

# Framework for the Identification of Rare Events via Machine Learning and IoT Networks

Pedro Nardelli<sup>1</sup>, Constantinos Papadias<sup>2</sup>, Charalampos Kalalas<sup>3</sup>, Hirley Alves<sup>4</sup>, Ioannis T. Christou<sup>2</sup>, Irene Macaluso<sup>5</sup>, Nicola Marchetti<sup>5</sup>, Raul Palacios<sup>6</sup>, Jesus Alonso-Zarate<sup>3</sup>

<sup>1</sup>LUT University, Finland

<sup>2</sup>Athens Information Technology, Greece

<sup>3</sup>Centre Tecnològic de Telecomunicacions de Catalunya, CTTC/CERCA, Spain

<sup>4</sup>6G Flagship, University of Oulu, Finland

<sup>5</sup>Trinity College Dublin, Ireland

<sup>6</sup>Sociedad Española de Automóviles de Turismo, Spain

**Abstract**—This paper introduces an industrial cyber-physical system (CPS) based on the Internet of Things (IoT) that is designed to detect rare events based on machine learning. The framework follows the following three generic steps: (1) **Large data acquisition / dissemination:** A physical process is monitored by sensors that pre-process the (assumed large) collected data and send the processed information to an intelligent node (e.g., aggregator, central controller); (2) **Big data fusion:** The intelligent node uses machine learning techniques (e.g., data clustering, neural networks) to convert the received ("big") data to useful information to guide short-term operational decisions related to the physical process; (3) **Big data analytics:** The physical process together with the acquisition and fusion steps can be virtualized, building then a cyber-physical process, whose dynamic performance can be analyzed and optimized through visualization (if human intervention is available) or artificial intelligence (if the decisions are automatic) or a combination thereof. Our proposed general framework, which relies on an IoT network, aims at an ultra-reliable detection/prevention of rare events related to a pre-determined industrial physical process (modelled by a particular signal). The framework will be process-independent, however, our demonstrated solution will be designed case-by-case. This paper is an introduction to the solution to be developed by the FIREMAN consortium.

**Index Terms**—Industrial IoT, rare-event detection, predictive maintenance, machine-learning.

## I. INTRODUCTION

Predictive systems are usually designed to find recurrent patterns while excluding outliers. This is also true when analyzing methods of machine learning applied in industrial settings or in Industrial Internet of Things (IIoT) more generally. This paper, which is based on the project FIREMAN<sup>1</sup> funded by CHIST-ERA, focuses on this point by building an IoT network in combination with specific machine learning algorithms capable of detecting rare events in industrial setups in an ultra-reliable way. This, however, is not a trivial task due to the interdisciplinarity needed by the proposed solution. We offer below the state-of-the-art in different topics that will be covered by our solution.

<sup>1</sup>FIREMAN stands for *Framework for the Identification of Rare Events via Machine Learning and IoT Networks*.

## A. Literature review

**Cyber-Physical Systems (CPS) and IIoT:** In [1], the authors provide many examples of CPS in industry discussing the challenges involved and specifying the needs of different industries. [2] provides a good review of high-level technical challenges.

**Modelling complex systems:** Different approaches to model 5G and IoT communication networks can be found in [3]. In [4], the authors discuss different ways of modelling modern power grids, including not only the physical system but its relations with communication and decision-making networks. In [5], the authors introduce an integrative, three-layer, methodology.

**Time-based vs. event-based data acquisition:** Sensors acquire raw data to monitor physical processes via periodic sampling or via events. Depending on how the parameters are set, traffic in the communication network is generated, and each packet will contain information depending on the sparsity of the signal. The sampling strategy and the reliability of communication have been assessed in [6], while [7] dealt with the impact of the combination of these sampling methods.

