

Distributing Intelligence among Cloud, Fog and Edge in Industrial Cyber-physical Systems

Jonas Queiroz¹, Paulo Leitão², José Barbosa² and Eugénio Oliveira¹

¹University of Porto, Faculty of Engineering - LIACC, Porto, Portugal

²Research Center in Digitalization and Intelligent Robotics (CeDRI), Instituto Politécnico de Bragança, Bragança, Portugal

Keywords: Cyber-physical Systems, Edge Computing, Artificial Intelligence, Data Analysis, Multi-agent Systems.

Abstract: The 4th industrial revolution advent promotes the reorganization of the traditional hierarchical automation systems towards decentralized Cyber-Physical Systems (CPS). In this context, Artificial Intelligence (AI) can address the new requirements through the use of data-driven and distributed problem solving approaches, such those based on Machine-Learning and Multi-agent Systems. Although their promising perspectives to enable and manage intelligent Internet of Things environments, the traditional Cloud-based AI approaches are not suitable to handle many industrial scenarios, constrained by responsiveness and data sensitive. The solution lies in taking advantage of Edge and Fog computing to create a decentralized multi-level data analysis computing infrastructure that supports the development of industrial CPS. However, this is not a straightforward task, posing several challenges and demanding new approaches and technologies. In this context, this work discusses the distribution of intelligence along Cloud, Fog and Edge computing layers in industrial CPS, leveraging some research challenges and future directions.

1 INTRODUCTION

The era of the 4th industrial revolution (4IR), usually called I4.0, promotes the digitalization of traditional factories towards intelligent ones that are more flexible, robust, efficient, adaptive and competitive (Lu, 2017). This revolution is based on several disruptive ICT technologies, namely Internet of Things (IoT), Cloud Computing, Big Data and Machine-Learning (ML), which are glued by Cyber-Physical Systems (CPS) (Bauer et al., 2015). CPS comprise a set of networked cooperating and autonomous entities, combining cyber and physical counterparts. They are suitable to solve problems in complex and large-scale systems that can be found in smart manufacturing, smart cities, smart electrical grids and smart health (Leitão et al., 2016; Khaitan and McCalley, 2015).

In this context, Artificial Intelligence (AI) is a key enabler for the realization of the envisioned CPS and their features, like autonomy, self-awareness and dynamic reconfigurability. In industrial environments, the use of AI can contribute to develop intelligent decision support and control systems, smart machines and products, and consequently intelligent production systems and factories (Wuest et al., 2016; Lee et al., 2018). For instance, AI provides ML algorithms that

are capable to learn and extract patterns from data, devising models for prediction and data-driven decisions. The increased computational processing power together with the availability of huge amounts of data, mainly powered by the IoT, have contributed to their increasing adoption in various applications.

However, AI goes beyond ML algorithms, also providing distributed approaches, such those based on Multi-Agent System (MAS) (Wooldridge, 2002), which works as vessels for AI algorithms. MAS completely fits the CPS due to its inherent characteristics of modularity, autonomy and cooperation (Leitão et al., 2016). In such approach, the system behavior emerges from the interaction among autonomous agents, where the decisions are taken in a decentralized way, in opposition to centralized structures that are not able to address requirements related to flexibility, robustness, and on-the-fly reconfigurability.

This raises pertinent questions regarding which, how and where these key enabling technologies should be deployed. Traditionally, Cloud infrastructures provide on demand storage and processing capabilities. However, CPS go beyond traditional IoT applications, envisioning devices with embedded processing capabilities, capable to make decisions and interact with other entities. In this context, the de-

velopment of Fog and Edge Computing (CISCO, 2015) have allowed to deploy data processing capabilities close to the data sources, thus attending important industrial requirements, such as, data-sensitive, responsiveness and constrained network bandwidth (Breivold and Sandström, 2015). While the physical systems are deployed at the Edge, the cyber systems are distributed along the Fog and Cloud layers.

This paper aims to provide a discussion and raise concepts, technologies, challenges and benefits regarding how intelligence, particularly related to data analysis, can be distributed and balanced along Cloud, Fog and Edge layers for the development of industrial CPS. While some straightforward benefits encompass the design and development of more autonomous and adaptive systems, the challenges arise with the increasing complexity to manage such distribution.

