

Enhancing Safety and Privacy in Industry 4.0: The ICE Laboratory Case Study

Original

Enhancing Safety and Privacy in Industry 4.0: The ICE Laboratory Case Study / Cunico, Federico; Aldegheri, Stefano; Avogaro, Andrea; Boldo, Michele; Bombieri, Nicola; Capogrosso, Luigi; Caputo, Ariel; Carra, Damiano; Centomo, Stefano; Cheng, Dong Seon; Cinquetti, Ettore; Cristani, Marco; Marchi, Mirco De; Demrozi, Florenc; Emporio, Marco; Fummi, Franco; Geretti, Luca; Germiniani, Samuele; Giachetti, Andrea; Girella, Federico; Martini, Enrico; Menegaz, Gloria; Muijs, Niek; Paci, Federica; Panato, Marco; Pravadelli, Graziano; Quintarelli, Elisa; Siviero, Ilaria; Storti, Silvia Francesca; Tadiello, Carlo; Turetta, Cristian; Villa, Tiziano; Zannone, Nicola; Quaglia, Davide. - In: IEEE ACCESS. - ISSN 2169-3536. - ELETTRONICO. - 12:(2024), pp. 154570-154599. [10.1109/access.2024.3479411]
This version is available at: 11583/2994055 since: 2024-10-31 17:34:56Z

Publisher:

Institute of Electrical and Electronics Engineers (IEEE)

Published

DOI:10.1109/access.2024.3479411

Terms of use:

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

Publisher copyright

(Article begins on next page)

RESEARCH ARTICLE

Enhancing Safety and Privacy in Industry 4.0: The ICE Laboratory Case Study

FEDERICO CUNICO¹, (Graduate Student Member, IEEE), STEFANO ALDEGHERI², ANDREA AVOGARO¹, (Associate Member, IEEE), MICHELE BOLDO¹, (Student Member, IEEE), NICOLA BOMBIERI², (Member, IEEE), LUIGI CAPOGROSSO¹, (Student Member, IEEE), ARIEL CAPUTO², DAMIANO CARRA¹, STEFANO CENTOMO¹, DONG SEON CHENG², (Member, IEEE), ETTORE CINQUETTI¹, MARCO CRISTANI², (Member, IEEE), MIRCO DE MARCHI¹, FLORENC DEMROZI³, (Member, IEEE), MARCO EMPORIO¹, (Student Member, IEEE), FRANCO FUMMI², (Member, IEEE), LUCA GERETTI¹, SAMUELE GERMINIANI², (Member, IEEE), ANDREA GIACHETTI², FEDERICO GIRELLA¹, ENRICO MARTINI¹, (Member, IEEE), GLORIA MENEGAZ², (Senior Member, IEEE), NIEK MUIJS⁴, FEDERICA PACI¹, MARCO PANATO², GRAZIANO PRAVADELLI², (Senior Member, IEEE), ELISA QUINTARELLI¹, ILARIA SIVIERO¹, (Member, IEEE), SILVIA FRANCESCA STORTI², (Member, IEEE), CARLO TADIELLO², CRISTIAN TURETTA¹, (Graduate Student Member, IEEE), TIZIANO VILLA¹, (Member, IEEE), NICOLA ZANNONE⁴, AND DAVIDE QUAGLIA¹

¹Department of Computer Science, University of Verona, 37129 Verona, Italy

²Department of Engineering for Innovation Medicine, University of Verona, 37129 Verona, Italy

³Department of Electrical Engineering and Computer Science, University of Stavanger, 4021 Stavanger, Norway

⁴Department of Mathematics and Computer Science, Eindhoven University of Technology, 5612 AZ Eindhoven, The Netherlands

Corresponding author: Federico Cunico (federico.cunico@univr.it)

This work was supported in part by the Italian Ministry of University and Research (MUR) through “Dipartimenti di Eccellenza” (2018–2022); in part by the Fondazione Cariverona through “Ricerca&Sviluppo;” and in part by the European Union–NextGenerationEU (Italiadomani, Piano Nazionale di Ripresa e Resilienza (PNRR)–Missione 4–Componente 2, Investimento 3.3–progetto M4C2).

This work involved human subjects or animals in its research. Approval of all ethical and experimental procedures and protocols was granted by the Comitato di Approvazione della Ricerca sulla Persona under Application No. 2.R1/2022.

ABSTRACT The revolutionary technologies behind Industry 4.0 have opened a new era for manufacturing: connected and autonomous machines, collaborative robotics, and monitoring techniques are spreading to increase productivity and sustainability. From the workers’ perspective, they bring new safety threats but also opportunities to solve old ones, while concerns about workers’ privacy arise due to the increase of data sensed and transferred from the shop floor. This paper presents the results of a research project addressing the prediction of dangerous conditions through workplace monitoring with privacy guarantees. This work is driven by a realistic approach starting from the fact that it is entirely centered on a real 14-meter production line equipped with an extensive array of top-tier devices, including robotic arms, autonomous mobile robots, a reconfigurable moving belt, a multi-camera system, and a highly efficient data transport and computation infrastructure. This project shows safety and privacy achievements over six representative use cases such as man-on-the-ground, environmental events (e.g., fire incidents), workers’ errors that can lead to potential accidents, compliance of Personal Protective Equipment (PPE), and gatherings restrictions. The benefits of this study extend to stakeholders such as manufacturers and workers offering safety systems that can be deployed in industrial settings while addressing privacy concerns and providing compliance with regulations.

The associate editor coordinating the review of this manuscript and approving it for publication was Wai-Keung Fung¹.

The industrial laboratory at the heart of this study represents with realism a dynamic and interconnected Industry 4.0 and 5.0 environment.

• **INDEX TERMS** Industry 4.0, occupational safety, privacy, human–robot interaction, industry 5.0.

I. INTRODUCTION

Industry 4.0 is transforming traditional manufacturing by integrating advanced automation, communication, and interaction technologies within production plants. The most striking example of this revolution is *smart factories*, in which employees, machines, processes, and products connect to give a more effective organization of all the productive means to achieve higher levels of efficiency and flexibility [1], [2]. Interconnection between devices is possible thanks to the Industrial Internet of Things (IIoT), which allows industrial objects to be sensed and controlled remotely across existing network infrastructure, creating opportunities for more direct integration of the physical world into computer-based systems and resulting in improved efficiency and accuracy [3], [4].

This Fourth Industrial Revolution heavily relies on cyber-physical systems (CPS) to make production architectures modular and adaptable. IIoT and CPS enable the synergy of many technologies, including advanced robotics, machine learning, artificial intelligence, and nanotechnology.

Industry 4.0 enables human-robot collaboration scenarios, even if they are not the focus of the 4.0 revolution, allowing for human-machine challenges to arise and begin to be addressed [5], [6]. Its evolution, the Industry 5.0 framework born from the advancement of the Industry 4.0 revolution, is much more oriented to addressing human-robot mixed environment challenges. Under both frameworks, it is not uncommon for smart factories to have human-in-the-loop scenarios, i.e., mixed environments in which humans and robots are present. This kind of scenario is known as *human-in-the-loop* [5], [6]. Therefore, Industry 4.0 is addressing the challenge of enhancing human-machine integration in manufacturing systems and augmenting workers' physical and cognitive capabilities.

Such integration brings inevitable safety concerns, particularly when considering moving parts and robots operating in a mixed environment. These concerns must be accounted for to guarantee the safety of workers and whoever interacts with the industrial environment. Safety concerns include possible harm by some moving robots or moving parts, e.g., manipulators, originating from their poor awareness of human behavior, but they can also occur in the case of workers' mistakes or wrong actions [7].

Assessing occupational safety means not only reacting to dangerous situations but also preventing possible adverse outcomes by forecasting the evolution of the human-environment system. For instance, predicting a future collision between a human and a moving robot may prevent potential injuries. To this purpose, in recent years, there has been a significant effort to leverage data analytics to assist

in decision-making that directly improves safety performance [8], [9].

However, the more specific data-driven technologies are employed to address safety concerns, the more privacy concerns emerge. Trying to compensate for occupational safety issues leads to the necessity of monitoring humans in the environment and collecting data that can be identified as personal information, hence creating privacy issues.

This work presents a case study featuring the creation of an Industry 4.0 and 5.0 laboratory with an emphasis on workers' safety and privacy. Our lab, the Industrial Computer Engineering (ICE) Laboratory¹ was designed to develop cutting-edge IT solutions and methodologies in Industry 4.0 and 5.0 concerning automation and efficiency of production lines. The goals of the laboratory are research, technology transfer, and teaching.

This project aims to show a case study for enhancing safety in an industrial environment, considering the privacy issues that inevitably take a role in the process. The technologies employed to improve safety will be presented, along with how privacy issues could be addressed using a context-aware manager that handles sensitive data from various technologies.

For this reason, six *use cases* are considered, representing possible threats to workers' safety and how the techniques developed by the authors can address them. In Table 1, the rows list these use cases which are described in Section III and their mapping to the used techniques that are listed in the columns and described in Section IV.

The paper is organized as follows. Section II presents the ICE Laboratory with its equipment and infrastructures. Section III presents the use cases considered with respect to the safety of workers. Section IV shows how the use cases are addressed from a scientific point-of-view, presenting the algorithms and solutions adopted with the results achieved. Section V presents the context-aware solution implemented to allow the technologies to be safety-enhancing while keeping sensitive data under control, providing privacy to the whole system. Finally, Section VI presents the conclusions of this study and the future works for the laboratory.

II. THE ICE LABORATORY

In 2018, the Italian Ministry of Education, Universities, and Research (MIUR) awarded the Department of Computer Science of the University of Verona a grant for being a Department of Excellence. One key aspect of the awarded project was the establishment of the Industrial Computer Engineering (ICE) Laboratory for Industry 4.0 with a modern production line, extended with equipment for augmented

¹<https://www.icelab.di.univr.it/>

TABLE 1. The use cases for privacy and safety considered in the Industry 4.0 scenario of the ICE Lab. For each use case, there are several technologies applied to take into account the potential threat.