**Machine-type communications and data aggregation:** The information exchange system is built upon ubiquitous wireless connectivity of machine-type communications (MTC) [8]. MTC may need to work in extreme cases: massive MTC (mMTC) and ultra-reliable low-latency communications (URLLC). The concept of data aggregation [9]-[11] is promising in this respect. Traffic from MTC devices is first transmitted to a special node called data aggregator that collects and processes the received data. Depending on the application, the aggregator can relay the processed data to the core network, use the data for a feedback control loop, or be used to monitor some metric.

**Big data mining for rare-event prediction and response:** Rare-event detection in time-series is a large field with many approaches [12]. Local Outlying Factor [13] is based on clustering multi-dimensional data points. Other techniques

borrow principles from speech processing [14]. Even more potent are Quantitative Association Rules (QAR) with high confidence & small support; recent results with a synthetic dataset of multidimensional vectors, with components drawn from Gaussian distributions have verified the effectiveness of this approach [15]. QARMA (standing for "Quantitative Association Rule Mining Algorithm") outperforms all other well-known machine learning methods in terms of both detection rate and false alarm rate, as shown in Table I. Finally, the detection of rare events can be also viewed as outlier detection in the context of sparse modeling and optimization techniques [16]. The GARD method in [17] and the KGARD in [18] were proposed as offline tools for robust linear and nonlinear estimation (supervised learning), respectively.

TABLE I

COMPARISON OF QARMA WITH STATE-OF-THE-ART CLASSIFIERS FOR RARE-EVENT DETECTION ON THE TEST PART OF A SYNTHETIC DATASET (70,200 POINTS, 200 RARE EVENTS, 50% TRAINING, 50% TESTING). ADAPTED FROM [15].

Method	True Detection in %	False Alarm in %
QARMA	100	0.002
JRip (Ripper)	97	0.06
MultiLayerPerceptron	98	0.017
AdaBoost.M1	91	0
Bagging	51	0
J48 (Decision Tree),	0	0
SMO (SVM)	0	0
Stacking	0	0
BayesNet	0	0
LogitBoost	0	0

**Predictive maintenance and rare events:** In classical periodic maintenance, parts may be replaced too soon, and supply chains may be disrupted, causing extra variability in the supply chain unnecessarily. [19] constitutes a review considering condition-based maintenance (CBM) as a decision-making strategy based on the detection of rare events. Yam *et al.* in [20] describe an integrated Decision Support System (DSS). In [21], the authors provide different performance trade-offs between frequency of unexpected breaks and unexploited lifetime.

### B. Structure of the paper

This paper will present an innovative framework to be proposed by FIREMAN by moving beyond the state-of-the-art in the above-mentioned fields in order to build an effective solution for rare event detection in industrial environments. The rest of this paper is organized as follows: Section II describes the general approach to be used by FIREMAN. Section III focuses on the method to solve specific problems and the progress beyond existing solutions. Section IV introduces the expected results to be delivered by FIREMAN and Section V provides the main impact insights of the project. Finally, Section VI concludes this paper.

## II. FIREMAN APPROACH

Our solution aims at providing a big-data-based optimized framework to predict and detect rare events in industrial

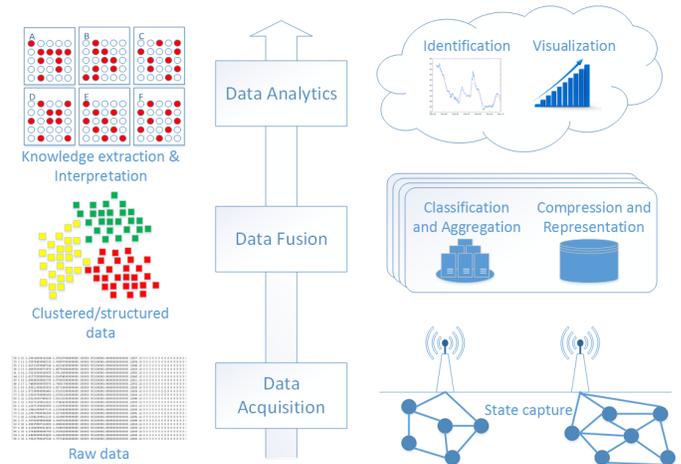


Fig. 1. The FIREMAN approach and concept illustration.

