An Initial Analysis of the Shortcomings of Conventional AI and the Benefits of Distributed AI Approaches in Industrial Use Cases

Anna Hristoskova, Nicolás González-Deleito, Sarah Klein, Joana Sousa, Nuno Martins, João Tagaio, João Serra, Carlos Silva, João Ferreira, Pedro M. Santos, et al.

An Initial Analysis of the Shortcomings of Conventional AI and the Benefits of Distributed AI Approaches in Industrial Use Cases

Anna Hristoskova[1], Nicolás González-Deleito[1], Sarah Klein[1], Joana Sousa[2,3], Nuno Martins[2,4], João Tagaio[2,5], João Serra[2], Carlos Silva[2], João Ferreira[2], Pedro M. Santos[6,8], Ricardo Morla[7], Luís Almeida[7,8], Barış Bulut[9], Sencer Sultanoğlu[10]

[1] Sirris, Brussels, Belgium
[2] NOS Inovação, Lisbon, Portugal
[3] Bold International, Lisbon, Portugal
[4] Caixa Mágica, Lisbon, Portugal
[5] KCSIT, Lisbon, Portugal
[6] Instituto Superior de Engenharia do Porto, Porto, Portugal
[7] Universidade do Porto, Faculdade de Engenharia, Porto, Portugal
[8] CISTER Research Center in Real-Time & Embedded Computing Systems, Portugal
[9] Enforma Bilişim, İstanbul, Turkey
[10] Eliar, İstanbul, Turkey
Author for correspondence: anna.hristoskova@sirris.be

Abstract. The centralised approach of IoT (Internet of Things) applications leveraging cloud infrastructures to address constraints at the level of end and edge nodes is no longer viable, especially for applications with hard real-time requirements and increasing AI (Artificial Intelligence) usage. This paper presents an initial analysis of the shortcomings of such centralised AI approaches applied to the five industrial use cases considered in the ITEA3 MIRAI project and discusses the expected benefits that distributed AI approaches will bring to these use cases, namely, to lift constraints such as computing power, bandwidth, latency, security and privacy.

Keywords: Artificial Intelligence (AI), Distributed AI, Secure Analytics, Access Control, Privacy.

1 Introduction

For the last two decades the usual approach for IoT applications has been to leverage on cloud infrastructures to address the computational and storage limitations and constraints at end and edge nodes. However, offloading processing capabilities to the cloud requires transferring data from edge devices to a backend cloud infrastructure, dealing with the limitations imposed by the underlying communication channels, depending on the availability constraints of both those communication channels and the cloud infrastructure, establishing proper mechanisms to secure data in transit and data at rest at the backend (esp. when privacy-sensitive data is considered) [1], dealing with

the costs incurred from the transmission of data and by the usage of the backend cloud infrastructure, etc. Depending on the application and underlying infrastructure, it might not be feasible to transmit all the data to the backend due to bandwidth limitations or cost constraints, resulting in only a fraction of the data being transferred, while the remaining data is discarded at a very early stage. In addition, the underlying communication channels' latency might be unacceptable for many industrial control application and applications with (hard) real-time constraints.

This has been changing in recent years. The advent of end and edge nodes with increased computational and storage capabilities makes it possible to perform a large range of increasingly resource-demanding computations (including AI-based tasks) locally at the edge, and only rely on a backend cloud for communicating results or performing computations requiring the combination and further processing of results from different edge devices or historical data sources. In addition, this increase in computation capabilities also enables neighbouring edge devices to perform tasks collaboratively, leveraging on each other's available computational resources, before offloading computations to a cloud backend. These new possibilities pave the way for the (further) development of Distributed AI.

In this regard, existing local computation platforms for AI are typically based on high-performance hardware [2]. They contain expensive and power-hungry high-speed processors and AI accelerators. These platforms offer sufficient computational power and fulfil most needs even in industrial application scenarios, but they often come with prohibitively high investments and energy costs. In addition, they usually do not scale very well and represent a single point of failure.

A promising alternative technology approach are embedded computing devices in a massively parallel and distributed architecture, which may overcome most of the difficulties mentioned above. With this approach, the computation and storage can be located very near to the data source, making latency and bandwidth non-critical issues. Data collection, machine learning (ML) and inference tasks of AI applications can be distributed in a federated architecture. Tasks can be migrated in case of component failure and continue their execution at other neighbouring devices, which offers robustness. Likewise, workloads can be shared among many different devices, which provides scalability.