Technology \ Use Case	Infrastructure for data transport and computation	Visual tracking	Pose estimation	Wearables	XR and 3D interaction	Data-driven safety assessment	BCI	Context-aware Privacy Management
Adverse event	✓	✓	✓	✓	✓	✓		✓
Mixed human-machine environment	✓	✓	✓	✓	✓	✓	✓	✓
Wrong action	✓	✓			✓		✓	✓
Environmental event	✓	✓		✓		✓		✓
PPE compliance	✓	✓	✓					✓
Gathering	✓	✓						✓

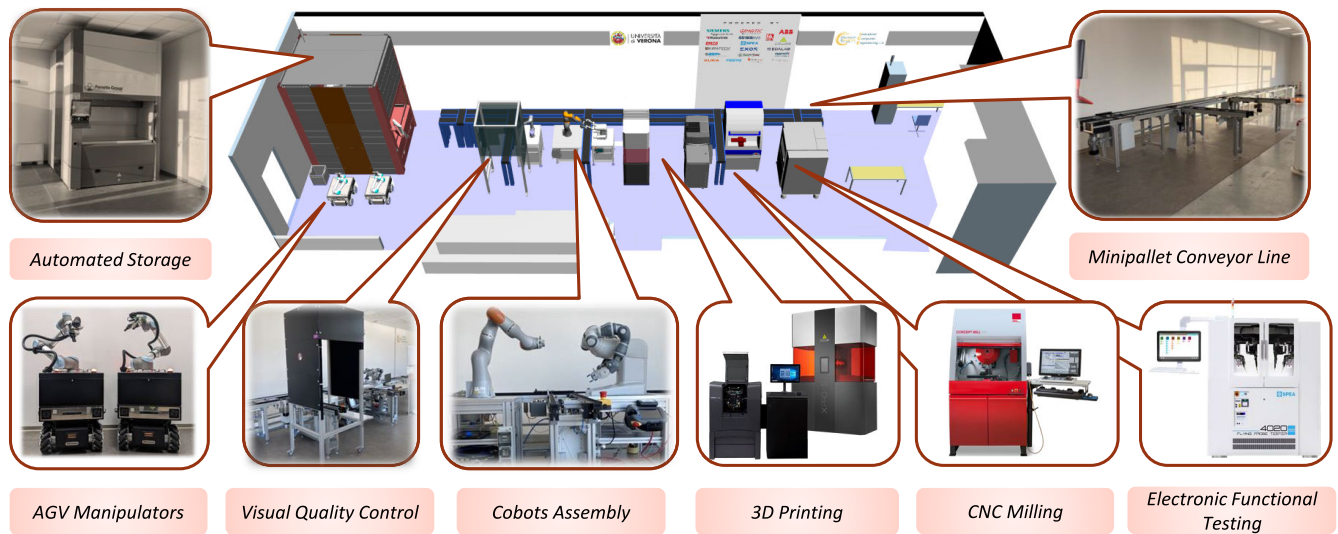


FIGURE 1. The ICE laboratory at the university of verona.

reality and digital production, and connected with the university’s high-performance computational platform.

The laboratory was built with three objectives in mind: education, research, and technological transfer. The lab has been used by the research groups of the affiliated departments for courses and research projects in areas like robotics, cybernetics, factory automation, predictive maintenance, workplace safety and privacy, digital twin technology, mixed reality, human tracking, pose estimation, and machine learning. It has also been used to incubate and showcase technologies being developed for future commercial purposes.

The facility serves as a technology demonstrator for a comprehensive set of computational technologies applied to the industrial manufacturing field; as depicted in Figure 1, the centerpiece of the laboratory is an entirely reconfigurable production line equipped with an assortment of heterogeneous devices grouped into manufacturing cells providing smart storage, mobile as well as conveyor-based

advanced logistics, visual quality control, robotic assembly, 3D printing, subtractive operations, and electronic functional test.

Every cell contains commercially available, and industrial-grade machines. This was an intentional design requirement to create a realistic, albeit synthetically compact, production environment, unlike a research prototype far removed from a typical industrial floor. The majority of the initial work in the lab required unifying communication protocols, automating each device’s operations, integrating physical and software aspects of the devices. This has resulted in the development of a unique service-oriented software architecture that handles communication and control of the machines in order to facilitate modularity and reconfigurability in production, which are key aspects of Industry 4.0.

The following describes the hardware equipment, computational infrastructures, and design specifications of the ICE laboratory. A comprehensive list of all ICE laboratory equipment is presented in Table 2 where the column “Work

Cell” refers to the operative sectors depicted in Figure 1 and, in case of “technology” attribute, to the equipment specifically used by the technologies of Section IV.

A. MANUFACTURING HARDWARE

The laboratory is a small-scale version of an advanced manufacturing environment and includes several pieces of equipment grouped into representative technological areas. Each of these manufacturing cells, and sometimes individual equipment, is monitored by energy meters to better understand correlations between usage and power consumption.

1) AUTOMATED STORAGE

Raw materials and product components are stored within the *Ferretto Group VERTIMAG EF65XS* vertical automated storage, with the Ergo-Tech patented system that allows incoming and outgoing trays to move simultaneously. The storage contains 10 trays of varying height and payload capacity (ranging from 300 to 990 Kg) and has a touch panel screen with a user interface for monitoring and control; the picking bay always extracts the trays at the same height for dependable access; there are safety size checks for incoming goods and safety barriers for the operators. The accompanying software organizes and updates the content of the trays and allows remote control.

2) ADVANCED LOGISTICS

The spine of the production environment is the *Bosch-Rexroth TS 2plus* conveyor system with RFID-enabled mini-pallets that circulate in a ring and can be directed to several side bays for loading/unloading and processing. The pallets can be flexibly re-directed to any position according to the current needs, and the system is compliant with manual repositioning. The side bays steer components to the manufacturing cells along the production line, but, unlike most factory sequential setups, it allows to connect any cell with any other cell without pre-configuration.

Augmenting the conveyor system is a fleet of mobile robots in the form of two *Robotnik RB-KAIROS 5* autonomous ground vehicles (AGVs) equipped with UR5 manipulators. These robots navigate the open spaces between the other equipment and bridge missing links in the line, e.g. from the warehouse to the conveyor. They are managed by the robotic operating system (ROS) and are thus open for quickly developing new adaptations required by production recipes.

3) ROBOTIC ASSEMBLY

The robotic assembly cell is a collaborative working area endowed with two robots, the dual-arm high-precision *ABB YuMi IRB14000* and the heavy payload *Kuka LBR iiwa 14 R820*, controlled using a *Siemens S7-1500* safety PLC.

Each arm of the ABB robot has 6DoF and is equipped with grippers and a suction tool for collecting small parts. The right arm also has a camera to facilitate operations requiring visual assistance. These characteristics make it an ideal solution

for fast and accurate operations on small objects in confined spaces that also enable cooperation with humans and other robotic operators.

The Kuka is a 7DoF robot arm equipped with dual torque sensors on each joint, accurately detecting contacts and collisions to instantly reduce the level of force and speed exerted. The robot arm is designed to support up to 14 kg payloads while maintaining excellent positioning accuracy. On its end, depending on the production recipes, there is either a *Gimatic MPLM3240P* gripper or a *Fiam eTensil* electric tightening device as its tool.

The PLC device communicates with the two robots using digital signals, is equipped with safety functions, and also handles area safety monitoring. It supports the OPC-UA communication protocol and can be directly integrated into the laboratory’s software infrastructure.

4) ADDITIVE MANUFACTURING

Regarding additive synthesis, the laboratory includes several 3D printing technologies: filament, resin, and multi-material. In the first group, there are two *Prusa i3 MK3S+* (one equipped with the MMU2) for fast and cheap prototyping.

In addition, the *DWS Systems X PRO S* is a stereolithographic printer that uses a laser source to solidify a photosensitive liquid resin to produce three-dimensional plastic models. The machine uses a class 3b UV laser, which, employing a moving axis and a system of oscillating mirrors, enables the creation of three-dimensional models.

Finally, the *Stratasys J826* features multi-material additive technology. It can create objects by composing up to eight materials with different physical (opacity and transparency) and mechanical (flexibility and rigidity) properties. The printer also supports more than 500,000 colors, eliminating the need to paint models. The process’s color quality is optimized thanks to Pantone validation, which makes PMS (Pantone Matching System) colors available for 3D printing.

5) SUBTRACTIVE MANUFACTURING

The laboratory has paired an *EMCO ConceptMill 105* CNC milling cutter with a UR5 robotic arm for loading/unloading pieces into/from the machine’s enclosure. This subtractive manufacturing cell allows the production recipes to include, for example, custom engraving operations on surfaces of interest in the products. Additional features of the milling machine are an engraving spindle with a grooved guide, a continuously adjustable central unit, a CNC divider as an optional fourth axis, and a robotic interface for integration in FMS / CIM systems.

6) ELECTRONIC FUNCTIONAL TEST

The laboratory has installed a *SPEA 4020S2* automatic flying probe tester to check the quality of electronic boards. It has four flying probes that provide highly accurate and fast contacting even on the smallest SMD components and a set of modules to allow a wide range of operations, from in-circuit test, short circuit, open pin, power on, functional test

and on-board programming. The machine is fully automatic, with a closed enclosure around the test area, and a conveyor belt to safely load/unload boards. Its role in the laboratory production line is to check the quality of the electronic boards and to flash the firmware into their MCUs. The laboratory has installed a custom-tailored OPC-UA server to integrate this machine into its infrastructure.

7) VISUAL QUALITY CONTROL

Product quality control is performed using a camera with *EV76C570 e2v* CMOS sensor and an edge computer equipped with a high-performance Nvidia GPU and a *Basler acA1600-60gm GigE*. The camera is placed in a darkened cabin to avoid light disturbances located in a bay of the production line.

Inside the cabin, there is also a *Gocator 3210* laser scanner, used for full-field three-dimensional inspection of large components using snapshot scanning functionality. The laser scanner has a scanning speed of 6 Hz with acceleration. The field of view extends to 100×154 mm. Parameters summarising the integrity of the sensor can be monitored on the user interface dashboard.

To expose all sides of the product to the cameras' view in order to complete the visual analysis, an UR5e collaborative robot is positioned at the camera system and is coordinated by the edge computer.

8) IoT NETWORK

To support wireless and wearable devices that don't rely on WiFi, the laboratory employs the *EDALAB BOX-IO IoT Platform*. Examples of devices supported with this platform are environmental sensors that track temperature, humidity, CO_2 and light levels. This allows us to trace comfort levels within the rooms.

