}

^{*} *Department of Control and Computer Engineering, Politecnico di Torino, Turin, ITALY, (e-mail: name.surname@polito.it)*

^{**} *Laboratory for Manufacturing Systems & Automation, Department of Mechanical Engineering and Aeronautics, University of Patras, Patras, GREECE, (e-mail: surname@lms.mech.upatras.gr)*

^{***} *COMAU S.p.A., Turin, ITALY, (e-mail: name.surname@comau.com)*

^{****} *Interuniversity Department of Regional and Urban Studies and Planning, Politecnico di Torino, Turin, ITALY, (e-mail: name.surname@polito.it)*

[†] *SynArea Consultants S.r.l., Turin, ITALY, (e-mail: name.surname@synarea.com)*

[‡] *Engineering Ingegneria Informatica S.p.A., Palermo, ITALY, (e-mail: name.surname@eng.it)*

[§] *Whirlpool EMEA S.p.A., Biandronno, ITALY, (e-mail: name.surname@whirlpool.com)*

[‡] *VDL Weweler, Apeldoorn, NETHERLANDS, (e-mail: name.surname@vdlweweler.nl)*

Abstract: In the last few years, manufacturing systems are getting gradually transformed into smart factories. In this context, an increasing number of information and communication technologies is incorporated towards facilitating management, production, and control processes. The introduction of advanced embedded systems with enhanced connectivity produces a vast amount of data, posing a challenge in terms of data analytics. However, the in-time collection and analysis of acquired data can create insight into the manufacturing process as well as its assets. One aspect of major importance for every production system is preserving its equipment in operational condition, and within those limits that could minimize unplanned breakdowns and production stoppages. This paper details the predictive analytics methodology integrated into the SERENA platform able to: (i) streamline the prognostics of the industrial components, (ii) characterize the health status of the monitored equipment, (iii) generate an early warning related to the condition of the equipment, and (iv) forecast the future evolution of the monitored equipment's degradation. To demonstrate the effectiveness of the proposed methodology, different use cases are discussed with results obtained on real-data collected in real-time from the industrial environments.

Copyright © 2020 The Authors. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0>)

Keywords: Data analytics, data management and analytics architecture, predictive analytics, production systems.

1. INTRODUCTION

The emergence of technological advances in computer systems and connectivity have given birth to concepts such as the Internet of Things (IoT) and Cyber-Physical systems Greer et al. (2019), which facilitate the acquisition and analysis of huge volumes of data Ahmed et al. (2017). In turn, those data can create insights over underlying

processes, which when combined with advanced machine learning techniques can help to identify similarities and underlying patterns that are not easily visible to the human operator or even human expert Yang et al. (2019).

Data analytics can be defined as the means by which raw data is converted into useful information and value Bumlauskas et al. (2017). In a factory, downtime can be costly and machine failures are also potentially dangerous for the operators. In this context, distributed sensors can capture multiple parameters from the shop floor and through a proper feeding mechanism. The analysis enables

^{*} This research leading has been partially funded by the European Commission under the H2020-IND-CE-2016-17 program, FOF-09-2017, Grant agreement no. 767561 "SERENA" project, VerSatilE plug-and-play platform enabling REmote predictive mainteNance.

prediction when a machine will need maintenance in advance of its failure, reducing the risk of a downtime or reducing the cost of repair Jin et al. (2016).

Nowadays, several companies across the world are turning to predictive analytics to increase their savings. It is worth mentioning that, according to Diab et al. (2017), predictive analytics and smart connected products are the two highest-ranked advanced manufacturing technologies. Tata, Schneider, American Electric Power are just some of the companies investing heavily in their tech infrastructure to enable early warning notification systems. In turn, this has allowed to enable the identification of an imminent failure and profit losses.

In this context, this work discusses on a lightweight industrial internet of things (IIoT) platform to facilitate data exchange and analysis across different stakeholders and users towards enabling predictive analytics on the shop floor and for versatile cases. In particular, the proposed architecture can identify the symptoms of imminent machine failure, through the characterization of the current dynamics of the process/machine using data collected in the factory. A scalable and modular approach has been taken in the design of the architecture, decoupling the overall design from any specific set of technologies. A prototype implementation has been tested in a use case related to the white goods industry and a discussion on its application on the steel industry is also included, focusing on the transferability of the proposed analytics methodology.