The remaining of this paper is organized as follows. Section 2 overviews the concepts of Cloud, Fog and Edge Computing, and Section 3 discusses the distribution of intelligence by these layers. Section 4 discusses the related challenges in industrial CPS and points out some research directions. Finally, Section 5 rounds up the paper with the conclusions.

2 CLOUD-FOG-EDGE FOR CPS

In the context of CPS, these computing paradigms define what, where and how computational resources are deployed. Although the terms Fog and Edge Computing have been used interchangeably (Chiang and Zhang, 2016), this paper follows the definition where Fog comprises the intermediary layer between Cloud and IoT devices, while Edge considers IoT-based end devices with embedded processing capabilities.

2.1 Cloud - Everything as a Service

Cloud envisions high performance computing and data storage, promoting the shift from local corporate computational resources to remote large scale data centers, which are provided as services over the Internet (Duan et al., 2015). Cloud offers a large pool of virtualized resources, managed dynamically and transparently, and following a multi-tenant and pay-use model (Mell et al., 2011).

In spite of its benefits, the Cloud presents some drawbacks that hinder its widely adoption as the business model for many companies. For example, information security and privacy are important concerns (Xiao and Xiao, 2013), since companies fear having third-parties controlling their sensitive data. The network connection and bandwidth are other concerns,

mainly in data intensive scenarios where transfer large amounts of data can be expensive and lead to latency issues and network bottlenecks (Bonomi et al., 2014; Xu et al., 2014; Breivold and Sandström, 2015).

Despite the Cloud technology has been widely used by web-based companies, its adoption in other domains, such as health-care, transportation, energy and manufacturing, has increased considerably. In this context, Cloud provides a highly suitable framework to handle the massive amounts of data, attending their growing demand for information management and data analysis services. In the 4IR, it has been seen as a new manufacturing paradigm (Ren et al., 2017; Bauer et al., 2015; Lu, 2017).

2.2 Fog - Cloud Closer to Data Source

Although the Cloud offers efficient solutions to manage the huge volume of data, there are an increasing demand for real-time, data-sensitive and constrained network applications (e.g., in industrial domain) (Xu et al., 2014; Breivold and Sandström, 2015). In this context, Fog Computing (Bonomi et al., 2014) was leveraged as a complementary paradigm to cover the limitations of centralized Cloud approaches, addressing the challenges to perform data analysis locally, close to the data sources (Chiang and Zhang, 2016). It comprises a layer between Cloud applications and IoT devices, providing a more direct, reliable, secure and fast link between them, also reducing the amount of data and transmitting more meaningful information to Cloud. At this level, the data processing can be performed by network equipment, computers, mobile devices, as well as local servers (Bonomi et al., 2014).

Fog inherits most of the Cloud technologies, but considering constrained hardware. Moreover, it enables the decentralization of data analysis, decision-making and control, increasing local components autonomy, thus contributing to the shift of the traditional centralized automation pyramid to a decentralized, flexible and self-organized approach. Besides the aforementioned benefits, it faces some challenges, e.g., the use of non-general purpose devices, lack of frameworks for the development of algorithms optimized for constrained devices, and for the management of Fog nodes and their services (Shi et al., 2016).

2.3 Edge - The Edge of IoT

The extreme edge of the network is represented by IoT-based devices that usually have constrained computational resources, limited to collect and send data or execute commands. However, there is a growing number of IoT devices and platforms with consider-

able embedded computational resources (Al-Fuqaha et al., 2015). Like in Fog Computing, the focus is mainly to support local data analysis, but now it is directly embedded in the device itself. In this domain, there is an increasing number of AI applications (e.g., online speech, image and video processing), which are pushing the development of dedicated AI chips for such devices and related equipment (Ota et al., 2017).

The Edge also faces several limitations (Al-Fuqaha et al., 2015), mainly regarding the heterogeneity of devices and platforms, which increases the complexity of application development, leading to proprietary and non-standardized solutions. Other issue concerns how to enable Edge Computing in the huge number of legacy equipment that cannot be simply replaced by modern ones. Regarding security, although the local processing can help with the data anonymization, connect such devices to the network demand extra security procedures.

The Edge devices are represented by advanced IoT devices, such as those based on System-on-Chip and single-board computer (e.g., FPGA, Raspberry Pi and smartphones). They also include device gateways, i.e., devices that connect a set of highly constrained devices to the network. In industry such devices include Programmable Logic/Automation Controllers (PLC/PAC), which already present advanced processing capabilities and network connection based on standardized protocols like OPC-UA.