processes (mainly in maintenance), also including possible interventions. Figure 1 illustrates the general FIREMAN concept, built to answer the following scientific questions:

(1) Many industrial processes are physical and can be virtualized following three generic, autonomous but interrelated, steps: (i) acquisition of raw data and its pre-processing related to the sensor network deployment, (ii) fusion to aggregate and represent the raw data and (iii) analytics to build an intelligent prediction/detection of rare events. **Research question:** What is the optimal system architecture (i.e., design of each one of the three above steps and their integrated deployment) for detecting/predicting rare events for an industrial process with its own characteristics? FIREMAN will answer this by providing a framework based on CPS to guide the optimal system architecture design that is process-independent. However, for demonstration purposes the system design will be determined case-by-case. Our proposal will be tested for predictive maintenance in various environments: an automotive industrial plant, a base-station factory and a wind turbine. Both simulations and tests in real environment will be employed.

(2) In the acquisition step, there are many ways to obtain raw data: single vs. many sensors, periodic vs. event-driven sampling, random vs. planned deployments, single- vs. multi-hop transmissions, etc. Each industrial process can be evaluated by one or a set of relevant signals that shall be monitored. **Research question:** For a given industrial process, what are the main signal characteristics (e.g., periodic or sparse), what are the most suitable ways to collect, store and distribute the raw and pre-processed data within such sensor/IoT network? How can the communication (sensor) network be dimensioned and configured to attain the required performance metrics (e.g., reliability and low-latency) for the detection of rare events? FIREMAN will answer this by proposing a framework to model the industrial physical processes to assess and classify the signal characteristics. For the different classes of signal, different sets of collection-storage-dissemination deployments shall be analysed. A network architecture that is tuned to

rare events will be proposed and the corresponding network protocols (e.g., from 5G-PPP<sup>2</sup>) will be accordingly configured.

(3) Depending on the acquisition, the aggregation and representation of raw/pre-processed data require different treatment (e.g., periodic samples require different treatment from event-based ones). **Research question:** What are the most suitable data fusion options to aggregate/represent the heterogeneous big data sets generated by a specific acquisition deployment? FIREMAN will answer this question by studying different (heterogeneous big) data aggregation methods and assess the performance of different machine-learning algorithms to compress and represent the raw/pre-processed data after the acquisition step.

(4) The already fused data still need to be processed through analytics so rare events can be predicted/detected, and possible reactions proposed. **Research question:** What is the most efficient algorithm to deal with rare events in industrial processes, particularly for predictive maintenance? FIREMAN will answer this question by proposing new big data algorithms to deal with rare events, and developing visualization tools to help their identification and provide possible actuator reactions to cope with them. As mentioned, our framework shall be general, but the actual implementation is carried out according to the specific process being considered.

### III. FIREMAN METHODOLOGY

FIREMAN will employ the following methodology: (1) Definition of the system specifications and requirements, based on the needs and key performance indicators (KPIs) related to the detection of rare events, as provided by the industry itself. For this, we will rely on the consortium’s industrial partner, SEAT, as well as on inputs provided by other collaborating industrial players from various vertical fields (e.g., telecoms, energy, etc.). (2) Research work to achieve the project’s technical goals in its three key areas: large-scale data acquisition, big data fusion and big data analytics. The emphasis will be both on the development of novel techniques in each of these constituent areas, as well as on their optimal inter-working; the latter is an important novelty of the proposed work compared to the state-of-the-art. (3) Development and conducting of proof-of-concept trials and pilots that will showcase the validity and value of the proposed approach. The emphasis here will be both on the usefulness of the overall approach, and on its validation over different industry sectors. The scientific results and experimental trials will be disseminated in the scientific and industrial communities.