2. RELATED WORK

Referring to some state-of-the-art methodologies that address the problems of predictive maintenance research, in Wang et al. (2016) the authors present an intelligent factory model that incorporates industrial networks, the cloud, and supervisory control terminals with intelligent objects such as machines, conveyors and products. In this way, the results are obtained thanks to a self-organized system that uses feedback and coordination of the central control in order to achieve high efficiency. In such a system, more flexible tools and platforms are needed to analyze the huge amount of data collected by the sensorized equipment, able to guarantee a good quality of the data collected, limiting the anomalies. Good data quality is a fundamental condition to estimate some important parameters such as the RUL, which can be estimated in different ways as described in Zhang et al. (2016), and Asmai et al. (2011). Methods to improve data quality are discussed in Chen et al. (2013), where the authors describe how data sets can be displayed, grouped, classified, and evaluated in order to detect and remove outliers. Furthermore, the importance of having properly structured data for training and testing Machine Learning (ML) models is discussed in Aremu et al. (2018). In the aforementioned study, a framework is proposed using entropy measurements to minimize information loss during data processing, while preserving important information about the asset life cycle.

Some of this data can also be analyzed in real-time, as in the architectures described in Apiletti et al. (2018), both distributed on state-of-the-art open-source frameworks (e.g. Apache Kafka, Apache Spark, Cassandra). The first of these architectures also provides the integration of

a visualization tool, while the second features an auto-tuning engine for predictive maintenance.

Lee et al. (2014) examines the trend of production transformation in industrial 4.0 environments and assesses the adaptation of IT tools in managing industrial Big Data and predictive maintenance operations. Moreover, to overcome the complexity of Big Data pipelines, in Ardagna et al. (2017) the authors propose a new methodology based on Model-Driven Engineering (MDE), intending to limit the skills needed to manage a Big Data pipeline. Similar efforts to try to reduce the need for domain experts have also been made in the Machine Learning context. Among the works to make ML solutions exploitable as a service, Ribeiro et al. (2015) propose an architecture to create a flexible and scalable Machine Learning service, while Yao et al. (2017) provide a detailed empirical comparison between MLaaS platforms, analyzing the effectiveness of fully-automated, turnkey and customizable systems.

3. DESCRIPTION OF INDUSTRIAL USE CASES

In this section three real use-cases are introduced requiring predictive analytics services to improve the efficiency of the production by avoiding production breakdowns. The data-driven methodology (either complete or only a portion of), which will be adequately described in section 5, can be fruitfully applied to all three use cases below. However, in this paper we only discuss the experimental results obtained in the context of a white-goods industry (3.1), since it is more complex than others thus requiring more analytics building blocks. In addition, a preliminary set of results obtained in the context of robotics use-case (3.2) has been discussed in Panicucci et al. (2020), and the validation of the methodology for the rolling-mill machine (3.3) is still running.

3.1 A white-goods industry

In this use case, a nozzle injecting a combination of two reacting chemicals has been sensorized in order to monitor the overall process. A model of the overall system conditions is desired to predict possible alarms and failures in the process Proto et al. (2019). Several signals were collected during the process: temperature of the chemicals involved, the pressure of the liquids before the injection, injection timing and quantity, ratio of the injected chemicals, etc. Moreover, a set of different alarms, each describing a specific failure, has been collected and associated with the original production cycle that presented such failure. In particular, in recent past, domain experts have observed that after an alarm on the piston, the change of the foam injection head has always been necessary for a short period (about a week). Unfortunately, although the importance of the piston alarm is known, the physical events that determine this kind of alarm are still not known.

So, based on these premises, we tried to propose a data-driven methodology in order to identify which variables could be good predictors of piston alarms. The final goal of the use case is to predict whether a given set of monitored production cycles (for example every 8 hours or every day), characterized through several signals (e.g., temperature, pressure, etc), will trigger an alarm on the piston at a given time horizon (for example, the next day).

3.2 A robotics industry

Nowadays automation means a huge complexity since the objective is to produce high volumes of product often without any stops. In typical production lines, there are numerous machines with an increasing number of robots.