3 AI DISTRIBUTION IN CPS

3.1 Intelligence at Cloud, Fog and Edge

Currently there is a hype of using the Cloud to store the huge volume of the real-time data collected by a vast number of IoT devices, and to run powerful data analysis algorithms to extract valuable knowledge related to prediction and optimization. In fact, several reports, namely McKinsey (Bauer et al., 2015) and PwC (Geissbauer et al., 2016), refer Cloud Computing as a disruptive technology for the implementation of the Industry 4.0. However, the Cloud is not the solution for everything, with some data requiring to be processed as close as possible to its source and in real-time. Additionally, sending, storing and processing all the data in the Cloud could be expensive, constrained by the network latency and connectivity.

In this sense, the challenge lies in distributing intelligence by the Cloud, Fog and Edge layers, as illustrated in Figure 1. In fact, this distribution is strongly advised, where we “*should not send all collected data to be processed in the Cloud, but instead to make*

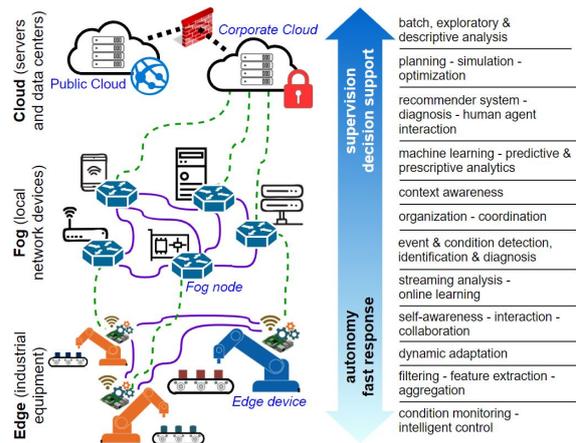


Figure 1: Intelligence distribution among Cloud-Fog-Edge.

analysis in the Edge” as sustained by J. Truchard from National Instruments at IFAC IMS’16, and “*analyzing data close to the device that collected the data can make the difference between averting disaster and a cascading system failure*” (CISCO, 2015).

Regarding the intelligence distribution, Figure 1 (right side) illustrates how different kinds of industrial related tasks that can be supported by AI approaches may be distributed from Cloud to Edge. This distribution is not unique, where the same task can coexist in different layers, handling different system functionalities. More details regarding the application of AI/ML in industrial environments can be found in the related surveys (Fei et al., 2019; Wang et al., 2018).

Table 1 summarizes several data-driven AI aspects of Cloud, Fog and Edge in industrial CPS. In general, at the Edge or Fog the intelligence is governed by rule-based and simple data analysis tasks. Since they comprise execution instead of development environments, i.e., they are not designed to build or train ML data models, which should be performed at the Cloud. For instance, classification and prediction data models can be easily deployed at these devices, since after the model creation, their execution comprises the simple verification of a set of rules or equations (e.g., decision tree or even Neural Network algorithms). For instance, regarding the Deep Learning, there are several works that propose the distribution of the neural network layers along Cloud to Edge, which besides offload the computation in central servers can also provide a local fast and partial response (Li et al., 2018).

Additionally, when considering Edge devices with their limited resource and very local and incomplete view of the system, the outcomes of the data analysis will be very simple and highly uncertain. On the other hand, at the Fog it is possible to reduce the levels of uncertainty, since the related equipment have more computational power and a wide view of the en-

Table 1: Data-driven AI aspects of Cloud, Fog and Edge in industrial CPS.