However, suitable frameworks and building blocks for smart and sustainable planning, deployment and operation of such distributed IoT and edge AI-computing applications is still lacking. Here, smartness refers to a degree of autonomy (AI performed at the edge level), robustness (such as against connection cut-offs leading to insufficient data; external disturbances; single point of failure at the centre), ability to incorporate some of the global criteria and constraints in a centralised-like fashion even when most decisions are made at the edge, whereas the sustainability refers to a carefully-designed allocation of resources (such as computing, storage, and data transmission) as well as their expansion capability under increased demand. The ITEA3 industrial R&D project MIRAI, aims at building such a framework and corresponding building blocks through which IoT and edge computing applications can horizontally scale among edge devices, in addition to the vertical scaling to the cloud.

In this paper, we present each of the five use cases being considered in MIRAI, along with their needs at the level of AI and how distributed AI is expected to contribute to achieve their objectives. These use cases are provided by five different problem owner companies from three different countries (Belgium, Portugal and Turkey), and cover the following domains: renewable energy management, Internet provisioning for households, road traffic management, water consumption management, and dyeing for the textile industry. The paper concludes with a brief wrap up that presents the team's anticipations to move towards a common reference architecture model, MIRAI Framework Building Blocks (MFBB), indicating the types of models or standards whose logic or even parts of design could be recycled.

2 MIRAI use cases and rationale for distributed AI

This section presents the five use cases of the MIRAI project and discusses the shortcomings of conventional AI-based approaches and the needs for distributed AI for each of them.

2.1 Use case 1: Distributed renewable energy systems

With more and more distributed renewable energy assets such as solar panels and wind turbines installed, the electricity production becomes more fluctuating than with an old-fashioned coal-fired power plant. The grid needs to be increasingly capable of adapting to an ever-fluctuating residual load on different timescales, from seconds to months. Solar and wind assets, as well as storage, also need to cater for a better grid quality. At the same time, renewables worldwide are entering a "post-subsidy era", in which investors and asset owners need to secure bankable revenue streams. The combination of these two trends creates a need and opportunity for flexibility services, from switching off plants at negative energy prices, to charging and discharging battery systems to help stabilise the grid.

3E is a well-established Belgium-based company existing for already 20 years and counting about 100 employees. 3E provides consultancy and software solutions for monitoring and improving the performance of sustainable energy installations and for optimising energy consumption. More precisely, 3E's SynaptiQ Asset Operations[1] solution enables to manage solar and wind assets, monitor, report and improve their performance, and organise their maintenance. Thanks to MIRAI, 3E aims to provide optimised control of renewable energy plant assets and real-time status updates.

Current AI shortcomings. The energy infrastructure is evolving from a hierarchical to a fully distributed architecture. Whereas in the past the lower-level nodes acted as pure consumers, an increasing number of nodes also produce energy and act therefore as producers, the so-called prosumers. This puts an additional stress level on the grid to guarantee stability and balance between electricity demand and supply (e.g. through

[1] https://3e.eu/solutions/digital-solutions/asset-operations

energy storage systems). Nowadays, most solutions offer a centralised grid optimization service. This can though lead to exchange of information between competitors or to privacy violations, for example in case of single households [3, 4].

Currently, single assets can mainly be monitored via 3E's SynaptiQ Asset Operations platform but there is no knowledge shared between different assets. In case the energy market price becomes negative due to too high availability of energy in the electric grid, the only solution is to switch off the asset for a given time. So far, the assets provide a very basic forecast for negative energy prices only.

Distributed AI expectations. 3E would like to execute (part of) the grid optimization service locally, on a prosumer's infrastructure. Being able to distribute the service in that way would enable a faster response to changes in the current status of the grid, as well as local energy production and consumption, and hence reduce latency. Further, it would enable to minimise the amount of actual data sent to the cloud for further processing in order to preserve prosumers' privacy.

Within MIRAI, single assets will be optimally controlled in real-time as the data will be used locally and with shared knowledge, without actually sharing data. One example is the case of negative energy market price: When the prices in the wholesale energy market are negative, in future, the assets shall not only be turned off to avoid losing revenue but by including flexible assets to shift energy, the produced electricity can be used in an optimal way. Hence, batteries can be charged during negative prices and this energy can be used to maximise self-consumption and/or to sell flexibility on short-term or balancing markets in order to increase revenue [5]. The decision on what to do and when to do it will be based on AI algorithms that forecast the energy production, consumption and market prices more precisely in a distributed way [6, 7]. This will lead to an optimised control of the single assets.