In order to preserve the equipment in operational condition, preventive maintenance procedures are applied, with particular attention to robotic assets.

One component monitored in the maintenance activity is the axes belt tension; indeed, belts are in charge of transmitting the power of the motor to the adaptor and so are a key element in robot precision. The less the belt tension, the more the skidding and decay of it; on the other hand, the more the tensioning, the more the stress on the mechanical component and thus the overheating.

Towards artificially creating a high quality and useful for predictive analytics dataset, a testbed has been designed and implemented, called *RobotBox* Panicucci et al. (2020), facilitating the study of all the physical behaviours without the entire complexity of an actual industrial robot. The *RobotBox* is made of a motor, whose position data is gathered thanks to an encoder, an adaptor, a rubber belt for the transmission and a five kilos weight to simulate a realistic application. In order to conduct realist experiments and achieve reusable results, all *RobotBox* components corresponded to the actual components of a 6-axis industrial robot and specifically from the sixth one.

It was possible to define a zero pre-configured status of the belt tension using the procedures used in a robotic line which enables to define as well the correct level of the belt tensioning. In order to simulate the ageing of the belt and incorrect setup of it, a slider with a centesimal indicator has been introduced in order to create different tensioning classes in a reproducible manner.

The purpose of this use case is to develop some predictive maintenance algorithms to predict the belt tension; gathering current data every 2 milliseconds from the robot controller.

3.3 A Rolling mill machine

The requirements stemming from this use case are related to the production of trailing arms and consider the condition monitoring of a rolling mill machine. The rolling machine is part of an automated line producing a set of trailing arms with different characteristics depending on the production orders.

In particular, the objective is to evaluate the wear of the segments of a rolling mill machine aiming at increasing its operational lifetime. In turn, this is expected to have a direct impact on the maintenance cost of the production line as well as its productivity.

In greater detail, the rolling mill machine has three top and bottom segments which are used to form the trailing arm by applying forces. The existing preventive maintenance plan involves a downtime of approx. 120 minutes for segments replacement every 2 weeks. In order to enable predictive analytics, different measurements are acquired including position, straightness, length, etc. associated

with a specific product and time. These measurements were selected after a discussion with domain experts, who advised on the most influential parameters to monitor based on the final business objectives. Based on the discussion, a set of IoT devices were deployed on the machinery to monitor the selected parameters. The analytics tasks rely on the complete set of measurements monitored over time.

The data analytics target the correlation of the different measurements to the monitoring of the machine's condition and its association to a predictive maintenance indicator suggesting when the replacement of the segments should occur.

4. PROPOSED IIOT ARCHITECTURE

The proposed IIoT platform, the SERENA system, is built upon a lightweight micro-services architecture based on Docker Swarm virtualisation technology. The environment modularity facilitates the replacement of a service with other containerised alternatives, allowing the platform to be characterised as a “plug-and-play” solution. The plug-and-play concept is also extended to the management of data repositories. The platform services do not access the data storage directly but using some intermediate layers as an interface to them. This decoupling allows replacing independently either the data interfaces (a broker, a modular set of REST APIs), the storage format and technologies (HDFS, relational DBMS, distributed data systems) or both.

The SERENA platform is composed of two sub-systems: Edge and Cloud. On the Edge side, one or more gateways fetch the data coming from the plant sensors, processing and enriching them with metadata and finally sending them to the Cloud sub-system. On the Cloud side, data is processed by several components. The Reverse Proxy Certification Authority (RPCA), implements a two-way TLS authentication rejecting or accepting the incoming requests based on a set of security policies. The legitimate requests are forwarded by the RPCA to the Message Broker and from there to the appropriate cloud services.

The cloud services communicate with each other through JSON-LD messages which format is partially based on the MIMOSA CRIS 3.2 standard. These messages carry the data gathered from the plant machinery along with their associated metadata. The data portions of the messages (measurements and alarms) are stored in a Hadoop Cluster living inside the cloud. The metadata portions of the messages are stored into the MIMOSA Metadata Repository that associates them to the plant elements which they refer to. The Metadata Repository can interact with any authorized services through the RESTful APIs provided by the Metadata Web Service.