	Cloud	Fog	Edge
Scope	Big Data and AI platforms and tools for development, management, planning and optimization	Computing platforms for low-latency data processing and integration of cyber and physical components	End devices enhanced with monitoring, data preprocessing and filtering capabilities
Data volume, variety and velocity (<i>persistence</i>)	Huge volume of heterogeneous external, historical and stream data (<i>very long-term</i>)	Mid-volume of multiple data streams and no or short-term historical data (<i>mid/short-term</i>)	Low volume of local raw stream data (<i>very short-term/transient</i>)
Data processing type (<i>coverage</i>)	Advanced batch and stream data analysis (<i>huge number of heterogeneous sources</i>)	Lightweight data integration, batch and stream preprocessing and analysis (<i>many local sources</i>)	Simple stream preprocessing, filtering and data analysis (<i>single/very few sources</i>)
Responsiveness	Long and mid-term with optimal solutions	Near real-time with reasonable/reliable solutions	Very fast (hard real-time), but simple/non-optimal solutions
Intelligent features (<i>algorithm complexity</i>)	Complex pattern detection, prescriptive decision support (<i>Big Data and advanced ML</i>)	Local context-awareness, predictive decision support and dynamic adaptation (<i>mid/advanced ML</i>)	Self-awareness, autonomy and collaboration (<i>lightweight, rule-based and built-in data analysis</i>)
Hardware devices (<i>organization / location</i>)	Remote powerful data centers (<i>centralized / multiple hops</i>)	Local network equipment, micro servers and computing devices / platforms (<i>distributed / one hop</i>)	IoT and non-IoT device gateways (<i>decentralized / embedded, direct connection</i>)
Examples of devices in industrial contexts	General or industrial, public or private Cloud infrastructures and servers	Industrial network routers, supervisory and control room servers, workstation computers	PLCs, PACs, machines with embedded computing boards (FP-GAs), AGVs, Robots
Examples of supported tasks and applications	Business Intelligence, simulation, scheduling and interactive tools to support decision/operational tasks	Augmented reality, predictive supervisory and control systems (asset condition, product quality, process status)	Self-adaptive monitoring and control, dynamic reconfiguration, M2M collaboration

vironment. The diversity of Fog components enables to host almost any kind of AI-based application. In this sense, Fog nodes are responsible for a variety of data intensive applications, including the integration and processing of multiple real-time streams.

Considering all these aspects, at the Edge and Fog levels some examples of application that use AI in industrial automation includes system real-time monitoring, early detection of abnormality and local diagnoses (Aazam et al., 2018).

Besides its virtually unlimited computational resources, Cloud provides the centralization of data combined with powerful processing capabilities, which can ease the complexity of data analysis. This makes Cloud suitable for the development of Big Data analysis solutions for both operational and business levels, and specially considering their integration. Different from Edge and Fog, which are highly suitable for the operational levels, Cloud fits better the needs of enterprise and business levels, since it can face latency, bandwidth and security issues. In this sense, most of the operational tasks will require a private corporate Cloud or Fog-based infrastructure.

In summary, Cloud provides optimization and continuous improvement, but lacks responsiveness, while Edge devices are resource constrained and can easily saturate with processing, but offer real-time monitoring and fast response. Fog can help to reduce the overload of both Cloud and Edge, and at the same time handle other issues. On the other hand, when compared with Edge and Cloud, Fog may present

higher levels of complexity, mainly regarding its gateway role, since it needs to support the integration of a wide number of heterogeneous devices, protocols and technologies. In this sense, it is clear that there should exist a balance of the computational resources among each one of these layers, and just as important as, an appropriate interconnection of them.

3.2 Connecting Computing Layers

There is no doubt that the decentralization of data analysis from Cloud to Fog and Edge layers is essential to attend the 4IR requirements. However, it does not depend only on the decentralization of tasks, but in the way the components are interconnected vertically and horizontally, and how they can support each other. In this context, AI goes beyond ML data-driven approaches, where MAS can support the design and management of such components, allowing to develop intelligent and self-organized distributed systems. MASs are based on the concept of autonomous and cooperative agents, that work as a vessel for different kinds of processing and control AI algorithms. In this sense, agents can be distributed along the Cloud, Fog and Edge layers, encapsulating the system functionalities and collaborating with each other (Pico-Valencia and Holgado-Terriza, 2018; Wang et al., 2016; Leitão et al., 2016).

Although MAS can provide standardized communication interfaces that can cope with the complexity posed by the devices heterogeneity, usually they do

not attend hard real-time constraints or run in constrained devices, such those found at the Edge layer (Calvaresi et al., 2017). In such scenarios, agents may be deployed at Fog and Cloud layers, while the Edge components should interact using standardized service-oriented or lightweight IoT communication solutions, e.g., MQTT, CoAP or OPC-UA.