2.2 Use case 2: Secure Internet provisioning

DDoS (Distributed Denial of Service) attacks are one of the major issues related to availability and security when customers are using cloud-based applications. A DDoS attack tries to deplete the resources of an application, rendering it unavailable to its legitimate users. DDoS attacks may be driven to any final public access point through Internet. Through MIRAI, NOS, one of the main telecom companies in Portugal, aims at developing services for DDoS protection with focus on the home network, combining the best practice of application design with the availability of a set of features for DDoS mitigation. The solution will be focused on customers' premises and the protection will be smart and simple, taking into account the following points:

1. All customers' devices, such as gateways/routers, set-top boxes, IoT devices must be usable and safe regarding internal or external DoS.
2. No single or multiple device(s) can cause exhaustion of network resources preventing other devices on the network to have access to the Internet for an amount of time higher than an identified goal (% time, etc.).

3. Any device trying to deny access or denying access to others, by consuming all available bandwidth, sending too many requests / packets per second or any other form of denial, should be set into a quarantine mode.
4. Depending on the type of device and its characterization, an IoT device has some periodic communication in the network (such as check for updates, send updated statistics, or some more random based on Human interaction, etc.). Leveraging this knowledge, it is possible to know what is the normal (usual) and abnormal (unusual) traffic for it.
5. Support for device quarantine should send a notification to customer notifying him/her that the device is not acting / behaving correctly for "amount of time" and is blocked from all communication.
6. Device quarantine will only block temporarily access to the device with unexpected behaviour. All devices connected to the router will remain with Internet access with no impact of isolation done in affected devices.

Based on the aforementioned points, smart and faster AI models are needed in order to easily detect, identify and mitigate attacks in customers' premises. This is also specifically critical with the exponential growth of IoT, but also with 5G delivery, where speed and latency will play an important role. But most importantly the high density of communication implies a huge amount of connected "things" at home that may be vulnerable to cyberattacks.

On the perspective of AI, IoT devices present network traffic patterns that differ from other types of devices (e.g. personal devices such as laptops and smart phones). As many IoT devices are meant to provide sensing and actuation over some physical process (e.g. opening of a door), much of their network traffic is event-triggered and can be directly traced to physical actions by the customer [8]. IoT devices also transmit periodic packets for keep-alive or logging purposes. In turn, DDoS attacks involve that a large number of distributed devices request a service, simultaneously or in short succession, from the target node, thus disrupting the ability of the node to serve legitimate clients. Such attacks involve considerable amounts of traffic targeting a single or few nodes; in turn, nodes contributing to the attack may also generate uncommon traffic patterns, either at low or high rates. Thus, DDoS detection is typically handled in the literature as a problem of anomaly detection over the 'typical' network traffic patterns; in turn, typical patterns need to be learned given that the customer's IoT ecosystem may exhibit a wide variety of legitimate traffic profiles. ML methods are particularly suited to this task. A considerable number of works have explored the problem in the context of software-defined networking (SDN), as the paradigm enables a bird's-eye view of the traffic in a network [9]; but there are also works taking a closer look at anomaly detection in customer networks and considering the specific characteristics of IoT traffic patterns [10, 11].

Current AI shortcomings. At the end of 2017, NOS launched a study in Portugal aiming at assessing attitudes and behaviours regarding Internet use and evaluating satisfaction with Internet Accesses Macro and Micro. 670 online interviews were carried out, targeting residents in the Portuguese territory, from 15 to 64 years old,

Internet users, whether or not they have Internet access (fixed or mobile). The study was focused on the usage of WiFi access and the corresponding concerns on several areas such as: e-mail, social network access, information search, chats, home banking and online purchases. In all instances, security was in the top 3 of the users' concerns. There are several products and services for security targeting DDoS attacks; however, they are focused on enterprise environment. With the growth of IoT and the rise of 5G as well as WiFi6, there is a need to shift the paradigm of DDoS. It is critical to see the customer's home as an enterprise environment, where several devices are connected, huge amount of data is exchanged (personal and behavioural), and people have low cybersecurity literacy and, consequently, are vulnerable to attacks.