An illustration of the proposed architecture with three of its supported services is illustrated in Fig 1. The Predictive Analytics Service fetches the stored data and metadata to feed them to data analytics algorithms, while The Visualization components and the Scheduler, working together on the predictions, support the maintenance operator(s). The Predictive Analytics service is discussed in greater detail in the following sections.

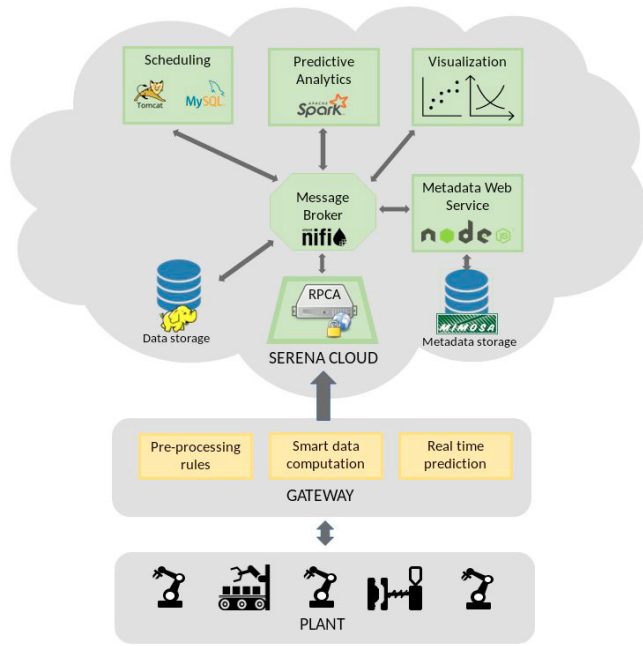


Fig. 1. Proposed IIoT platform architecture and supported services.

5. THE SERENA DATA-DRIVEN METHODOLOGY

The purpose of the machine-learning-based predictive-maintenance service is to predict alarms or failures in multi-cycle industrial processes, by creating a data-driven model on historical data. In order not to involve data scientists and domain experts in manual interventions, different strategies are exploited to describe cyclic time series data, aggregating them over multicycle horizons. In order to do this, two important analytic steps are implemented: Data-preprocessing and Supervised Learning. In the following subsections, each building block is described.

5.1 Data-preprocessing

In this first building block, outliers are identified and removed, for cleaning the data collection under analysis. To do this, the cycle length deciles are used, to remove the cycles belonging to the first or last decile. This approach was also validated by domain experts, who knew that some cycles were not real production measures but test cycles, for example. These cycles were successfully discarded. Additionally, to address the cyclic nature of the industrial processes under exam, an alignment task is performed to make the data fit a fixed-time structure through padding (the last value of the cycle is repeated until the cycle time slot is filled). Moreover, since the data collected by the sensors are raw time series, in the preprocessing phase this data are transformed into time-independent features set. This operation is performed through two main steps: *Statistics Computation* and *Smart Data computation*.

Statistics Computation. First of all, the original time series is split into contiguous portions, with the split size being a fixed parameter of SERENA. Then, for each portion, statistical features able to summarize the time series trend are computed, such as mean, standard deviation, quartiles, Kurtosis, Skewness, sum of absolute values,

number of elements over the mean, etc. It is important to notice that having portions of the same size is a choice of simplicity that has proven to work in our case, however the proposed approach can be successfully applied to splits of different sizes.

Finally, from the complete set of the numerous statistical characteristics, the most informative ones are selected and the unnecessary ones are discarded. To this aim, two techniques are exploited: (i) *multicollinearity-based*, that removes attributes whose values can be trivially predicted by a multiple regression model of the other attributes and (ii) *correlation-based*, which removes all the most correlated attributes, on average over all the (other) attributes.

Smart data computation. Often, in cyclic industrial processes, there is no interest in predicting the alarming-condition of a specific cycle, but that of a longer period, such as hours or days, which span over many cycles. The single cycle is too short for the target degradation phenomena and its prediction horizon. For this reason, this step aggregates the statistics data cycle-related features over longer, multi-cycle, time windows. The aggregation is performed separately for each split. SERENA captures the degradation of each statistics data feature by computing a linear regression on the aggregated multi-cycle period, and records the slope and intercept coefficients. Furthermore, for each attribute, the min, max, the mean and standard deviation of the values within the multi-cycle time window are recorded. It is worth noting that both the feature selection and the feature aggregation preserve the meaning of the measurements in terms of human readability, hence keeping the approach transparent and its decisions easily accountable.