In this context, while the Edge covers the physical systems, the Fog and Cloud encompass an intermediate and a high-level layer of cyber systems. Fog directly supports Edge, mainly regarding time constrained tasks, while offloading Cloud systems. On the other hand, Cloud applications have a global view of the system, which enable them to support the lower layer components. Therefore, Cloud applications take advantage of the operational data combined with business information to provide actionable information for business level personnel and engineers. Besides supporting strategic decision-making, or operational engineers (e.g., enhancing their skills during the planning and execution of interventions), it can generate knowledge to be used in the different layers (e.g., updated rules for monitoring systems at the Edge).

4 CHALLENGES

As previously discussed, AI is an enabler to develop industrial CPS, playing two main roles: provide data-driven approaches to endow cyber and physical components with data analysis capabilities, and provide distributed system and knowledge management approaches to endow such components with autonomy, reasoning and collaborative capabilities. In this context, several cross concerns can be leveraged along the development of industrial CPS, such those illustrated in Figure 2. Based on that, some scientific and technological challenges can be identified, such those listed in Table 2. Their expected development difficulty and business priority are also presented, based on the 4IR visions, requirements and technologies.

4.1 Decentralization

The distribution of data processing capabilities contributes with the decentralization of decision-making and control, which is crucial for the fully fulfillment of industrial CPS, mainly regarding the system adaptability and autonomy. However, it is not a straightforward task, presenting several issues, for example the performance of conventional centralized data-driven AI approaches can be drastically affected when considering a small set of attributes or number of samples. Thus, this multi-level (glocal - global and lo-

cal) data analysis, demands special concerns regarding what kind of data processing capabilities should be performed in each computational layer, as well as the mechanisms to manage the components (horizontal and vertical) interaction and other communication and interoperability issues.

Additionally, mechanisms and policies should be designed to cope with the system complexity and emergent behavior, enforcing the components to respect a set of rules and constraints, thus also assuring the QoS and Trust, regarding the reliability and security of shared information and resources.

4.2 Intelligent Machines and Products

The development of intelligent machines and products is another key aspect in industrial CPS. The challenge is not limited to the use of the next generation of intelligent equipment that natively support AI features, but mainly to adapt traditional approaches and embed them in the existing constrained device platforms. This requires the need to consider smooth migration approaches.

The intelligence in such components will be achieved through the development of decentralized collaborative self-reconfigurable and coordination algorithms and policies. These mechanisms can increase the autonomy, flexibility and response time, but also the complexity of the system and its components, thus comprising a main trade-off when compared with centralized approaches. The system adaptation nervousness is another issue in such self-reconfigurable and distributed environments that should be managed.

On the other hand, enable these components to autonomously take decisions and control raises several security concerns, e.g., when considering their interaction with people. In this context, the lack of explicitly explanation for the outputs of ML approaches is one of the main reasons that prevents the fully autonomous AI-based control systems. Thus, requiring further research to assure the deterministic behavior of AI approaches and in what kind of industrial applications they could be safely applied. Also in this context, the certification of AI software is an important issue that should be addressed to assure the security and proper operation of such systems.

Another essential requirement in industrial domain is related to the development of predictive analytics. However, given the impact of prediction in industrial decision making, it is only valuable when presenting high accuracy and proper response time.