Distributed AI expectations. MIRAI aims to develop services for DDoS protection with focus in the home network, combining the best practice of application design with the availability of a set of features for DDoS mitigation. The solution will be focused on customer's premises. The protection will be smart and simple. Thus, the service developed under the MIRAI framework will support features such as: active traffic monitoring & detection, automatic attack mitigations, availability guarantee, mitigation policies tuned to customers applications, metrics & alerts, mitigation flow logs and DDoS rapid response support. The DDoS mitigation should be always-on in terms of monitoring and analysis. This will be done through smart algorithms and filters before considering an anomaly. For that, data from traffic behaviour, consumed bandwidth per period and protocols normally used by the user will be analysed to classify the attack. This flow allows to reduce the false positives and negatives and to improve the efficiency of the mitigation actions. The development, implementation and scaling of an IoT framework based on smart auto-discover, smart interoperability and smart cybersecurity will be deployed.

2.3 Use case 3: Traffic management

Like many other aspects of our lives, cities and mobility are becoming smarter. Cameras, induction loops and smartphones, among others, monitor how we move and how road infrastructure is used, inform us about possible problems that we might encounter in our way, and route us through the most optimal path. Still, the safety of vulnerable road users (i.e. pedestrians and cyclists nearby railway crossings, in school streets and at complex intersections) remains an important problem with ample possibilities for improvement.

Macq, the provider of this use case, is a well-established Belgium-based family company existing for almost 100 years and counting more than 130 employees. Nowadays, most of Macq's solutions and research activities focus on smart mobility. In that area, Macq offers products ranging from sensor solutions for traffic monitoring, including advanced Smart Mobility cameras (able to operate on multiple lanes, distinguish between different types of vehicles, detect the number of passengers, estimate driving speed, etc.), to controllers for traffic light management at road intersections, and software packages for managing mobility-related data and extracting

valuable insights to different types of decision makers (such as police and road authorities).

Through MIRAI, Macq mainly aims at extracting valuable road safety analytics at the edge (specifically from their Smart Mobility cameras) and at reacting quickly when a potentially dangerous situation is detected. This requires in turn being able to deal with bandwidth-intensive image data, to extract insights even in the presence of noise and missing data, and to rely on other operational devices in the immediate surroundings, while considering their processing capabilities, current load, and connection link to them.

Current AI shortcomings. Especially for vulnerable road users, a given situation can look safe from a single point of view, but this risk assessment can easily change if additional aspects are considered. Imagine a child crossing a tram line with a tram in its back. One camera sees only the child and another one only the tram; both cameras have each a partial view of the situation and cannot detect an increased risk. Hence, several sensors and cameras are needed to provide a complete as possible view of complex road intersections, and an approach leveraging all the resulting information is needed, for which only little research has been conducted so far [12, 13].

While the most advanced of Macq's Smart Mobility cameras perform most of their computations locally, performing more advanced computations such as understanding a given situation (possibly based on information from several other nearby data sources) and to detect potentially dangerous situations when they occur, would require relying on a distributed infrastructure. On the one hand, the intrinsic nature of Macq's Smart Mobility camera data makes it challenging to merely transfer that data between other edge devices or to the cloud to perform the required computations. On the other hand, Macq's Smart Mobility cameras are deployed outdoors and are hence exposed to possibly harsh weather conditions. This results in incorrect readings, in communication problems, and, sometimes, in broken devices that need to be replaced.

Advanced ML and data mining approaches suitable for edge computing enable to leverage the computational power offered by edge devices. However, determining how to split tasks and perform computations between different neighbouring edge devices available and the cloud infrastructure remains a challenge to address ahead of fully exploiting the available edge and cloud infrastructure [14].

Distributed AI expectations. The issues described above can only be overcome in case the information about a situation is shared between several devices and an AI algorithm judges from this collective knowledge. Within MIRAI, a collectively acting algorithm will be developed, additionally respecting every road user’s privacy.

Being able to perform more advanced ML and data mining computations at the edge would enable a better exploitation of the computational edge capabilities, a more balanced usage of computational resources and faster response times, as bandwidth-intensive image data would not need to be transferred to the cloud. Additionally, the algorithms proposed will be able to operate even in the presence of noise and missing data to better ensure service continuity.

2.4 Use case 4: Water management

Water damages in buildings represent 31% of the operational costs of property & casualty insurance companies due to claims. In addition to reducing unnecessary water consumption and preserving drinking water resources, being able to detect water leakages as soon as they occur is crucial to prevent water damages in buildings.