5.2 Supervised Learning

This building block consists of two steps: *model building* and *real-time prediction*. The *real-time prediction* step simply labels the new data, based on an already built model. The *model building* instead performs the training on the historical data and extracts the relations among the data and the prediction labels (alarming conditions registered in the past). SERENA exploits two different classification algorithms: Decision Tree Classifier and Linear Support Vector Machine, automatically selecting the best performing one according to a score metric (defaulting to F-Score). However, other classification algorithms can be easily integrated and tested in the project, in order to compare more results. To assess the validity of the algorithms, a *Stratified K-Fold Cross Validation* is performed, that assures good robustness of the evaluation. This strategy equally divides the dataset into K folds (keeping the proportions of the original label distribution in each fold) and, in K iterations, alternatively uses a fold as the test set and the other $K - 1$ as the training set. In this way, at each iteration, a model is used to test a set of cycles always different than the ones used to train the model.

The most performing parameters of each algorithm are chosen by SERENA itself, thanks to a self-tuned strategy based on a grid optimization search over the classification algorithm parameters.

Finally, the predictive model is validated by calculating some metrics such as the precision, the recall and the F-Score for the class of interest (i.e. the alarming conditions or failures).

6. EXPERIMENTAL RESULTS

In this section the first set of experimental results obtained in the context of a white-goods use case (Section 3.1) are discussed. The data were collected in a WHIRLPOOL production plant and the objective of the analytics task is to predict in advance possible alarms and failures in a given phase of the production line. In particular, the data collected in the period from 13/09/2019 to 07/02/2020 have been considered, and the analysis aims to predict an alarm on the piston. The number of cycles analyzed is 44,544 and the greatest difficulty in effectively predicting a fault on the piston is the fact that only 14 alarms occurred on the piston during the time in question. Besides, these 14 alarms have always occurred in pairs, with a time distance of a few minutes from each other.

Each cycle immediately preceding an alarm has been labeled as 1, while all other cycles have been labeled as 0. All cycles are then aggregated daily and each day is labelled as 1 if it contains at least one cycle preceding an alarm. Of the 148 days obtained after aggregation, only seven of these are labeled as 1. In particular, each of these seven days contains two cycles that immediately preceded an alarm signal on the piston. Figure 2 shows the number of cycles in each day. The seven red vertical lines are positioned at the days on which the two alarm signals occurred.




Fig. 2. Number of signals in each day

In order to effectively predict the alarm signals on the piston, different signals and combinations of them have been tested as input, evaluating which of these would give better results. Towards this, a Decision Tree and a Linear SVM have been employed as classifiers, and the accuracy and F1-Scores values for the various input signals were calculated. It is important to note that, since the labels are strongly unbalanced, the most significant result is considered to be the F1-Scores relative to the minority class (F1-Scores(1)). Table 1 and table 2 report the results found, highlighting in bold the major F1-Scores(1) values between the two classifiers. In the tests, seven different signals were used as input and some combination of them.

Table 1. Decision Tree metrics.

Input	Accuracy	F1-Scores(0)	F1-Scores(1)
Signal1	0.97	0.99	0.71
Signal2	0.91	0.97	0.29
Signal3	0.92	0.97	0.43
Sig1+Sig2	0.91	0.98	0.14
Sig1+Sig3	0.92	0.95	0.57
Signal4	0.97	0.98	0
Signal5	0.96	0.97	0.17
Signal6	0.97	0.99	0
Signal7	0.97	0.99	0.17

Table 2. Linear SVM metrics.