The BOX-IO gateway also has the capability of working with many communication technologies (e.g., Modbus, LoRaWAN, Zigbee, WiFi, 4G LTE), and it can serve as a WiFi Access Point (AP). Thus, the BOX-IO can act as a "translator", enabling communication between devices using different interfaces. The gateway itself is visible to the rest of the infrastructure as an OPC-UA server.

B. INFRASTRUCTURE FOR DATA COLLECTION AND DISTRIBUTION

As important as the collection of manufacturing equipment is the infrastructure that allows advanced communication and control. The following section describes the networking equipment, the local processing cluster with the Kafka-based data collection and distribution platform (see Figure 2), and the line management system.

1) NETWORK INFRASTRUCTURE

The manufacturing equipment detailed in the previous section is connected to the local data center using industrial Cat7 Ethernet cables. The wired network allows huge and efficient data transfers at 1 Gb/s. Kafka, a distributed software for

event handling, is used to enable concurrent communication between various systems, including robotic cells, conveyor belts, moving robots, camera systems, and wearable devices. The messages, delivered with Kafka, are JSON-structured files, making them a good compromise between human-readability (and therefore maintainability) and efficiency.

The data center itself houses the main network infrastructure composed of two *Siemens SCALANCE SC636-2C* firewalls and four *Siemens SCALANCE XR324-4M PoE* managed Ethernet switches. Each manufacturing cell is isolated in its own virtual LAN for safety and security reasons and access from outside is granted using a local *Siemens SINEMA Remote Connect* VPN server.

In addition to the wired infrastructure, there is also an industrial *Siemens SCALANCE W1750D* access point to provide WiFi networks for the mobile devices, including the two AGVs.

2) DATA INTEGRATION HUB

Developed completely internally, the Data Integration Hub (DIH) is the custom platform that monitors the equipment, stores the gathered data, and propagates to the machines the commands requested by the line management system. It acts as a service platform that provides a variety of functionalities: machine monitoring and control through OPC-UA servers and clients, communication between Kubernetes nodes with Industrial Internet of Things (IIoT) brokers (RabbitMQ, Kafka, MQTT), storage and retrieval of time series (InfluxDB). Kafka in particular (see Figure 2), is used to provide high-throughput, low-latency channels for handling real-time data streams from the machines to the processing applications, and vice-versa.

3) LINE MANAGEMENT SYSTEM

Production management is handled at the highest level by the *Siemens OPCenter Execution Discrete* manufacturing execution system (MES) running on a virtual machine in the HPC. To ensure the production activities are correctly set up, tracked, and managed on the equipment, the laboratory has developed an intermediate platform, the Automation Manager [10], that implements the required production recipes.

Furthermore, there is a Digital Twin [11] for the production line developed using the *Tecnomatix Plant Simulation* software. Digital Twin is the virtual representation of a process, product, or service. In this case, it is the model of the production line that allows both data analysis and system monitoring. Specifically, there are two Digital Twins: *Connected* and *Autonomous*.

The *Connected Digital Twin*, called Digital Shadow, communicates with the machines' OPC-UA servers to monitor their status and visually replicate it in real-time. It connects to the line with sensors and replicates the physical plant, allowing real-time control over the entire system.

The *Autonomous Digital Twin* is asynchronous and is designed to simulate the plant as a whole. In particular,

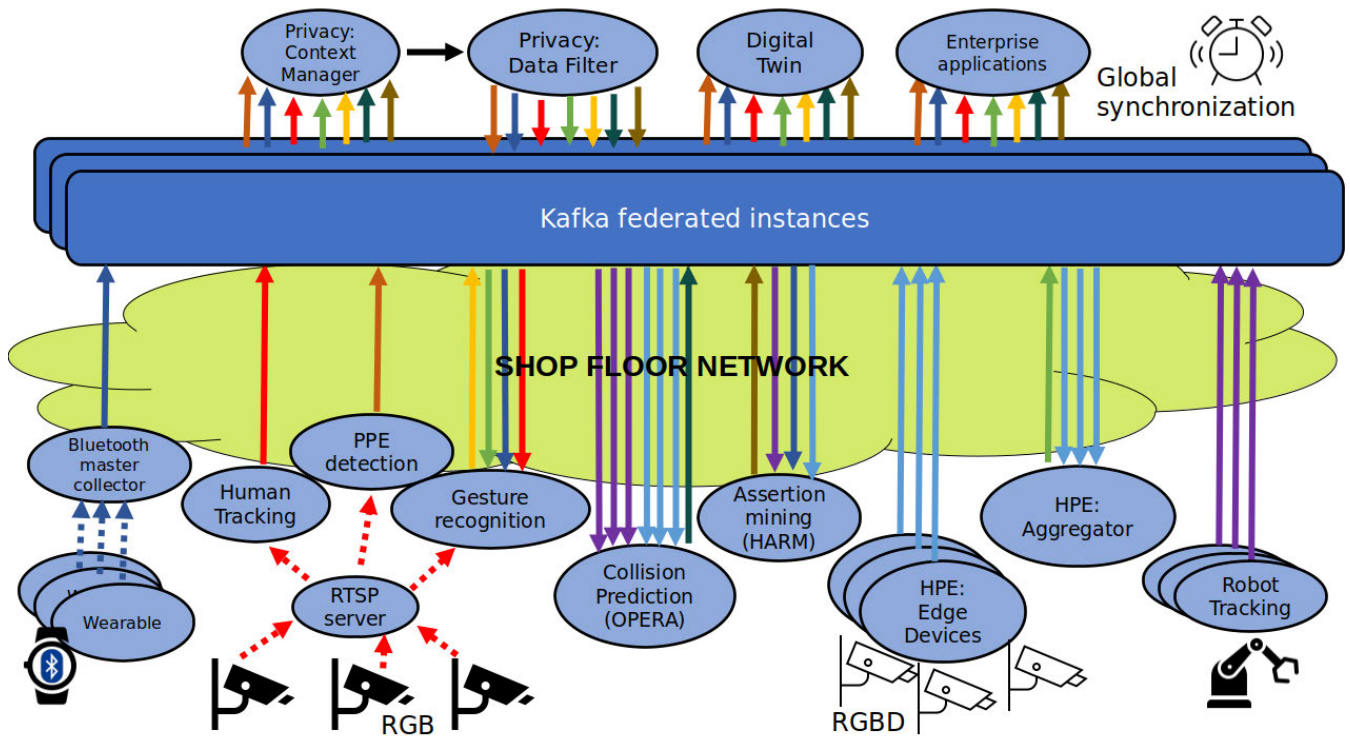


FIGURE 2. Kafka-based infrastructure for data collection and distribution.

it incorporates accurate machines’ models, replicating their functionality and can be used to test machinery that could be integrated in the future. It is useful to obtain real-time statistics, to study alternative configurations of the production line, and optimize its timing; moreover, through data analysis, it is possible to carry out predictive maintenance, detecting a possible failure before it occurs.

4) DATA CENTERS

Our local processing power is provided by a *Rancher Kubernetes Engine* (RKE) cluster running on nine Ubuntu 22.04 LTS nodes. RKE allows easy installation and operation of the Kubernetes cluster and supports all the major technologies for data collection, distribution, and storage. Every application described below relies on a collection of Kubernetes nodes that define and manage containerized services. This local cluster, named *ice-edge*, is complemented by a cloud-based cluster, named *ice-cloud*, hosted on the University’s High Performance Computing (HPC) platform. While the Data Integration Hub and the Automation Manager reside on the *ice-edge*, the various MES and the management tools reside on the *ice-cloud* where higher quality resources can be requested as needed and provide further scalability for future expansions.

5) COMPUTATION OFFLOAD

Some technologies used in the ICE laboratory are based on deep learning algorithms which cannot be executed

completely on edge devices because of their heavy computational requirements, whereas current optimization techniques (e.g., simplifying models, compression, distillation, etc.) may reduce the inference quality [12].

Currently, a common practice is to transfer sensed data to servers with high computing capabilities through the network. This approach preserves full accuracy but increases communication traffic and end-to-end latency, and it exposes privacy concerns.

One alternative implemented in this project is Split Computing (SC), a technique that intelligently divides a deep neural network between an edge device and a remote server. This approach improves privacy and reduces latency as well as required transmission bandwidth since the tensors are compressed through the use of an autoencoder.

Thanks to the contribution presented in [13], the prevalent methods for identifying potential splitting points have evolved from architecture-based techniques to more refined neuron-based methods. The method, I-SPLIT, demonstrates that both the architecture of the layers and the saliency of individual layers are crucial factors. Specifically, a neuron’s saliency is determined by its gradient in relation to the accurate decision.

Then, the *Split-Et-Impera* framework [14] is proposed as a fast and user-friendly solution that simplifies the design of a distributed architecture for executing one or more DNNs. Other than accurately mimicking diverse communication protocols and application requirements,

TABLE 2. Summary of all the equipment present in the laboratory, with the model, quantity, operative sector of the laboratory (work cell), and type of equipment. The “technology” attribute in the work cell column refers to the equipment specifically used in the project as explained in Sec. IV.