Shayp is a young Belgium-based company providing solutions for the monitoring and management of water leakages for different types of buildings, ranging from small buildings such as individual homes to larger buildings such as schools, hospitals and office and administration buildings. Shayp collects water consumption information through a connected water metering device installed in those buildings. Measurements are collected every 6 minutes and are sent to the cloud through NB-IoT. From this data, Shayp extracts water consumption analytics and detects potential leakages present in a building.

Through MIRAI, Shayp aims at (i) reducing and optimising communication bandwidth, as it drains the most battery power; (ii) identifying anomalies already at the edge, enabling a faster detection of water leakages; and (iii) adding remote control features to its meters, enabling future updates and remote calibration.

Current AI shortcomings. In the current setup, water consumption data is hosted in the cloud where it is processed and stored. This is where the leakage detection analysis is performed. The data is transferred via the NB-IoT communication protocol every 6 minutes. If the data is sent with a higher temporal granularity, the battery lifetime of the device decreases [15]. Currently, the leakage detection is based on statistical analysis and a high accuracy is reached within a time window of 3 hours. In order to detect leakages earlier, the temporal resolution has to be higher. With the intention to keep the battery lifetime as long as possible, a trained ML model is planned to predict leakages on the device directly without the need of sending all data on a high temporal resolution.

Distributed AI expectations. Within MIRAI, Shayp will take a first step from a purely statistical model to a distributed ML model in order to improve and speed up leakage detection [16], improve building classification [17, 18] and increase battery lifetime. This model will run locally on the edge devices and send water consumption to the cloud only if necessary. With an unsupervised approach on the edge, the leakage detection time will be reduced from 3 to 1 hour with an at least stable if not increased battery lifetime.

2.5 Use case 5: Continuous auto configuration of industrial controllers at edge

The main challenge in this use case is to auto-tune PID (Proportional-Integral-Derivative) controllers at the edge level in a textiles production site, using data streamed from different edge devices, using ML algorithms. This use case is constructed by Enforma (a data analytics company) and Eliar (an industrial electronics hardware and

service company), which have access to numerous pilot sites among Eliar's installation base in Turkey and more in EU & other countries.

The textile dyeing process is a batch process which takes 5-12 hours depending on various process parameters such as fabric to be dyed, desired colour, chemicals and dye. Eliar produces textile machine process control devices and PLCs (Programmable Logic Controllers) which control the machines according to the desired recipe and process steps. In the textile dyeing process, one of the important criteria that will ensure "right first time" is the correct temperature control of the machine. Currently, PID parameters are tuned by technicians according to their personal experience during installation of the dyeing machine. This may cause inconsistency in the process control and sustainability issues due to inefficient use of resources such as energy, steam, water, chemical, dye, and time.

Through MIRAI, the goal is to tune PID parameters adaptively with the output of the AI algorithm working on the process controllers and PLCs, which are IoT devices operating at the edge. The impact of the project can be described as follows:

- Resource-efficient deployment will be provided by AI algorithms.
- AI algorithms will be able to exchange data and information with each other.
- Each dyeing process is a batch process with unique definition. AI algorithms will be adapted to different processes.
- Noisy and uncertain data from a machine operating in an industrial environment will be studied.
- The installation of the machines will be easier, faster and guarantee trust.
- Results of these algorithms will be stored in the central database for data analysis on factory basis.
- A typical textile dyeing factory has 30 dyeing machines and control devices. Generally, all machines run simultaneously. The main tank liquid temperature is controlled by using PID, taking about 70% of the duration of the dyeing process.
- During the process, an average of 400 different data are collected. These data are various process values such as analogue input and output values, digital input and output values, alarms, operator interventions, running commands. In a typically slow dyeing process, it suffices to exchange data with PLC at 10 Hz.
- PID parameters will be set automatically during installation of the machine. After installation, the system is planned to continue learning and tune the PID parameter values when necessary.
- Daily production amount is around 30 tons in a 30-dyeing machine factory.

Current AI shortcomings. Continuous auto configuration of industrial controllers at the edge makes another challenging business use case. This case is introduced by Eliar (an industrial electronics company) and Enforma (a data analytics company). The main challenge is to provide a continuous stream of data to local (i.e. edge) industrial controllers that can then auto-tune their PID (proportional, integral, derivative) controller parameters. This tuning is currently done at the outset by the human operator based on personal know-how. The tuning parameters are later updated at coarse intervals again by a human operator. All this is despite of factory-wide data is collected

at a central server at the production facility. While auto-tuning the controller parameters, an important challenge will be to ensure that some supervisory aspects are included so that the edge devices have a more general view and therefore they do not end up configured in local optimums.