Input	Accuracy	F1-Scores(0)	F1-Scores(1)
Signal1	0.92	0.99	0.14
Signal2	0.92	0.98	0.29
Signal3	0.95	0.99	0.43
Sig1+Sig2	0.94	0.97	0.57
Sig1+Sig3	0.94	0.99	0.29
Signal4	0.98	0.99	0
Signal5	0.98	0.99	0
Signal6	0.98	1	0
Signal7	0.98	1	0

The two tables clearly show how the last four signals have less impact on alarm prediction than Signal1, Signal2 and Signal3. That's because these last 3 signals are collected by sensors on the head, while the other signals are collected by sensors positioned near the tanks containing the products to be mixed and for this reason the results are less accurate.

6.1 Changing horizon and time aggregation

After finding these first results, additional experiments were conducted to predict the alarm signal earlier (1 day, 2 days, 3 days before), allowing technicians to have more time to intervene on the machinery and prevent the alarm. The current model allows predicting whether a particular day contains a cycle leading to an alarm signal. Now instead, our goal is to predict whether a particular day will lead to an alarm on the following days. To implement this, the labels were shifted by (0,1) indicating the presence of an alarm by one day. In particular, if previously it was labeled with 1 the day in which an alarm occurred, now instead it is labeled as 1 the day before this (if the horizon is 1 day).

Table 3 shows the results obtained using the Signal1 as input, trying to modify the horizon up to three days in advance. In addition to the horizon, in these experiments, the time of aggregation of the signals is also modified, trying to group the signals in windows of two or three days.

Table 3. Horizon and time aggregation.

	F1-Scores(0)	F1-Scores(1)
Agg=1, H=0	0.99	0.71
Agg=1, H=1	0.95	0
Agg=1, H=2	0.90	0.57
Agg=1, H=3	0.96	0
Agg=2, H=2	0.92	0.17
Agg=2, H=3	0.82	0
Agg=3, H=3	0.92	0.33

The table shows the values of F1-Scores(0) and F1-Scores(1), for each type of aggregation and horizon. The

number of days of aggregation is indicated by 'Agg', while 'H' indicates the horizon. As we can see in Table 3 the performance in terms of F1-Scores(1) drastically changes from a configuration set to another, this is mainly due to the complexity of the analytics task and the amount of data under analysis. Based on the available data, only the configurations (i) Agg=1, H=0 and (iii) Agg=1, H=2 provide quite good results.

7. CONCLUSION AND FUTURE WORK

A prototype IIoT platform is proposed and implemented based on a lightweight micro-services architecture, making it scalable and not dependant to specific technologies, enriched with an innovative certification mechanism, the RPCA.

Moreover, a data-driven analytics methodology has been presented as enabled through the aforementioned platform and concerning the need for development and deployment in versatile industrial cases with a significantly different set of datasets and requirements. The results obtained in the analyzed use case are very promising, especially considering the high imbalance class distribution and limited data availability. Despite these difficulties, the proposed methodology has proven to be able to make good predictions about the alarm signal. In addition, some acceptable preliminary results have also been obtained by extending the horizon-time, thus trying to predict the failure with more notice.

In conclusion, future work will focus on integrating additional functionalities to the predictive analytics services for smart manufacturing including (i) innovative strategies to address concept-drift detection Cerquitelli et al. (2019), (ii) general-purpose techniques to derive the Remaining Useful Life of the machine under analysis by analyzing the frequency of relevant alarms over time, and (iii) self-configuring data analytics workflow able to automatically identify the predictive analytics functionality needed by the real-life application under analysis.