FIREMAN will seek radical new approaches to rare-event detection for predictive maintenance and other types of actuation. To reach this goal, the steps of acquisition, fusion and analytics are understood as constituent parts of an industrial CPS, modelled following the three-layer approach proposed in [5]. For the first step, we will focus on the characterization of the physical process data by employing traffic modeling techniques that account for event-based traffic

activation patterns [6, 22]. Furthermore, we will develop methods for data pre-processing, classification and storage that will enhance the process of rare-event detection. Finally, the process is enriched by the investigation of data transmission techniques that achieve scalable (mMTC), ultra-reliable low-latency (URLLC) and energy-efficient transmission, taking into account existing and ongoing work towards 5G standards [23, 24]. In the fusion step, we will investigate intra-cluster discovery algorithms, cluster-head selection policies, no-SQL databases and indexing from time-series databases, as well as dimensionality reduction techniques for big data, such as neural network-based approaches [25] and compressed sensing techniques [26, 27].

The big data analytics, in turn, will use the pre-processed big data from the previous step and will constitute the final stage for detecting/predicting/preventing rare events. The proposed solution will be based on new results from Quantitative Association Rule Mining (QARM) applied to very big datasets comprising multi-dimensional time-series (preliminary results illustrated in Table I). Using recent results in QARM (see [15]), as well as using results from the seemingly unrelated area of application of FOREX market direction prediction (see [14]) for detecting useful (even if rare) patterns of time-series move directions, as well as from robust regression, we shall build an ensemble of rare-event detectors that will be capable of detecting conditions that signal the need for maintenance operations even when the frequency of these signals does not exceed say, one occurrence per millions of sensor readings.

This capability is due to the ability of QARM to mine association rules with any required support the user defines that simultaneously maintain sufficient confidence, lift and any other combination of interestingness measures for the derived rules, and to the algorithm’s completely scalable (parallelizable in shared/distributed memory clusters/clouds) nature. As shown in [15], it is possible to mine all rules of the form:  $i_1[v^{ai1,1} \in r_{i1,1}] \wedge \dots \wedge i_n[v^{ain,m} \in r_{in,m}] \xRightarrow{s,c} j[p \geq l_j]$ , where the notation  $i[v \in r]$  denotes the “restriction” of item  $i$ ’s  $v$  attribute value to the range  $r$ . The rule is interpreted as follows: “whenever in a “user history” the item  $i_1$  appears with value for the attribute  $a^{i1,1}$  within the range  $r_{i1,1}$ , and the item  $i_n$  also appears with value for its attribute  $v^{ain,m}$  falling within the range  $r_{in,m}$ , then, with support at least equal to  $s$  and “interestingness metrics” at least equal to  $c$ , the item  $j$  also appears in the same user’s history with a value for  $p$  being greater than (or, if specified, with value exactly equal to)  $l_j$ .”

When detecting rare events in a manufacturing plant’s operations, a “user history” may simply correspond to a set of simultaneous sensor readings (each sensor reading corresponding to an “item” with a single quantitative attribute, its value) or any other derived feature from these readings (e.g., the immediately previously read sensor value). The consequence that is sought for may simply be the maximum reading of a particular sensor in the next 10 sampling periods or any other desired quantity such that, when its value exceeds a threshold

<sup>2</sup><https://5g-ppp.eu/>

this indicates the need to take some specific action e.g. halting operation and performing predictive maintenance.