Distributed AI expectations. With successful application of distributed AI algorithms at edge controllers that continuously update themselves, by also accounting for constraints such as insufficient steam sources and changing priorities of goods across neighbouring devices, we expect to have improved results.

The sheer volume of energy, chemicals and textiles involved in the process translates into considerable improvements considering a typical factory with ~30 dyeing machines working, together producing 30 tons of textiles. Heating takes up about 70% of the overall duration during dyeing. Also accounting for the amount of time, dye, other chemicals, water and energy consumed during this process, it becomes clear why the whole industry is in search of getting things "right first time".

2.6 A first step towards a common reference architecture model

Table 1. Consolidation of the high-level requirements of the five MIRAI use cases.

Requirement	**UC1**	**UC2**	**UC3**	**UC4**	**UC5**
Support real-time monitoring of edge devices	x	x		x	
Support remote control of edge devices	x			x	
Provide processing/learning capabilities at the edge	x	x	x	x	x
Leverage other edge devices (horizontal scaling)			x		x
Ability to learn even with noisy and missing data			x		x
Ability to handle private data and ensure privacy	x		x	x	
Provide fast response and avoid negative consequences	x	x	x	x	x
Reduce communications to the cloud	x		x	x	
Ensure optimal usage of other (natural) resources	x			x	x

Table 1 presents a first consolidation of the high-level requirements distilled from the MIRAI use cases, as described earlier. We observe, on the one hand, that both providing processing and/or learning capabilities at the edge and providing a fast response to avoid negative consequences are concerns shared by all use cases. On the other hand, each use case requires different functionalities (e.g. horizontal scaling, dealing with noisy and missing data, handling private data) specific to the application domain and context that it targets. Overall, the MIRAI project provides a rich and complementary set of five use cases, each within its own industrial context, covering a broad and diverse range of capabilities necessary to build distributed AI applications. These capabilities form the starting point of the MFBB, i.e. the different components that will enable the realization of distributed AI applications in a resource efficient, fully configurable and

composable way. In this design, we intend to bring in concepts from the EECC RAMEC[2] and the IEC 61499[3] models (Figure 1), which address the design and use of function blocks for industrial process, measurement and control systems.

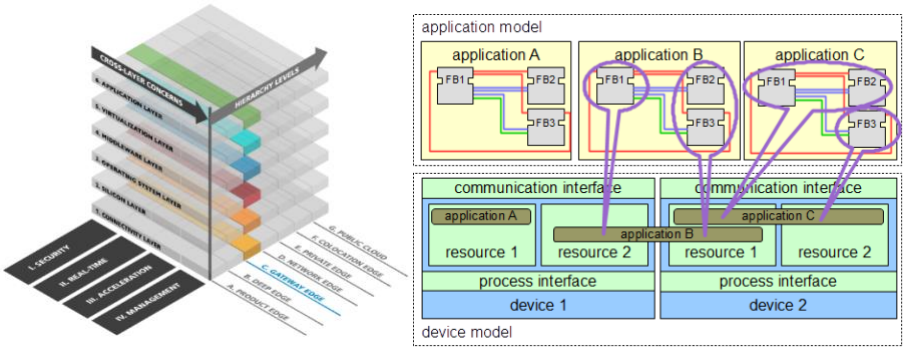

Figure 1. European Edge Computing Consortium's Reference Architecture Model (EECC RAMEC) (left) and International Electrotechnical Commission's International Standard for Distributed Systems (IEC 61499) Function Blocks (right).

3 Conclusions

This paper presents a first attempt to understand the need for migrating from conventional, cloud-based AI to distributed AI applications, hence enabling an innovative horizontal scaling across edge devices, on top of existing vertical scaling to the cloud. Thanks to the five industrial cases, a common argument has been made as to the shortcomings of the existing paradigms and the need to test out distributed AI methodologies. Through the ITEA3 MIRAI project, we plan to illustrate our points by devising, implementing and validating a framework of distributed AI building blocks, together comprising the MFBB, across a set of industrial use case scenarios.