Equipment	Model	Quantity	Work Cell	Type
Vertical Automated Storage	<i>Ferretto Group VERTIMAG EF65XS</i>	1	Automated Storage	Manufacturing Hardware
Conveyor System	<i>Bosch-Rexroth TS 2plus</i>	1	Advanced Logistics	Manufacturing Hardware
AGV	<i>Robotik RB-KAIROS 5</i>	2	Advanced Logistics	Manufacturing Hardware
UR5 cobot for logistics	<i>Universal Robot UR5</i>	2	Advanced Logistics	Cobot
Dual-arm high precision cobot	<i>ABB YuMi IRB14000</i>	1	Robotic Assembly	Manufacturing Hardware
Heavy load cobot	<i>Kuka KBR uiiwa 14 R820</i>	1	Robotic Assembly	Cobot
Filament 3D printer	<i>Prusa i3 MK3S+</i>	1	Additive Manufacturing	3D Printer
Stereolithographic 3D printer	<i>DWS Systems X PRO S</i>	1	Additive Manufacturing	3D Printer
Multi-material 3D printer	<i>Stratasys J826</i>	1	Additive Manufacturing	3D Printer
Milling cutter	<i>EMCO ConceptMill 105 CNC</i>	1	Subtractive Manufacturing	Manufacturing Hardware
UR5 cobot for milling support	<i>Universal Robot UR5</i>	1	Subtractive Manufacturing	Cobot
Automatic flying probe tester	<i>SPEA 4020S2</i>	1	Electronic Functional Testing	Manufacturing Hardware
Industrial-grade camera	<i>Basler acA1600-60gm GigE</i>	1	Visual Quality Control	Manufacturing Hardware
Laser scanner	<i>Gocator 3210</i>	1	Visual Quality Control	Manufacturing Hardware
Gateway for wireless devices	<i>EDALAB BOX-IO IoT</i>	1	IoT Network	Network
Surveillance cameras	<i>Bosch DINION IP 3000i IR</i>	7	Technology	Camera
PTZ Surveillance camera	<i>Bosch AUTODOME IP 5100i</i>	1	Technology	Camera
RGB-D cameras	<i>Zed 2</i>	7	Technology	Camera
VR Headset	<i>Meta Quest 2</i>	1	Technology	AR/VR
AR Headset	<i>Microsoft Hololens 2</i>	1	Technology	AR/VR
Wearable device	<i>Nordic Thingy52</i>	4	Technology	Wearable
BCI Headset	<i>EBNeuro BeMicro</i>	1	Technology	Human-Machine Interaction

Split-Et-Impera introduces a unique feature: it suggests the proper configuration to match the application’s quality of service (QoS) requirements and provide optimal performance in terms of accuracy and latency time. Since manipulating diverse SC configurations may require days of computation, *Split-Et-Impera* allows eliminating several configurations through communication-aware simulations.

At the same time, current state-of-the-art approaches in different deep learning applications rely on advanced learning procedures, such as Multi-Task Learning (MTL). In particular, MTL is a paradigm in which multiple related tasks are jointly learned to improve the generalizability of a model by using shared knowledge across different aspects of the input. As a result, in [15], a method is presented to partition multi-tasking DNN to be deployed within an SC framework. The proposed *MTL-Split* design can handle multiple tasks concurrently instead of the current focus on Single-Task Learning (STL) in SC, and through MTL, increase task performance, overcoming the challenge of preserving only the performance of the main task.

Moreover, in [16] is presented the effect of predefined sparsity within the SC paradigm. This approach, demonstrably effective for the first time in an SC scenario, allows us to further reduce the computational, storage and energy demands significantly during training and inference, regardless of the hardware platform.

Similarly, the work presented in [17] introduces a novel hardware-aware split-computing platform, targeting real-time human pose estimation. The goal of this platform is to consider the bandwidth of the transmission network and the load of the edge board hardware, to ensure a minimum working frequency. The platform, through a micro-benchmarking phase, extracts the latency and the output size of each layer of the neural network. To minimize traffic on the network, transmitted data are quantized to 4 bits and compressed with lossless encoding, which allows exploiting the large number of “0” numbers generated by *ReLU* activation functions. Thanks to a performance monitor on the board and an ad-hoc developed UDP-based network protocol, the platform is able to monitor the instantaneous




FIGURE 3. Camera locations. The yellow triangles are the RGB cameras, and the blue triangles are the RGB-D cameras. The blue dot with the number 9 is the PTZ camera location, and the dot with the number 8 is the fisheye camera.

condition of the infrastructure, going from time to time to change the split point in order to ensure the nonfunctional constraints of the maximum required latency.

C. MULTI-CAMERA SYSTEM

In addition to the manufacturing equipment, the laboratory is equipped with a multi-camera system, with both RGB and RGB-D cameras. There are 7 RGB Bosch DINION IP 3000i IR surveillance cameras, 7 RGB-D ZED-2 cameras, 1 PTZ Bosch AUTODOME IP 5100i camera, and 1 fisheye camera. Figure 4 shows the cameras mounted in the laboratory while the exact position of each camera is reported in Figure 3. The synchronization between the various cameras is maintained using a local NTP server (but any public server could be used). This is an important part of the system since wrong synchronization could drastically affect the performance of any system relying on cameras.

The surveillance cameras and the RGB-D cameras have different purposes. The surveillance cameras are designed to have a good resolution and wide opening to achieve good coverage over the whole laboratory. These cameras aim to demonstrate that a normal video-surveillance system, possibly found in existing industrial plants, can be turned into a powerful Industry 4.0 application by sending video fluxes to a central machine learning system. At the opposite, RGB-D cameras represent the case in which a new setup is needed and privacy concerns are the main issue. In fact, each camera is part of a smarter system with its own processing unit with GPU capabilities (Nvidia Jetson Xavier NX). In this way, the camera is completely privacy-preserving because no video stream is allowed to be sent over the network channel, and the only data transmitted is aggregated anonymized information or metadata. Furthermore, these cameras provide a depth channel, i.e., 3D information.

The use of one or the other set of streams (RGB surveillance or RGB-D) depends on the application and the computation required. Furthermore, the camera subnetwork is isolated and all the data running under this subnetwork is monitored and can only be accessed via access control.

D. WORKERS' EQUIPMENT

Besides the general equipment of ICE both infrastructure and manufacturing-related, each worker can have particular



FIGURE 4. The cameras mounted in the laboratory. On the left: the Zed RGB-D camera with the Jetson board in its enclosure. On the right: the RGB Bosch camera.

equipment depending on the task they have to accomplish. Each type of equipment has its own motivation and specific task that it helps to achieve.

1) XR HEADSETS

HoloLens 2 head-mounted displays (HMD) are available in the ICE Lab to support workers. Such devices allow for augmented reality with spatial understanding thanks to the built-in sensors, good autonomy in terms of battery, and non-tedious portability for long periods of time. This device is used to guide workers during the manufacturing process and the inspection of the final product, and can be used to provide important information for safety and error prevention.

Furthermore, thanks to the HoloLens 2 finger-tracking capabilities, it is possible to implement advanced interfaces based on real-time 3D gesture recognition algorithms [18], [19]. This is particularly important for human-robot collaborative scenarios, to enable faster and more effective collaboration.

VR headsets (Meta Quest 2) are available as well to allow training of workers and remote monitoring and collaboration.

2) WEARABLE DEVICE

The Nordic Thingy52 sensors, whose dimensions are comparable to those of commercial smartwatch devices like Tick Watch or Apple Watch, are worn by workers in the production line and play an essential role in worker identification and safety. Binding the device's unique identifier and the worker's identity allows for the unique identification of each worker, ensuring that individual movements and locations can be tracked with precision. This capability is particularly crucial in environments where worker safety and accountability are paramount. In conjunction with the Human Pose Estimation (HPE) system, the wearable device is used to associate the 3D set of key points representing the human skeleton of the worker with his/her identity. Furthermore, the inertial data provided by the wearable sensor allows fall detection. By leveraging the underlying infrastructure of ICE Lab,

immediate alerts can be sent to supervisors or emergency response teams, ensuring rapid intervention and minimizing potential injuries.

3) BRAIN COMPUTER INTERFACE DEVICE

Brain computer interface (BCI) represents an innovative field of research within the broader domain of human-machine interaction (HMI). A BCI system acts as a bridge enabling a direct connection between the human brain and an external device bypassing the neuromuscular pathway. This opens to a wide range of applications including assisting individuals in daily life activity and enabling the monitoring of the cognitive state of subjects engaged in repetitive tasks for preventing potential risks related to physiological stress and fatigue. Mitigation of these two mental conditions would improve productivity and reduce work accidents caused by lowered reaction time, mind wandering, and sleepiness, all phenomena that may occur while working on an assembly line. Electroencephalography (EEG) is of particular interest in the BCI field since it is a non-invasive technique that allows recording the electrical activity of the brain by means of electrodes distributed along the scalp. In the ICE Laboratory, BCI experiments were performed using a portable BeMicro EEG amplifier² equipped with a cap with 21 Ag/AgCl electrodes. To facilitate data acquisition and storage, a proprietary software, Galileo (EBNeuro S.p.A, Firenze, Italy), was used.

4) PPE

Personal protective equipments (PPEs) are specifically designed to protect a person in an environment that presents potential hazards to safety. Examples of these devices are helmets, gloves, face shield visors, etc., and are common in industrial environments. Depending on the task the worker has to achieve, different PPEs are required to be worn.

During the COVID pandemic, a particular type of PPE, the facial masks, was also enforced in workplaces [20]. It was mandatory to verify the compliance of the workers with such devices [21].

III. CONSIDERED USE CASES

In order to enhance the safety and privacy of workers in the 4.0 industrial environments, such as the ICE Laboratory explained in Section II, this work details six use cases that highlight possible threats to either safety or privacy (or both).

Each use case was designed to solve issues and incidents that can happen in a real-world industrial facility. In Industry 4.0, and even more so in 5.0, the shared space between humans and robots gets increasingly important. Such shared space requires the robots to be aware of the humans to prevent injuries.

For instance, this may happen when humans work together with robots, and such robots need to know what humans are

²EBNeuro S.p.A, Firenze, Italy, <http://www.ebneuro.com/en/ambulatory-egg/be-micro>

doing. If a human is posing himself in an unsafe manner, e.g. being tired or distracted, the system must react to preserve his/her safety. There is also the case in which a human makes a wrong action that is not planned, e.g., a security system is disabled.

At the same time, the human can be unsafe for many reasons, some of which may be independent of the facility. This could be the case of an unexpected fall that requires human intervention (e.g., due to health reasons). In that case, the system must act to avoid worsening the situation, e.g., notifying the facts to supervisors. A similar example can be applied to situations involving multiple people, like a fire in the facility, an earthquake, or any other environmental threat to the lives of the workers.

Moreover, in various places of the facility, it could be mandatory to wear personal protective equipment (PPE), but some workers may forget or use them improperly due to some distractions. A safety system should acknowledge this and enforce compliance. Finally, some areas may be subject to controlled access for crowds, such as during the COVID-19 pandemic. In that scenario, the number of people in the room should be monitored and restricted.

This section presents the use cases as they were designed, presenting for each of them the consequences with respect to safety concerns.