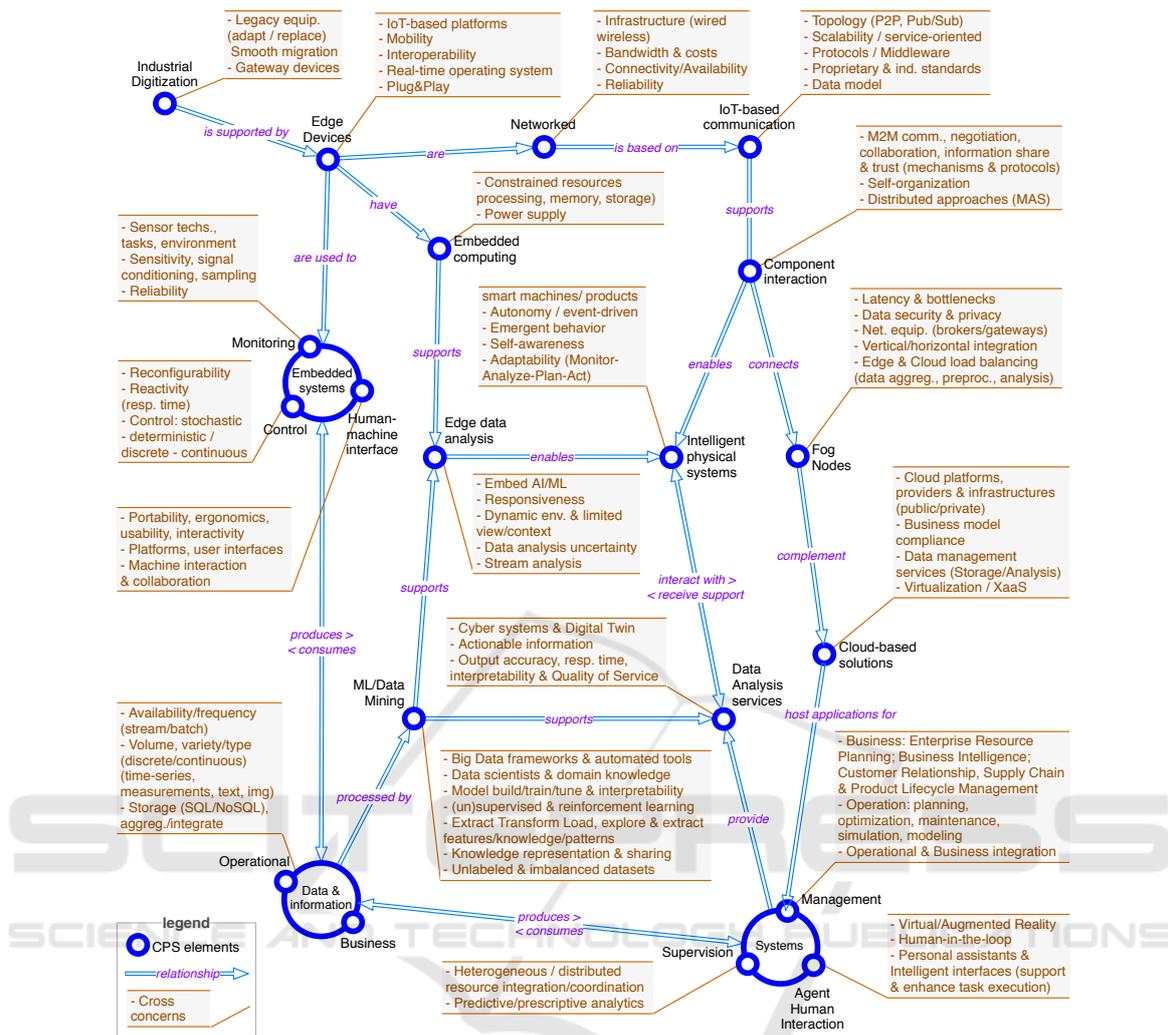


Figure 2: Cross concerns of AI decentralization towards industrial CPS.

4.3 Industrial Big Data

Industrial Big Data mostly comprises raw time-series data from hundreds to thousands of heterogeneous devices, complemented with business data (e.g., logistics, sales and customers). In this context, besides removing the barriers between business and operational layers, the challenge is to develop approaches capable to proactively use business information to redefine the course of operations and in the same way, adapt the business strategies based on operational data.

As aforementioned, data has been considered a key asset in any industry, however collecting and storing everything only increases the costs and complexity to extract value from it. In spite of the benefits of preprocess the data close to the data sources, this should be handled carefully, since some algorithms can mistakenly remove or cover important val-

ues. Additionally, some applications may require to transfer raw data, e.g., Digital Twin. This reinforces the need to harmonize the data processing at the Edge.

In industrial domain, mainly regarding the operational level, besides overcoming data uncertainty, reliability and incompleteness, the ML algorithms should be able to handle imbalanced or unlabeled datasets. Both are very common aspects in such environments, given that abnormalities are not so common, and usually are not proper labeled. This restricts the use of supervised learning algorithms, raising the need to adopt techniques to enforce the user feedback in the annotation of data samples. However, the user interaction at the Edge and Fog layers may not be feasible, thus leveraging the importance of unsupervised learning approaches in the operational levels.

Table 2: Challenges for implementing intelligence in industrial CPS.