The QARM framework is therefore theoretically very well suited for predictive maintenance based on rare event detection, because it can automatically discover all non-dominated rules that indicate that a set of (rarely occurring) pre-conditions triggers a maintenance need; we have also provided initial experimental evidence supporting the theoretical intuition as well: the preliminary results in Table I show that QARMA is indeed superior to all current competitors. By specifying sufficiently low support levels to capture the rarity of the events' occurring frequency, and at the same time specifying sufficiently high required confidence and lift or conviction values for the derived rules, the algorithm can produce a (possibly large) set of rules that trigger rarely, but when they do, they indicate with high likelihood the need for maintenance. Once produced, such rules are continually and automatically tested to see if they fire as new observations arrive. When they fire, they signal an alert for a maintenance need. In FIREMAN, we will further innovate by turning the QARMA algorithm into an online algorithm that continuously learns to adapt its rules as new evidence becomes available (tackling Big Data). Furthermore, we will investigate how to choose in a near-optimal manner what values to consider for each attribute when their domains become prohibitively large, without serious sacrifices in the quality of the results.

Similarly, the pattern mining system described in detail in [14] is based on a highly parallel algorithm that allows the prediction of the movement of any component of a multi-dimensional time-series above or below a certain threshold within an upcoming time interval, and again, the support levels for the patterns can be arbitrarily small, defined by the user. The idea is to "fit" a wave-form of a time-series in a so-called template-grid, and derive a similarity value between the current time-series fragment, and a "prototype" template-grid pattern that has the property that with significant support and confidence, to indicate that the time-series will exceed or fall below a certain percentage of its current value within a defined number of sampling periods in the near future, which should be an alert for maintenance.

A third approach for the online (real-time) detection of rare-events is to use machine learning techniques based on sparse representations. In FIREMAN, we intend to develop novel online algorithms for outlier detection, based on the state-of-the-art methods [17, 18, 28], which have been used as offline (batch) detection tools for robust regression over relatively small data sets (supervised-learning). The novelty of our learning approach is summarized as follows: a) develop online methods based on the aforementioned sparse optimization directions, b) investigate the best possible sampling techniques (stochastic versus mini-batch) and c) extend the regression tasks to their classification counterparts (which can be also viewed as regression with integer output variables, e.g., for the binary case +1,-1 - outlier versus no-outlier). As already mentioned, in FIREMAN, we shall build a pool of rare-event detection tools, consisting of robust regression/classification,

QARM, and FOREX market prediction techniques. The goal is to use these methods collaboratively, considering parallel or sequential use over different layers for the identification and prediction tasks. This approaches will be demonstrated in three sites: SEAT car manufacture, University of Oulu together with Nokia base-station manufacture, and LUT automation laboratories. Experiments with well-defined blocks will be considered to guarantee reproducibility. The setup will consist of at least one controlled machine with sensors, a communication network, and a workstation including data acquisition, management, analytics blocks, as well as a visualization block (refer to Figure 1) for at least one industrial process. Real users will be considered to define the performance metrics as reliability, low latency and usability with well-defined criteria, and to validate the proposed approach. From their feedback, it will be refined until users' validation.

#### IV. EXPECTED RESULTS

The overall objective of FIREMAN is to design, develop and showcase a novel big-data-based framework that encompasses all steps from sensing and data acquisition to statistical analysis and operational decisions, to accurately identify, detect, forecast and prevent rare events in a pre-determined industrial physical process. FIREMAN aims to build an architecture with a strong interplay among several research areas towards a highly-integrated CPS design at all data-processing levels. In this context, FIREMAN will provide breakthroughs in all three key pillars, i.e., large-scale data acquisition, big data fusion and big data analytics, considering the developments of relevant ongoing projects, but it will contribute far beyond state-of-the-art by pursuing the following objectives:

- Model the industrial physical process to retrieve the signal characteristics, e.g., time- vs. event-triggered sampling, from the physical data captured/recorded using sensors.
- Design a flexible network topology for the scalable deployment of many sensors to securely collect data at a low cost with a proper utilization of the scarce shared resources.
- Optimize the communication protocols to efficiently handle mission-critical transmission with stringent QoS requirements (ultra-high reliability, very low-latency) in industrial setups.
- Perform preliminary data analysis and processing to transform multi-stream raw data generated by various sources (heterogeneous, high-dimensional) into usable formats.
- Extract useful information from the monitoring of data and guide short-term operational decisions related to the physical process monitored by the sensors.
- Achieve proactive (predictive) maintenance via data-driven learning and mining techniques that identify and unlock the potential value of data.
- Offer system situational awareness on the occurrence of rare events for timely detection, effective decision-making and possible issuing of actuator commands.