A. ADVERSE EVENT

Usually, the term *adverse event* is used in medicine to represent any untoward medical occurrence in a subject exposed to a pharmaceutical product. Similarly, as defined here, an *adverse event* is any untoward occurrence that can happen to a subject that requires immediate action and attention. An example of this is the “man on the ground”. In this scenario, the worker suddenly fell unconscious due to a sudden unknown cause, i.e., heart attack, stroke, or illness.

This is an extreme case for safety intervention. In particular, the worker may not be conscious or able to move and could be injured by moving vehicles. In this case, the system must acknowledge the situation and immediately notify supervisors.

Implementing technologies to address such adverse events can yield significant benefits. These include faster incident detection and response, enhanced worker safety, and a potential reduction in workplace accidents and fatalities. Moreover, these technological interventions can improve data collection for ongoing safety enhancements and potentially lower costs associated with workplace incidents.

For an in-depth analysis of the technologies involved in this use case, refer to the works [22], [23], [24], [25].

B. MIXED HUMAN-MACHINE ENVIRONMENT

Cobots (collaborative robots) are designed and intended for direct human-robot interaction, in shared spaces where humans and robots are in close proximity, operating for cooperation and collaboration. In manufacturing, and in

in a workspace. PPEs, as described in Section II-D4, can protect the worker from both environmental hazards and injuries from distraction or incorrect usage of machinery. In most states, the safety of workers in working environments is enforced with the introduction of specific laws regarding the usage of determined PPEs.

Monitoring and enforcing PPE usage can yield significant benefits in workplace safety, reducing workplace injuries and accidents, improving compliance with safety regulations, and enhancing overall safety culture.

For an in-depth analysis of the technologies involved in this use case, refer to the works [42], [43].

F. GATHERING

During the COVID-19 pandemic, along with PPEs, people gathering detection and control was critical [44], [45]. The restriction for people gatherings is not only suggested for fighting the pandemic and safety in any environment: the presence of too many people can cause bottlenecks, and potentially block evacuation routes [46].

Detection of people gathering can provide significant safety and operational benefits. Expected outcomes include improved crowd management, reduced risk of disease transmission (important, but not limited to, the COVID-19 pandemic), and enhanced emergency evacuation procedures. Technologies can help maintain optimal occupancy levels, ensure social distancing when necessary, and identify potential bottlenecks in real-time. Furthermore, the data collected can be used to design improvements and optimize space utilization, leading to safer and more efficient environments.

For an in-depth analysis of the technologies involved in this use case, refer to the works [32], [40], [47].

IV. TECHNOLOGIES FOR SAFETY ENFORCEMENT

To address the safety concerns that each of the use cases presented in the previous section may introduce, different technologies are employed. This section presents the several technologies adopted and the results obtained by each technology in order to address the safety issues underlying the use cases. Finally, details regarding the scalability and integration with existing infrastructure are presented. Table 1 summarizes the technology that presents countermeasures in order to take into account each specific use case.

A. VISUAL TRACKING

Visual tracking is useful to enable continuous human localization and tracking inside the ICE Laboratory facility. This tracking system uses solely the RGB surveillance system, consisting of 7 cameras mounted in different locations inside the facility.

The ICE laboratory is a perfect example of a highly occluded indoor environment: the presence of several machines and moving parts can create severe occlusions for human operators who are moving and working in the environment. A visual tracking system must take into account this high level of occlusion. To achieve this, we argue that a

multi-camera system is fundamental and that it's necessary to associate different camera views to cover as many occluded areas as possible. Therefore, we present a multi-camera human tracking system that leverages appearance-based re-identification features for multi-camera matching.

Visual tracking guarantees two main objectives: the tracking of the people in the laboratory and compliance with PPE requirements. In this sense, this technology addresses all the use cases, since can provide meaningful information for any use case we consider.

1) HUMAN TRACKING

The tracking algorithm is implemented as a multi-camera tracker, that uses the frame of each camera to extract some tracklets (the track of a person across the sequential frames) and perform aggregation in the global reference system. This is possible using a robust single-camera tracker [48] in junction with a planar 2D homography between the camera views and the planimetry of the facility. Moreover, the localization and matching of the people tracked by the single cameras is made possible by the homography transformation. Homography is useful because it is easy to set up, since only 4 correspondences are required to create the mapping between the image and the plane, and can be performed with just the camera view. This means it is easily generalizable in any kind of similar environment.

Homography was also used to calculate the actual visual coverage of each camera since it is possible to map the camera view directly to the planimetry.

Furthermore, to take into account the occlusions caused by the laboratory hardware equipment, a novel open-world re-identification (ReID) approach has been added to the tracker. The re-identification module, based on [49], is designed to keep track of each camera tracklet and provide a global tag that represents a unique identifier of the person. This is necessary to perform correct counting as well as potentially provide effective access control. Note that there is no association between this unique identifier and the actual identity of the person.

TABLE 3. The performances in terms of accuracy, precision, and recall, of the open-world occlusion resistant tracking system implemented inside the ICE laboratory for achieving reliable human tracking. The visual coverage column represents the percentage of the areas of the laboratory covered.


# Cam	Precision	Recall	Accuracy	Visual Coverage
1	99.9%	100.0%	99.9%	45%
2	69.4%	100.0%	69.4%	65%
3	83.0%	100.0%	83.0%	90%
4	72.2%	100.0%	72.2%	100%

The system runs in real-time over all the cameras, with a messaging output frequency on the laboratory's Kafka backbone of 10Hz. Performances, w.r.t. the metrics of

TABLE 4. Results in the open-set scenario. The results are provided with the closed-set Market1501 [50], extending it as open-set following [51], and the Facility-ReID dataset [40].

Dataset	Precision \uparrow	Accuracy \uparrow	Recall \uparrow
Market1501 [50]	0.888	0.720	0.770
Facility-ReID [40]	0.908	0.790	0.950

tracking and re-identification are schematized in Table 3. The system was tested on an 8k frames benchmark dataset (2k per camera from 4 different cameras) acquired in the laboratory with 6 people moving in the laboratory creating occlusions with devices, between each other, and hiding from the camera views, in order to stress the multi-camera tracking capabilities of the system [40]. Furthermore, to simulate more people in the dataset, they were requested to exit from the laboratory and change clothes to simulate the entrance of the same person but with a different appearance in order to stress the re-identification module. An example of this dataset is shown in Figure 6. The performance of the system with respect to the state-of-the-art benchmarks is reported in Table 4. Results are reported in terms of precision, accuracy, and recall. More details can be found in [40].



(a)



(b)

FIGURE 6. Example of the frames from the surveillance camera views of the ICE laboratory. These frames are used for visual tracking.

The visual coverage was calculated using homography with planimetry. Since not all camera locations are explicitly designed for human tracking (such as the PTZ or the cameras

observing the robotic arms), the cameras selected for this task were chosen to balance the field of view and the occlusions caused by the machinery in the laboratory. As a result, the use of three (3) partially overlapping cameras is performing best, also considering the coverage of the environment. This is because having more coverage requires the use of cameras that have a very skewed point of view, leading to a distorted appearance of the people, which directly affects the performance of the tracking system.

The system has been containerized and distributed in the computing architecture of the laboratory enabling a fast and scalable way to include more cameras. Furthermore, to ease the computation, split computing has been used to deploy large neural networks on more constrained edge devices [13], [14].

Finally, in an effort to detect the intention of interaction with the laboratory devices, we studied the people's visual selective attention within the production line. For this reason, we created a dataset with a set of baselines that allowed us to get ground truth data related to people's tracks and gaze trajectories [32].

2) PPE DETECTION

With regards to compliance with PPEs equipment, the visual tracking capabilities allow for a dense verification of the law requirements. In this sense, during the period of the COVID-19 pandemic until May 2023 [52], the system was enabled to verify the compliance of the facial masks inside the whole laboratory. The system was implemented using a facial mask detector using a novel dataset, specifically designed to enable face mask detection in crowded environments from surveillance cameras, also created with images acquired from the ICE laboratory [42]. The classifier performs with an overall accuracy of 0.884 on the test partition. The complete PPE detection system implements a combination of a 2D pose estimation pipeline, to ensure the people are facing toward the camera before verifying the compliance of PPEs, and a facial detector that allows the verification of the compliance [43], [53]. This can be, depending on the situation, integrated with the multi-camera visual tracking, i.e., if an alert should be raised, the information on the person's location could be provided.

Now that the COVID-19 pandemic has ended, the requirement for facial masks has also been loosened. Nevertheless, given that this system was designed to consider small (60×60 pixel) faces from a surveillance point-of-view, it could be easily extended to any wearable PPE equipment with slight modifications to the pipeline.

The visual tracking system offers a versatile solution for workplace safety across all identified use cases. It excels in real-time monitoring, providing immediate detection and response capabilities without physical contact. This makes them particularly effective for adverse event detection, cobot interaction monitoring, and crowd management. At the same time, in crowded scenarios, occlusions may reduce the capabilities of the system for PPE compliance or wrong

action prevention. While efficient in providing comprehensive coverage, these systems face challenges such as potential privacy concerns, dependence on lighting conditions, and the need for advanced interpretation in complex scenarios. Their effectiveness in environmental event detection varies based on the nature of the threat.

B. HUMAN POSE ESTIMATION

Visual Human Pose Estimation (HPE) plays an increasingly significant role in the context of Industry 4.0, by simplifying the analysis of human motion through images and video data. Operator position awareness plays a pivotal role in enhancing human-robot collaboration, a fundamental aspect of Industry 4.0 [54], and even more in Industry 5.0 [55], [56], while also serving as a critical safety measure to prevent collisions with moving components, collaborative or mobile robots, and to detect potential injuries and faints.

As mentioned in Section III-B, the industrial scenarios are rich in machinery which shares the workspace with the operators. This can cause occlusions and absence of detection if a single-view system is used. Therefore, for the ICE laboratory, we developed a multi-camera distributed HPE system to have complete coverage and redundancy of information [57]. For each view, an RGB-D camera is connected to an embedded board equipped with a GPU, which processes the images locally, thus reducing data traffic on the communication network, ensuring scalability, and avoiding the transmission of sensitive information. In this type of architecture, there are several challenges to overcome. The first is limited computational resources on edge devices, and since the HPE software has to ensure real-time constraints, it has to be lightweight sacrificing the accuracy needed for human-robot interaction. The second challenge is related to the transmission network which, being standard and shared with other applications, uses industrial protocols. This results in bandwidth fluctuations, packet losses, and delays, making the synchronization of poses from the same scene acquired by different edge nodes difficult. The third main challenge is related to multi-person scenarios, specifically associating information about the same operator from various views (i.e., edge devices) that are not hardware synchronized.