Key Aspect	Major Challenges	Difficulty	Priority
<i>Decentralization</i>	Decentralized and collaborative data analysis	high	medium
	Synchronization and management of glocal data analysis	medium	high
	Control the system emergent behavior	high	low
<i>Intelligent machines and products</i>	Embed AI in constrained device platforms	medium	medium
	Self-reconfigurable and collaborative mechanisms	high	high
	Completely autonomous decision-making	high	low
	Predictive and proactive analytics	medium	high
<i>Industrial Big Data</i>	Synergy between operational and business information	medium	low
	Efficient and reliable data acquisition and management	low	high
	ML algorithms dedicated to industrial Big Data	medium	medium
	Support data uncertainty, reliability and incompleteness	high	medium
<i>Infrastructures and tools</i>	Enhanced automated data analysis tools	medium	low
	CPS frameworks and testbeds	medium	high
	Knowledge representation and sharing	low	medium
	Multi-disciplinary knowledge and expertise in AI and CPS	high	high
	Human Agent Interaction	high	high

4.4 Infrastructures and Tools

There are several challenges, not only related to frameworks, tools and standards for the development and testing, but also in terms of professionals with multi-disciplinary skills capable to develop and use such systems. In spite of the existence of several frameworks to assist and automate several data analysis tasks, usually being able to find ML models that outperform the ones developed by hand, most problems still require domain knowledge which refrains the complete process automation. This illustrates that even considering the higher level of automation, the human presence will be still required to develop, maintain and also operate such complex systems.

In the same sense, Human Agent Interaction approaches aim to provide dynamic, interactive and intelligent interfaces for personnel interacting software and physical systems. These approaches take advantage of computer vision, speech recognition and augmented/virtual reality, to create a partnership between human and system that can enhance their productivity and efficiency, as well as the quality of the outcomes.

Another important aspect is the need to test and evaluate the CPS solutions before the deployment in real environments, allowing to demonstrate their feasibility to industrial stakeholders. In this context, a major challenge lies in the development of dedicated CPS frameworks and testbeds, considering benchmarks and real case studies, not only capable to support the integration of the related technologies but also to test, monitor and evaluate the distributed and emergent behaviors.

5 CONCLUSIONS

In the 4IR, AI-based approaches cannot be seen anymore as just promising solutions, but instead as a must for any industrial system. On the other hand, in spite of the large potentialities offered by Cloud-based AI approaches, it is clear that they are not suitable to attend some important requirements of industrial CPS, namely the fast response, limited bandwidth and data sensitive. Moreover, sending all the raw data to the Cloud is costly and makes no sense, since most of it is meaningless and carries no significant value. However, the processing capabilities at or near the physical levels are very limited, not only regarding the constrained computational resources, but also the lack of context-awareness. In this context, this paper highlights the need and importance to achieve a balanced distribution of intelligence and computational capabilities among Cloud, Fog and Edge computing layers, aiming to achieve the desired levels of customization and dynamic adaptability envisioned by the 4IR.

The envisioned solution considers to take advantage of the technologies provided by each one of these layers to create a decentralized multi-level and collaborative data analysis infrastructure for the design and development of intelligent machines, products and systems. In this context, it is raised and discussed the emerging requirements, trade-off, issues and some major challenges that need to be properly addressed. It is also pointed out some research directions related to key concepts, approaches and technologies for the successfully exploitation of such industrial CPS.

ACKNOWLEDGEMENTS



This work is part of the GOOD MAN project that has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement N^o 723764.