- Develop experimental test-beds to enable the validation and evaluation of the proposed solutions and verification of the overall system performance.

## V. IMPACT

The results from FIREMAN will clearly have a strong impact on the field of predictive maintenance for IIoT systems; on the other hand, the specific challenges posed to the new generation of wireless networks will lead to important advances for reliable and massive-scale connectivity. From social perspective, FIREMAN will lead to a more connected society and to the implementation of radically new services based on a sustainable industrial value creation. Specifically, FIREMAN will achieve the following impacts:

- 1) Poor maintenance in production lines can reduce their productivity by 5 to 20% and unplanned downtime is estimated to cost 50 billion dollars each year [29]. The proposed predictive maintenance approach may eliminate such productivity losses, reducing the time required to plan maintenance by 20 to 50%, increasing equipment availability by 10 to 20%, and reducing overall maintenance costs by an estimate of 5 to 10%.
- 2) Rare events related to industrial plants result in serious risks to the working force and poor maintenance may cause accidents and permanent damage. FIREMAN's approach to identify and predict rare events will help to diminish such issues, improving the working safety in industrial plants.
- 3) More efficient communication and information processing in industrial plants may increase the efficiency of the plant in general, following the concept of Industry 4.0. The integrative proposed approach based on IIoT and CPS may create 10 million new qualified jobs in Western Europe and may provide useful benchmarks for European industry.

## VI. CONCLUSIONS

This paper provides an overview of an IoT-based framework to design an ultra-reliable rare-event detector using machine-learning, mainly focused on industrial settings. The proposed approach is holistic, from the physical process modeling to the end-user data visualization. Differently from existing solutions, FIREMAN will design general guidelines of how to jointly sample, communicate, fuse and perform analytics so that the desired rare events are reliably detected.

## ACKNOWLEDGMENTS

This work is supported by CHIST-ERA (call 2017) via FIREMAN consortium, which is funded by the following national foundations: Academy of Finland (n. 326270, n. 326301), Irish Research Council, Spanish and Catalan Government under grants TEC2017-87456-P and 2017-SGR-891, respectively. This work is partially funded by Academy of Finland 6Genesis Flagship (n. 318927) and ee-IoT (n.319009) and was supported in part by the Research Grant from Science Foundation Ireland and the European Regional Development Fund under Grant 13/RC/2077.