To overcome these challenges first, we customized a 2D HPE framework designed for low-powered boards. Thanks to the 2D keypoints that represent the human body joints in the image space (in pixels) and the depth information given by the sensor, we are able, through deprojections and rototranslations and using parameters obtained during a calibration phase, to obtain for each view a set of 3D skeletons in a global reference system (in meters). Once this information is obtained, it is encapsulated in a pub/sub message labeled with a timestamp. The timestamps at the time of sending can be assumed to be synchronized across edge devices, relying on the NTP protocol. Finally, the messages are sent to an aggregator that is responsible for

merging the contributions from the various views. First, the messages are temporally synchronized thanks to an internal virtual clock, and those that are too old are discarded. Then clusters of points from the various views (in the same reference system) that belong to the same person are identified through a clustering algorithm. Using an association algorithm, a temporal continuity is given between the current instant to the previous, which is useful for temporal filtering. Finally, the centralized unit, according to fusion metrics, merges the data from the multiple views to obtain complete information robust to occlusions and of higher quality. This information is finally forwarded to Kafka to become accessible to all applications that make use of it. These applications use 3D positions to monitor operator operations on the production line and detect irregularities that can result from collisions with mobile and collaborative robots [24], [41], [58], as well as fainting and illness. They also are useful for locating operators in case of dangerous situations. In addition, this information can be integrated with inertial sensors to obtain even more robust and complete information [23].




FIGURE 7. Scalability of the proposed platform (green bar), compared with a full-centralized approach (red bar) and an RTSP-based centralized system (cyan bar).

Figure 7 shows the scalability of the proposed system, by varying the number of views. For the comparison, we used two different approaches. The first is a centralized approach, in which the views are connected directly to a server that executes the several pose estimators. The second approach, on the other hand, involves the transmission of information about the views from the edge devices via the RTSP protocol, with the pose estimators always executed on a centralized server. As you can see, initially, the approach with a higher working frequency is the centralized one, as there is no transmission latency, and the server's hardware is much more performant. In contrast, the RTSP-based approach suffers from significant transmission latencies due to image compression and transmission. Both methods fail with more than five views because of memory saturation caused by the neural network models in the pose estimator. The proposed methodology, on the other hand, is shown to be scalable with a variable number of cameras, maintaining a constant frame rate of about 20 fps.




FIGURE 8. Accuracy, expressed in mean absolute error (a) and correlation (b) of the proposed platform (green bar), compared with a full-centralized approach (red bar) and an RTSP-based centralized system (cyan bar).

Figure 8 shows the accuracy of the proposed system compared to the other two methodologies previously presented. The accuracy is expressed in terms of Mean Absolute Error (MAE) and Correlation compared to an accurate and heavyweight software (i.e., Openpose [59]), customized to operate in multi-view, but not made for edge computing devices. This test was done to demonstrate the degree to which the proposed system can achieve comparable accuracy results with centralized systems, overcoming issues due to synchronization and the use of edge computing software. As can be seen, by increasing the number of cameras, MAE decreases. In particular, it is useful to note how the proposed system (green bar) achieves similar results to the centralized system (red bar). This means that the system manages to be temporally synchronized, as opposed to the solution implemented using the RTSP protocol (cyan bar). The same conclusions can be drawn from the results in terms of correlation.

Human pose estimation technology, applied to adverse events, mixed human-machine environments, and PPE compliance, offers a targeted approach to workplace safety. This technology excels in detecting unusual body positions, crucial for identifying adverse events like falls or accidents. In human-machine interactions, it helps maintain safe distances and predict potential collisions with cobots. For PPE compliance, it can effectively detect whether workers are wearing the required equipment since the localization of body parts can help to enhance PPE detection performance.

The strength of pose estimation lies in its ability to provide detailed, real-time information about worker positioning and movement without relying on wearable devices. However, it may struggle with occlusions or complex environments. While generally effective, its accuracy can be affected by clothing, lighting, and crowded scenes. Privacy concerns are reduced compared to full visual tracking, but implementation costs and computational requirements can be significant.

C. WEARABLES

To further strengthen the recognition capabilities of the HPE platform and reduce the recognition errors in the multi-person scenario, we introduced a dedicated wearable device into the context of the ICE/Opera lab. The wearable is a

technological solution that leverages the nRF52832 System on Chip (SoC). This wearable device is equipped with diverse communication protocols, most notably Bluetooth 5 and Bluetooth Low Energy (BLE). Its intrinsic on-chip power management system ensures minimal energy consumption, a critical aspect for prolonged operational effectiveness. Beyond its communication capabilities, the wearable integrates essential sensors, including motion (i.e., 3-axis accelerometer, gyroscope, and compass) and environmental (i.e., temperature, pressure, humidity). It presents on-device (usually known as on-the-edge computation) computation capabilities for calculating fundamental metrics like Euler angles, quaternions, gravity vectors, and rotation matrices. Moreover, to tailor its functionality to specific needs, the sensor's sampling frequency is fully adjustable, offering a range spanning from 5Hz to 200Hz. This adaptability caters to the granularity required to fulfill the intricacies of the monitoring objectives. Within the broader context, the fusion of different individual wearable devices coalesces into establishing a coherent Wireless Body Area Network (WBAN) [35]. This WBAN, meticulously designed, enables seamless real-time monitoring and tracking, unburdened by the constraints of physical wiring.

The wearable presents a unique identifier (ID) that serves as a distinctive marker. Its primary function encompasses gesture recognition, user identification, and synchronization between the wearable and the pose estimation system. The overarching objective is to ensure the continuity of the association between users recognized by the pose estimation system and the corresponding individuals uniquely identified by their wearable IDs. This linkage is pivotal in maintaining a robust connection between the two realms. The system achieves this by aligning the wrist movements detected by the wearables with those captured by the pose estimation system, using a numerical correlation mechanism to establish a meaningful connection. This correlation process effectively bridges the two sources of data. Furthermore, the wearable assumes a crucial role in monitoring operator status within areas of the facility that lie beyond the coverage of the tracking or pose estimation systems. In these scenarios, the wearable becomes an indispensable tool, facilitating the ongoing surveillance of the operator's well-being and activities [23].

The system underwent testing involving four male individuals with diverse body characteristics. The test protocol entailed the execution of six repetitions of the following two actions for a minimum of 1 second: a) the subject raises the hand above the head and moves it from left to right or right to left, and b) the subject raises the hand above the head and executes a vertical rotation. Table 5 provides an overview of the recognition tests performed over the wearable base hand waving recognition and correlation accuracy of the hand waving recognized by the HPE system. The first "Gesture" column presents the wearable device's accuracy in detecting the first type of gesture, which it proficiently recognizes. Additionally, the system exhibits a strong correlation in

TABLE 5. Wearable-based recognition performance on the two cases. (a) the subject raises the hand above the head and moves left to right. (b) the subject raises the hand above the head and executes a vertical rotation.

Subject info			Performance with wearable (a)		Performance with HPE (b)	
ID	Height	Weight	Gesture (%)	Correlation (%)	Gesture (%)	Correlation (%)
1	195	110	100%	98%	96%	36%
2	180	105	100%	96%	97%	41%
3	178	72	100%	99%	96%	34%
4	177	79	100%	95%	95%	30%

recognizing handshakes via the wearable device and tracking hand movements, as indicated in the first “Correlation” column of Table 5. However, it is noteworthy that the system’s performance is less satisfactory regarding the hand movement described in type b). Specifically, the wearable device excels at accurately identifying vertical hand rotations (second “Gesture” column), whereas the HPE system struggles to recognize such movements (second “Correlation” column). This limitation is primarily attributed to the HPE’s inability to discern rotation information at extended distances, as it can only identify the subject’s wrist as a singular point in space, offering insights into lateral and depth movements but falling short in capturing rotational aspects.

Table 6 shows the results in terms of the accuracy of the wearable solution of specific tasks on public benchmarks.

Finally, the software managing the wearable sensors within the environment is encapsulated within a Docker container and executed on the NVIDIA Jetson board, which is also employed for the HPE, as outlined in Section IV-B. Additionally, each wearable sensor transmits its data to Kafka through the Jetson board, making it accessible for further applications, such as the collision prediction system [24], [41]. Depending on the configuration, the data can either be streamed based on specific events (e.g., detection of a fall or sensor shaking) or in its raw form (e.g., every captured snapshot of inertial data). Figure 9 shows the wearable employed.

1) USABILITY AND ACCEPTABILITY

To investigate the usability and acceptability of the proposed wearable, we surveyed the users via an anonymous questionnaire. The questionnaire features ten questions about user-friendliness, training requirements, reliability, and comfortability. Users’ feedback scores range from 0 to 3, where 0 means “not at all”, and 3 means “very true”. We surveyed eight users, and the survey outcome shows how users found the wearable device helpful. The questionnaire results are reported in [35].

Wearable devices, combined with human pose estimation, offer an enhanced solution for adverse events, mixed human-machine environments, and environmental events. This combination significantly improves accuracy in multi-person scenarios, addressing a key limitation of standalone systems.




FIGURE 9. The wearable device. On the top: the front and rear view of the Nordic device. On the bottom: the enclosure case of the wearable device.

In adverse event detection, the integrated approach enables more precise identification of falls or unusual movements. For human-machine interactions, it allows for nuanced gesture-based control of cobots, enhancing safety and efficiency. During environmental events, the system can interpret evacuation-related gestures, improving response coordination.

This integrated approach reduces false positives and enhances overall system reliability. It maintains the benefits of continuous individual monitoring while mitigating the challenges of occlusion and complex environments. Privacy concerns remain low compared to visual-only systems. The main trade-offs are increased complexity and potential cost. Overall, this combined wearable and pose estimation system offers a highly effective, privacy-conscious solution for the specified use cases, particularly excelling in crowded or dynamic work environments.