REFERENCES

- Ahmed, E., Yaqoob, I., Hashem, I.A.T., Khan, I., Ahmed, A.I.A., Imran, M., and Vasilakos, A.V. (2017). The role of big data analytics in Internet of Things. *Computer Networks*, 129, 459–471.
- Apiletti, D., Barberis, C., Cerquitelli, T., Macii, A., Macii, E., Poncino, M., and Ventura, F. (2018). istep, an integrated self-tuning engine for predictive maintenance in industry 4.0. In *IEEE ISPA/IUCC/BDCloud/SocialCom/SustainCom 2018, Melbourne, Australia, December 11-13, 2018*, 924–931.
- Ardagna, C.A., Bellandi, V., Ceravolo, P., Damiani, E., Bezzi, M., and Hebert, C. (2017). A model-driven methodology for big data analytics-as-a-service. In *2017 IEEE International Congress on Big Data*.
- Aremu, O.O., Palau, A.S., Parlikad, A.K., Hyland-Wood, D., and McAree, P.R. (2018). Structuring Data for Intelligent Predictive Maintenance in Asset Management. *IFAC-PapersOnLine*, 51(11), 514–519.
- Asmai, S.A., Basari, A.S.H., Shibghatullah, A.S., Ibrahim, N.K., and Hussin, B. (2011). Neural network prognostics model for industrial equipment maintenance. In *2011 11th International Conference on Hybrid Intelligent Systems (HIS)*, 635–640.
- Bumblauskas, D., Nold, H., Bumblauskas, P., and Igou, A. (2017). Big data analytics: transforming data to action. *Business Process Management Journal*, 23(3), 703–720.
- Cerquitelli, T., Proto, S., Ventura, F., Apiletti, D., and Baralis, E. (2019). Towards a real-time unsupervised estimation of predictive model degradation. In *Proceedings of the International Workshop on Real-Time Business Intelligence and Analytics, BIRTE 2019, Los Angeles, CA, USA, August 26, 2019*, 5:1–5:6.
- Chen, Y., Zhu, F., and Lee, J. (2013). Data quality evaluation and improvement for prognostic modeling using visual assessment based data partitioning method. *Computers in Industry*, 64(3), 214 – 225.
- Diab, W.W., Harper, K.E., and Lin, S.w. (2017). Industrial Analytics : The Engine Driving the IIoT Revolution. 0–16.
- Greer, C., Burns, M., Wollman, D., and Griffor, E. (2019). Cyber-Physical Systems and Internet of Things NIST Special Publication 1900-202 Cyber-Physical Systems and Internet of Things. Technical report, National Institute of Standards and Technology.
- Jin, X., Weiss, B.A., Siegel, D., and Lee, J. (2016). Present Status and Future Growth of Advanced Maintenance Technology and Strategy in US Manufacturing. *International journal of prognostics and health management*, 7(Spec Iss on Smart Manufacturing PHM), 012.
- Lee, J., Kao, H.A., and Yang, S. (2014). Service innovation and smart analytics for Industry 4.0 and big data environment. *Procedia CIRP*, 16, 3–8.
- Panicucci, S., Nikolakis, N., Cerquitelli, T., Ventura, F., Proto, S., Macii, E., Makris, S., Bowden, D., Becker, P., O'Mahony, N., Morabito, L., Napione, C., Marguglio, A., Coppo, G., and Andolina, S. (2020). A cloud-to-edge approach to support predictive analytics in robotics industry. *Electronics*, 9(3), 492–513.
- Proto, S., Ventura, F., Apiletti, D., Cerquitelli, T., Baralis, E., Macii, E., and Macii, A. (2019). Premises, a scalable data-driven service to predict alarms in slowly-degrading multi-cycle industrial processes. In *2019 IEEE International Congress on Big Data, BigData Congress 2019, Milan, Italy, July 8-13, 2019*, 139–143.
- Ribeiro, M., Grolinger, K., and Capretz, M.A.M. (2015). Mlaas: Machine learning as a service. In *2015 IEEE 14th International Conference on Machine Learning and Applications (ICMLA)*. doi:10.1109/ICMLA.2015.152.
- Wang, S., Wan, J., Zhang, D., Li, D., and Zhang, C. (2016). Towards smart factory for industry 4.0: A self-organized multi-agent system with big data based feedback and coordination. *Computer Networks*, 101, 158–168.
- Yang, H., Kumara, S., Bukkapatnam, S.T., and Tsung, F. (2019). The internet of things for smart manufacturing: A review. *IIEE Transactions*, 51(11), 1190–1216.
- Yao, Y., Xiao, Z., Wang, B., Viswanath, B., Zheng, H., and Zhao, B.Y. (2017). Complexity vs. performance: Empirical analysis of machine learning as a service. In *Proceedings of the 2017 Internet Measurement Conference, IMC '17*, 384–397. ACM, New York, NY, USA.
- Zhang, C., Lim, P., Qin, A.K., and Tan, K.C. (2016). Multiobjective deep belief networks ensemble for remaining useful life estimation in prognostics. *IEEE Transactions on Neural Networks and Learning Systems*, (99), 1–13.