## REFERENCES

- [1] A. W. Colombo et al. "Industrial CPSs: A Backbone of the Fourth Industrial Revolution." *IEEE Ind. Electronics Mag.*, 2017
- [2] D. Xu, et al. "Internet of things in industries: A survey." *IEEE Trans. Ind. Inform.*, 2014.
- [3] N. Marchetti et al. "Complex Systems Science meets 5G and IoT." arXiv:1710.11548, 2017.
- [4] P. H. J. Nardelli et al. "Models for the modern power grid." *EPJ-ST*, 2014.
- [5] F. Kühnlenz, P. Nardelli. "Dynamics of complex systems built as physical, communication and decision layers." *PloS one*, 2016.
- [6] M. Tome, P. Nardelli, H. Alves, "Long-range Low-power Wireless Networks and Sampling Strategies in Electricity Metering." *IEEE Trans. Ind. Elect.*, 2018.
- [7] C. Kalalas, F. Vazquez-Gallego, and J. Alonso-Zarate. "Handling mission-critical communication in smart grid distribution automation services through LTE." *IEEE SmartGridComm*, 2016.
- [8] H. Shariatmadari et al. "MTC: Current status and future perspectives toward 5G systems." *IEEE Commun. Mag.*, 2015.
- [9] A. Laya, L. Alonso, and J. Alonso-Zarate. "Is the Random Access Channel of LTE and LTE-A Suitable for M2M Communications? A Survey of Alternatives." *IEEE Commun. Surv. Tutorials*, 2014.
- [10] Z. Dawy, et al. "Toward Massive Machine Type Cellular Communications." *IEEE Wirel. Commun.*, 2017.
- [11] J. Guo, et al. "Massive MTC with Data Aggregation and Resource Scheduling." *IEEE Trans. Commun.*, 2017.
- [12] B. Van der Paal, "A comparison of different methods for modelling rare events data." PhD thesis, Ghent University, 2014.
- [13] J. Lee, et al. "Integrating independent component analysis and local outlier factor for plant-wide process monitoring." *J. Process Control*, 2011.
- [14] N. Goumatianos, et al. "An Algorithmic Framework for Frequent Intraday Pattern Recognition and Exploitation in Forex Market." *Knowledge and Information Systems*, 2017.
- [15] I. T. Christou, E. Amolochitis, Z. H. Tan. "A Parallel/Distributed Algorithmic Framework for Mining All Quantitative Association Rules", *arXiv preprint*, arXiv:1804.06764, 2018.
- [16] A. M. Bruckstein, David L. Donoho, Michael Elad. "From sparse solutions of systems of equations to sparse modeling of signals and images." *SIAM review*, 2009
- [17] G. Papageorgiou, P. Bouboulis, S. Theodoridis, "Robust Linear Regression Analysis - A Greedy Approach." *IEEE Trans. on Signal Process.*, 2015.
- [18] G. Papageorgiou, P. Bouboulis, S. Theodoridis, "Robust Nonlinear Regression: A Greedy Approach Employing Kernels with Application to Image Denoising." *IEEE Trans. on Signal Process.*, 2017.
- [19] Y. Peng, et al. "Current status of machine prognostics in condition-based maintenance: A review." *Intl. J. Adv.d Manufact. Tech.*, 2010.
- [20] R. C. M. Yam et al. "Intelligent Predictive Decision Support System for Condition-Based Maintenance." *Intl. J. Adv.d Manufact. Tech.*, 2001.
- [21] G.A. Susto et al. "Machine learning for predictive maintenance: a multiple classifier approach." *IEEE Trans. Ind.Inform.*, 2015.
- [22] M. Neuts. "A versatile Markovian point process." *Journal of Applied Probability*, 1979.
- [23] A. Laya, C. Kalalas, F. Vazquez-Gallego, L. Alonso and J. Alonso-Zarate, "Goodbye, ALOHA!" *IEEE Access*, 2016.
- [24] P. Popovski et al. "Ultra-Reliable Low-Latency Communication (URLLC): Principles and Building Blocks", *IEEE Network*, 2018.
- [25] A. Stateczny, M. Włodarczyk-Sielicka. "Self-organizing artificial neural networks into hydrographic big data reduction process." *International Conference on Rough Sets and Intelligent Systems Paradigms*. Springer, Cham, 2014.
- [26] E. J. Candes, J. K. Romberg, T. Tao. "Stable signal recovery from incomplete and inaccurate measurements." *Communications on pure and applied mathematics*, 2006
- [27] A. Bhowmik, J. Ghosh, and O. Koyejo. "Sparse parameter recovery from aggregated data." *International Conference on Machine Learning*, 2016.
- [28] J. Liu, P. C. Cosman, B. D. Rao, "Robust Linear Regression via  $\downarrow_0$  Regularization," *IEEE Trans. on Signal Process.*, 2018.
- [29] C. Coleman, et al. "Making maintenance smarter: Predictive maintenance and the digital supply network." *Deloitte Insights*, 2017.