D. DATA-DRIVEN SAFETY ASSESSMENT

Risk prediction and avoidance in industrial scenarios are of utmost importance. In the era of Industry 5.0, the interest in this topic has been steadily increasing within smart factories, where autonomous robots strongly interact with humans. Contacts between humans and robots must be managed carefully to avoid injuries and to prevent stops in the production chain, especially in the case of collaborative robotics. In this context, several techniques have been developed to avoid collisions, particularly based on RGB cameras and human pose estimation software. In all existing solutions, the common challenge consists of implementing a platform for the assessment of risky situations and the prediction of human-robot collisions that guarantees both

TABLE 6. Results of the wearable technology w.r.t. to state-of-the-art dataset and approaches.

Technologies	Task	Dataset	Accuracy
Wearables and HPE	Identification and tracking	Self-collected [36] + TotalCapture [60]	0.97
Wearables	Recognition of the phases of motor tasks	Self-collected [36]	0.90

accuracy and computational efficiency. These are, in fact, conflicting requirements; accuracy generally comes at the cost of very intensive computing activities lasting for long periods (i.e., up to 24/7 in some cases), since large sets of variables (i.e., 3D human joints and 3D robot joints) must be monitored and their evolution simulated in realtime; when resource-constrained computing platforms are used, the cost of risk assessment and prediction can steal computing power from ordinary tasks, negatively affecting the system’s functionality. Section presents a Human-Robot Interaction (HRI) framework that combines a lightweight assertion mining-based engine, with a heavier collision prediction system to pursue both accuracy and computational efficiency in human vs. robot collision avoidance. A preliminary description of the methodology can be found in [24].

This allows the HRI framework to coexist with ordinary applications in a resource-constraint computational platform; these applications implement the actual functionality of the target system, such as the monitoring and controlling of tasks in a smart manufacturing line for Industry 4.0.

Figure 10 shows an overview of HRI. The inputs are the temporal-synchronized traces representing the positions of the human and robot joints in the 3D space.

A distributed real-time multi-view 3D human pose estimation system extracts the traces of the human joints. This system implements inference on RGB streams through convolutional neural networks (CNNs) to extrapolate the human joints in 2D. It combines such information with the one extrapolated by a depth camera and implements a subsequent pose aggregation step to reconstruct the 3D skeletons representing the people in the scene. For robot trace extraction, the platform exploits the geometrical position of the robot joints and a set of discrete variables that allow determining its discrete configuration. These two parts are temporally and spatially synchronized through a network time protocol [61] and a calibration phase [62].

The traces are analyzed at runtime by a risk detection software based on the HARM [33]³ assertion miner. It continuously monitors the system traces, searching for occurrences of unsafe states by verifying position-based assertions. When an assertion fires, the system is considered to have entered a risky situation. Then HARM activates a collision prediction module called OPERA,⁴ whose goal is to precisely generate the temporal range within which a possible collision is expected to occur. In parallel, HARM provides the user with the causes that led to the unsafe state (e.g., the robot arm and the human body were too near).

OPERA is the most computationally and memory-consuming block inside HRI. It generates accurate results in terms of collision prediction, as it provides spatial and temporal information for possible future collisions between body parts. On the contrary, HARM is sensibly lighter in detecting unsafe states expressed as LTL formulas. However, LTL assertions do not provide a detailed representation of the risky state as OPERA does through its geometric representation and prediction model. This approach implements an effective and efficient continuous risk detection system, activating OPERA on demand when the target system enters an unsafe condition and deactivating it once the system returns to a safe state.

³<https://github.com/SamueleGerminiani/harm>

⁴<https://github.com/ariadne-cps/opera>




FIGURE 10. Overview of the HRI risk assessment and prediction system.

The first goal of the proposed HRI framework is to predict the risk of collisions in an unsafe human-robot environment and to explain the causes leading to such a risk.

The second goal is to prevent the HRI framework from depleting the computational resources of the whole system.

In this combined solution, the risk detection task (i.e., *HARM*) acts as a trigger for the collision prediction and risk explanation. For this reason, it is always active. While the risk detection task recognizes a safe state of the environment, the collision prediction *OPERA* and the risk explanation modules are not executed even if *OPERA* still acquires robot messages to derive a history of its behavior.

The result of this switched model is a multi-tasking framework that efficiently manages the CPUs and memory resources, providing a comprehensive interpretation of a risky scenario that considers the prediction of collisions between human and robot parts while determining the cause of the hazardous scenario.

Data-driven safety assessment, as a risk prediction and collision avoidance system, offers a proactive approach to safety in adverse events, mixed human-machine environments, and environmental events. For adverse events, it can identify high-risk situations or individuals. In human-machine interactions, it can anticipate and prevent potential conflicts or collisions. For environmental events, it aids in predicting and preparing for various scenarios.

The system's strength lies in its ability to leverage historical and real-time data for predictive insights, enabling preemptive action. It's highly effective in complex, data-rich environments. However, its accuracy depends on the quality and quantity of available data (such as for the linked services like HPE and robot states). While it enhances overall safety, it may struggle with unprecedented situations. Privacy concerns are moderate, as it requires extensive data collection. Finally, the implementation can be complex and resource-intensive.

E. BRAIN COMPUTER INTERFACE

Brain Computer Interfaces (BCIs) encompass two distinct categories: active and passive. Active BCI is a closed-loop system, translating the user's thoughts into actions in real time. In contrast, passive BCI monitors the user's mental state over time. Passive BCI systems could hold value in a manufacturing context, assessing the operators' vigilance state for a safer mixed human-robot environment. This helps prevent potentially harmful actions impacting operators' health and production capabilities. Vigilance, defined as the ability to sustain attention over an extended period of time, is crucial in contexts involving repetitive stimuli [63]. Accidents are often the results of impaired vigilance capabilities [64].

A correlate of vigilance can be studied by relying on EEG. The electrical activity captured reveals an increase in a specific frequency band (8-12 Hz), known as the α rhythm when the subject is getting more tired and thus less vigilant.

In the pursuit of extracting relevant brain wave patterns from EEG recordings for BCI applications, various signal processing pipelines were devised using advanced signal processing techniques. These strategies were implemented to enhance the performance of active BCI systems during motor imagery tasks. This involved integrating the empirical

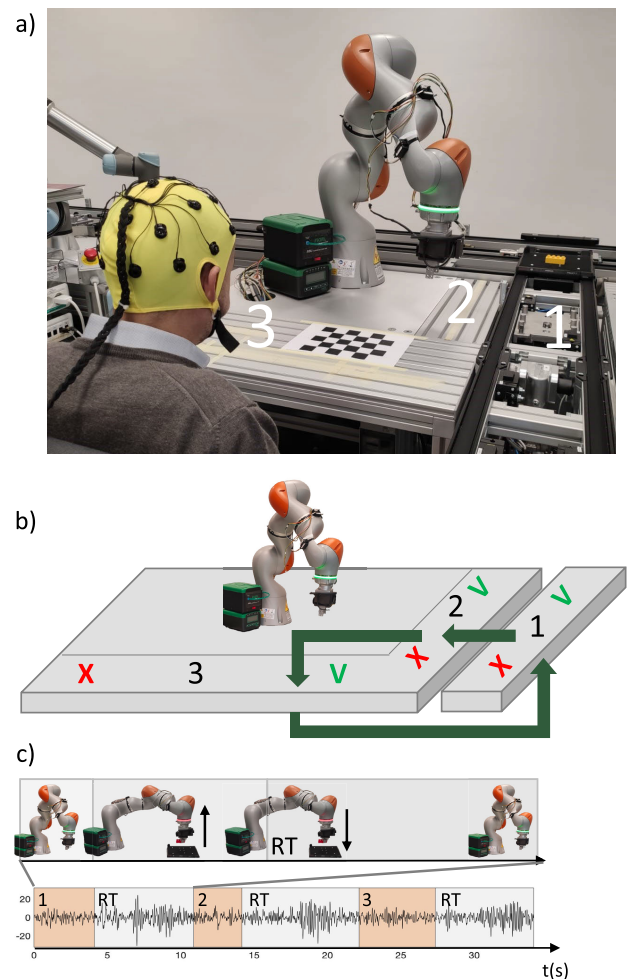


FIGURE 11. Representation of the experimental procedure and signal acquisition. a) A volunteer in the test environment; b) Robot workspace in which three different movements are continuously repeated (i.e., from position 1 to 2, from 2 to 3, from 3 to 1, etc.); c) Timing schedule of the robot movements and an example of processed EEG signals acquired during an experimental trial.

wavelet transform and scattering convolutional networks within an innovative framework that embraces brain connectivity measures. In comparison with existing methods, both binary and multi-class pipelines have proven their effectiveness in the classification tasks considered [37], [38], [39].

Moreover, this research involved acquiring data sets in the ICE Laboratory with the specific aim of monitoring operators' vigilance in the manufacturing pipeline. To this end, 10 subjects performed a surveillance task, observing the KUKA robotic arm while it moved cyclically a block between "right" and "wrong" positions and assessing the positions as such (Figure 11). This task lasted 23 minutes. A single trial of the experiment coincides with the total duration of a pick-and-place operation and lasts approximately 11 seconds. In total, 120 trials were performed. A consistent increase in α frequency band activity was observed across all subjects, discerned from the time-frequency behavior of the signal,




FIGURE 12. a) α wave amplitude trend monitored along the task, for a single subject [light blue], average α power of subject's resting state [red], fitted line [black] and the 95% prediction interval [gray]. Graph are placed as to mimick the original electrodes position on the head, which can be seen in Fig. b).

which was extracted through the continuous wavelet transform. Figure 12 shows an example of increasing α oscillations with time-on-task, especially in the centro-parietal channels (C3, C4, P3, and P4).

On self-reported questionnaires, subjects express a general sense of “Frustration” toward the task (mean value of 16.4 on a scale from 1 to 20, with 1 being “no frustration” and 20 being “very frustrating”), and the sensation of the task being slow and boring (“Temporal Demand” has mean value of 2.9 on a scale from 1 to 20, with 1 being “too slow” and 20 being “too fast”).

This analysis, an extension of what was presented in [34], shows that modeling and forecasting the α frequency band activity is a feasible strategy that can improve the safety of operators working in a mixed human-robot environment. Real-time prediction of a state of lowered vigilance can be used to mitigate the effects of boredom by modulating the task intensity, or by requiring additional user feedback. Studies on air traffic controllers [64] demonstrate that having the monitored system engage with the operators reduces the total number of mistakes. The forecast would allow the planning of the incoming tasks beforehand, improving productivity across the whole assembly line.