REFERENCES

- Aazam, M., Zeadally, S., and Harras, K. A. (2018). Deploying fog computing in industrial internet of things and industry 4.0. *IEEE Transactions on Industrial Informatics*, 14(10):4674–4682.
- Al-Fuqaha, A., Guizani, M., Mohammadi, M., Aledhari, M., and Ayyash, M. (2015). Internet of things: A survey on enabling technologies, protocols, and applications. *IEEE Communications Surveys Tutorials*, 17(4):2347–2376.
- Bauer, H., Baur, C., Camplone, G., and et. al. (2015). Industry 4.0: How to Navigate Digitization of the Manufacturing Sector. Technical report, McKinsey Digital.
- Bonomi, F., Milito, R., Natarajan, P., and Zhu, J. (2014). *Fog Computing: A Platform for Internet of Things and Analytics*, pages 169–186. Springer.
- Breivold, H. and Sandström, K. (2015). Internet of things for industrial automation – challenges and technical solutions. In *2015 IEEE Int'l Conf. on Data Science and Data Intensive Systems*, pages 532–539.
- Calvaresi, D., Marinoni, M., Sturm, A., Schumacher, M., and Buttazzo, G. (2017). The Challenge of Real-time Multi-agent Systems for Enabling IoT and CPS. In *Proc. of Int'l Conf. on Web Intelligence*, pages 356–364. ACM.
- Chiang, M. and Zhang, T. (2016). Fog and IoT: An Overview of Research Opportunities. *IEEE Internet of Things Journal*, 3(6):854–864.
- CISCO (2015). Fog Computing and the Internet of Things: Extend the Cloud to Where the Things Are. Technical report, Cisco, white paper.
- Duan, Y., Fu, G., Zhou, N., Sun, X., Narendra, N., and Hu, B. (2015). Everything as a service (xaas) on the cloud: origins, current and future trends. In *8th IEEE International Conf. on Cloud Computing*, pages 621–628.
- Fei, X., Shah, N., Verba, N., Chao, K.-M., Sanchez-Anguix, V., Lewandowski, J., James, A., and Usman, Z. (2019). Cps data streams analytics based on machine learning for cloud and fog computing: A survey. *Future Generation Computer Systems*, 90:435 – 450.
- Geissbauer, R., Vedso, J., and Schrauf, S. (2016). Industry 4.0: Building the digital enterprise. Technical report, PwC.
- Khaitan, S. K. and McCalley, J. D. (2015). Design techniques and applications of cyberphysical systems: A survey. *IEEE Systems Journal*, 9(2):350–365.
- Lee, J., Davari, H., Singh, J., and Pandhare, V. (2018). Industrial artificial intelligence for industry 4.0-based manufacturing systems. *Manufacturing Letters*, 18:20–23.
- Leitão, P., Colombo, A., and Karnouskos, S. (2016). Industrial automation based on Cyber-Physical Systems technologies: Prototype implementations and challenges. *Computers in Industry*, 81:11–25.
- Leitão, P., Karnouskos, S., Ribeiro, L., Lee, J., Strasser, T., and Colombo, A. W. (2016). Smart agents in industrial cyber-physical systems. *Proceedings of the IEEE*, 104(5):1086–1101.
- Li, L., Ota, K., and Dong, M. (2018). Deep learning for smart industry: Efficient manufacture inspection system with fog computing. *IEEE Transactions on Industrial Informatics*, 14(10):4665–4673.
- Lu, Y. (2017). Industry 4.0: A survey on technologies, applications and open research issues. *Journal of Industrial Information Integration*, 6:1–10.
- Mell, P., Grance, T., and et al. (2011). The NIST definition of cloud computing.
- Ota, K., Dao, M. S., Mezaris, V., and Natale, F. G. B. D. (2017). Deep learning for mobile multimedia: A survey. *ACM Trans. Multimedia Comput. Commun. Appl.*, 13(3s):34:1–34:22.
- Pico-Valencia, P. and Holgado-Terriza, J. A. (2018). Agentification of the internet of things: A systematic literature review. *International Journal of Distributed Sensor Networks*, 14(10).
- Ren, L., Zhang, L., Wang, L., Tao, F., and Chai, X. (2017). Cloud manufacturing: key characteristics and applications. *International Journal of Computer Integrated Manufacturing*, 30(6):501–515.
- Shi, W., Cao, J., Zhang, Q., Li, Y., and Xu, L. (2016). Edge computing: Vision and challenges. *IEEE Internet of Things Journal*, 3(5):637–646.
- Wang, J., Ma, Y., Zhang, L., Gao, R. X., and Wu, D. (2018). Deep learning for smart manufacturing: Methods and applications. *Journal of Manufacturing Systems*, 48:144–156. Special Issue on Smart Manufacturing.
- 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.
- Wooldridge, M. (2002). *Introduction to Multiagent Systems*, volume 30. John Wiley and; Sons, Inc.
- Wuest, T., Weimer, D., Irgens, C., and Thoben, K.-D. (2016). Machine learning in manufacturing: advantages, challenges, and applications. *Production & Manufacturing Research*, 4(1):23–45.
- Xiao, Z. and Xiao, Y. (2013). Security and privacy in cloud computing. *IEEE Communications Surveys Tutorials*, 15(2):843–859.
- Xu, L. D., He, W., and Li, S. (2014). Internet of things in industries: A survey. *IEEE Transactions on Industrial Informatics*, 10(4):2233–2243.