Acknowledgements. The work in this paper is based on MIRAI, a project labelled by ITEA3 under project no 19034, with funding support from Agência Nacional de Inovação in Portugal, Innoviris in Belgium, and TÜBİTAK in Turkey.

Bibliography

1. S. Kewei, T. A. Yang, W. Wei and S. Davari, "A survey of edge computing-based designs for IoT security," *Digital Communications and Networks*, pp. 195-202, 2020.
2. M. A. Talib, S. Majzoub and D. Jamal, "A systematic literature review on hardware implementation of artificial intelligence algorithms," *The Journal of Supercomputing*, pp. 1897-1938, 2020.

[2] https://ecconsortium.eu/
[3] https://www.iec61499.de/

3. J. Lines, A. Bagnall, P. Caiger-Smith and S. Anderson, “Classification of household devices by electricity usage profiles,” in *International conference on intelligent data engineering and automated learning*, 2011.
4. C. Beckel, L. Sadamori and S. Santini, “Automatic socio-economic classification of households using electricity consumption data,” in *Proceedings of the fourth international conference on Future energy systems*, 2013.
5. C. W. Potter, A. Archambault and K. Westrick, “Building a smarter smart grid through better renewable energy information,” in *2009 IEEE/PES Power Systems Conference and Exposition*, 2009.
6. W. Kong, Z. Y. Dong, Y. Jia, D. J. Hill, Y. Xu and Y. Zhang, “Short-term residential load forecasting based on LSTM recurrent neural network,” *IEEE Transactions on Smart Grid*, vol. 10, no. 1, pp. 841--851, 2017.
7. R. Bessa, A. Trindade, C. S. Silva and V. Miranda, “Probabilistic solar power forecasting in smart grids using distributed information,” *International Journal of Electrical Power & Energy Systems*, vol. 72, pp. 16-23, 2015.
8. N. Apthorpe, D. Reisman, N. Feamster, “A smart home is no castle: Privacy vulnerabilities of encrypted iot traffic,” *Workshop on Data and Algorithmic Transparency*, 2016.
9. J. A. Pérez-Díaz, I. Amezcua Valdovinos, K.-K. R. Choo; D. Zhu, “A Flexible SDN-Based Architecture for Identifying and Mitigating Low-Rate DDoS Attacks Using Machine Learning,” *IEEE Access*, vol. 8, pp. 155859-155872, 2020.
10. R. Doshi, N. Apthorpe and N. Feamster, “Machine Learning DDoS Detection for Consumer Internet of Things Devices,” *IEEE Security and Privacy Workshops (SPW)*, pp. 29-35, May 2018.
11. M. Eskandari, Z. H. Janjua, M. Vecchio and F. Antonelli, “Passban IDS: An Intelligent Anomaly-Based Intrusion Detection System for IoT Edge Devices,” *IEEE Internet of Things Journal*, vol. 7, no. 8, pp. 6882-6897, August 2020.
12. S. Maurya and A. Choudhary, “Deep learning based vulnerable road user detection and collision avoidance,” 2018.
13. J. Schleusner, L. Neu and H. Blume, “Deep Learning Based Classification of Pedestrian Vulnerability Trained on Synthetic Datasets,” in *2019 IEEE 9th International Conference on Consumer Electronics (ICCE-Berlin)*, 2019.
14. M. Nitinder and J. Kangasharju, “Edge-Fog cloud: A distributed cloud for Internet of Things computations,” in *2016 Cloudification of the Internet of Things (CIoT)*, 2016.
15. F. Forooghifar, A. Aminifar and D. Atienza, “Resource-aware distributed epilepsy monitoring using self-awareness from edge to cloud,” *IEEE transactions on biomedical circuits and systems*, vol. 13, no. 6, pp. 1338-1350, 2019.
16. Y. Liu, X. Ma, Y. Li, Y. Tie, Y. Zhang and J. Gao, “Water pipeline leakage detection based on machine learning and wireless sensor networks,” *Sensors*, vol. 19, no. 23, p. 5086, 2019.
17. C. de Souza and A. Kalbusch, “Estimation of water consumption in multifamily residential buildings,” *Acta Scientiarum. Technology*, vol. 39, no. 2, pp. 161-168, 2017.
18. R. M. Almeida, N. M. Ramos, L. M. Simões and V. P. de Freitas, “Energy and water consumption variability in school buildings: review and application of clustering techniques,” *Journal of Performance of Constructed Facilities*, vol. 29, no. 6, p. 04014165, 2015.