While forecasting of the EEG signal has shown success in other fields, such as predicting the onset of a crisis in epilepsy [65], its application in the industrial context is still

unexplored, and it is a novel field of research. The α fluctuations over time were estimated using a vector autoregressive (VAR) mode, representing a linear multivariate time series by incorporating the past n values [66]. The model order selection was based on the Akaike information criterion (AIC) for each subject; the final fitting was performed with a model of order 6 for all subjects. Only four channels, namely C3, C4, P3, and P4, which exhibited more significant variations in activity within the frequency band of interest, as illustrated in Figure 12, were included in the model. Mean absolute error (MAE) between the forecast horizons (1/44, 1/4, 1/2, 3/4, 1 fractions of a trial) and the true signals was computed for each data point of subjects' α activity. Results show that forecasting data points up to 2.75 seconds ahead (1/4 of a trial) gives satisfactory performances (avg. median over channels: 0.068, avg. upper whisker bound over channels: 0.169), which degrade with further horizon extensions (Figure 13). Still, the VAR model suffers from limitations, as it is a linear method that may not fully capture the oscillatory nature of attention-influenced EEG signals. A more sophisticated and up-to-date approach to forecasting will be adopted in future work, with a focus on using recurrent neural networks [67] and temporal convolutional networks [68], taking inspiration from successful forecasting applications on biological time series, as in the case of continuous glucose monitoring [69] or epilepsy prediction

[70]. The vast majority of EEG-related works focus on the classification of mental states over collected data rather than on forecasting [71]. Hence, due to the original nature of this work, and the specificity of the collected and processed data, a direct comparison with EEG forecasting application yields few basis. Nevertheless, the proposed approach has the potential to anticipate risky situations for operators when implemented within a manufacturing pipeline, especially if combined with forecasting techniques proven effective in similar domains.

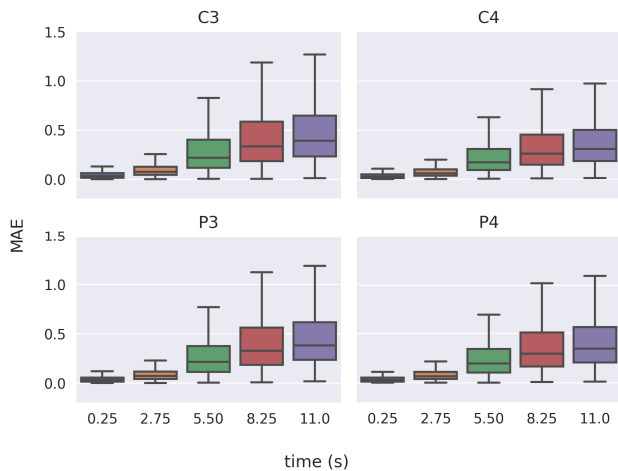


FIGURE 13. VAR MAE results aggregated over the 10 subjects. Results are reported on a log scale. Y-axis: MAE values for the predictions. X-axis: forecast time horizon. Outliers were not plotted for clarity of representation.

The constricted setting of the experiment is a consequence of practical limitations: the equipment does not allow for the subject's movements, since even tiny oscillations of the cables may cause heavy artifacts; furthermore, the equipment intended use foresees a controlled environment with electromagnetic shielding, since the extremely low amplitude of the signal can easily be disrupted by external interference. Another obstacle to the adoption of EEG techniques in the industrial context is the need to wear the headset for a prolonged time, which would require the re-applying of the gel, causing annoyance.

Recent works offer possible solutions to the aforementioned issues; in particular, the adoption of portable EEG headsets allow for effective monitoring and classification of construction workers' mental state with a reduced amount of channels [72], [73], [74], a paradigm that mitigates the challenges of a practical implementation.

On a final note, the extraction of a vigilance index must be treated with extreme caution. The user's informed consent is necessary before performing any acquisition. Data acquired in real-time must be inaccessible, anonymous, and deleted at the end of the work shift. To study the difficulty in maintaining focus during a specific task, an offline application must have the collected data password-protected and identified exclusively through a serial number. This

offline approach provides a reliable measure of operators' engagement that could be used to optimize the working shifts, tasks, and environment. Regarding the experiments, participants were asked to give informed consent before participating in the study. Data were collected according to the protocol approved by the local ethical committee of the University of Verona.⁵

Brain-Computer Interface technology using EEG for predicting user alert states, fatigue, and stress offers a unique approach to safety in mixed human-machine environments and preventing wrong actions. This system provides real-time insights into a worker's cognitive state, potentially preventing accidents due to fatigue or stress-induced errors.

In human-machine interactions, BCI can adjust cobot behavior based on the operator's mental state, enhancing safety. For preventing wrong actions, it can trigger alerts or interventions when detecting decreased alertness or high stress levels.

The technology's strength lies in its ability to detect subtle cognitive changes before they manifest in observable behavior. However, its effectiveness is currently limited by the lack of portable devices and sensitivity to magnetic fluctuations in industrial settings. While potentially highly accurate for individual monitoring, it faces significant implementation challenges. Privacy concerns are high due to the nature of brain data collection. The research-stage status also implies high costs and limited real-world validation.

Overall, BCI presents a promising but nascent approach to workplace safety, offering unique insights but requiring substantial development to overcome practical limitations and privacy concerns.

F. MIXED REALITY AND ADVANCED INTERACTION

Mixed Reality technologies in the ICE lab are aimed at (i) creating spatially registered augmentation for workers providing relevant information to support tasks and prevent errors; (ii) interacting with machines by manipulating virtual objects or performing gestures; (iii) Collaborating with other local and remote users through shared virtual artifacts; (iv) support training in immersive virtual reality and remote assistance. For this reason, it was realized: (a) a visually accurate 3D model of the lab, with walls, ceiling, windows, light sources, and including optimized digital models of lab assets: mobile robots, transport line, and all the other machines, optimized and available at multiple Levels of Detail for fast rendering and with realistic materials applied; (b) an AR app for the Hololens 2 (Figure 14) using selected markers in the scene to create a good alignment of the virtual and real lab, allowing the visualization of spatially registered augmentations; (c) a VR app for the Meta Quest 2, where subjects can navigate inside the virtual lab (Figure 15) where the status of the machines can be simulated or updated in real-time exploiting the Kafka messages.

⁵Comitato di approvazione della ricerca sulla persona [CARP]: protocollo n. 2.R/2022.




FIGURE 14. The HoloLens 2 AR app allows workers to monitor the status of the production line with spatially-referenced augmentations (left) and to show alerts and prevent errors, showing instructions for emergency actions (right).




FIGURE 15. The graphical twin of the ICE lab allows remote monitoring and simulation of tasks in immersive VR.

These tools are currently ready for testing and several experiments are planned to demonstrate safety-related use cases in the industrial domain. The combination of mixed reality and digital twin technologies is not new, but, as shown in [75] the readiness level of effective AR-assisted digital twins is low, few works showing complete implementations of systems and/or user evaluation are available and typically they are not focused on safety-oriented tasks.

Typical safety-related use cases for VR and AR are related to training and monitoring specific tasks on the machines. For this reason, a specialized framework is currently being developed to support task definition, training, and monitoring. The framework will allow experienced workers to record sequences of actions defining an example task simply by performing it in the real or in the virtual work environment. The application will record it as a standardized sequence of events. In the next step, a potential new worker will then be able to be trained to complete the task in the VR environment or in the real one with AR assistance. The VR/AR application will guide the user and provide hints on the next action required to complete the task. It will also monitor the correct execution and raise warnings in case of mistakes that could result in harmful behaviors.

An original aspect of this work is that both the AR and VR applications feature custom solutions allowing the users to interact “naturally” with the system using gestures.

This technology is particularly interesting for industrial applications, as it supports an intuitive and quick command execution in a potentially noisy environment but is challenging due to the low-latency constraints and the risk of false detection. The custom gesture-recognition technology [18] is specifically designed to recognize the gestures of a complex dictionary [19], avoiding false positives and achieving state-of-the-art classification accuracy. Table 7 presents the results against existing competitors.

TABLE 7. Classification results of SHREC22 benchmark. The metrics are the detection rate (DR), false positives scores (FP), the Jaccard Index (JI), the delay in frames, and per-frame processing time. In brackets are the standard deviations.

Method	DR ↑	FP ↓	JI ↑	Delay(fr.) ↓	time(ms) ↓
Stronger [19] (2022)	0.72 (.11)	0.34 (.26)	0.59 (.18)	14.79	100.0
DSTA [76] (2020)	0.73 (.07)	0.24 (.13)	0.61 (.12)	8.00	9.2
2ST-GCN SF [19] (2022)	0.74 (.12)	0.23 (.05)	0.61 (.11)	13.28	2.1
TN-FSM+JD [19] (2022)	0.77 (.06)	0.23 (.12)	0.63 (.03)	10.00	4.6
Causal TCN [19] (2022)	0.80 (.15)	0.29 (.22)	0.68 (.24)	19.00	28.0
DDNet [77] (2019)	0.88 (.06)	0.16 (.18)	0.78 (.14)	8.00	2.2
OO-dMVM (2023)	0.92 (.06)	0.09 (.09)	0.85 (.11)	8.00	4.1

Preliminary tests demonstrated the effectiveness of the gestural interfaces in controlling the task flow in the ICE lab, but more complex experiments are needed to assess the impact of its regular use on the network performance. The current work investigates methods for creating optimal gesture dictionaries for various tasks, focusing on maximizing the ease of gesture execution, their learnability, and their recognizability by AI-based systems.

Gesture recognition systems could significantly enhance workplace safety. They can be used to detect/forecast operators’ errors and to deliver related warnings in the AR application. Gestural interfaces can also be an effective solution in work environments where human operators should act on potentially harmful machines. In fact, by controlling the machines at a distance by using gestures, it is possible to eliminate the risk of injuries caused by the direct interaction of the workers with possibly dangerous parts. The challenge in this scenario is to design a gesture-based interface that can keep the workers safe without hindering their ability to effectively control the machines. Another significant advantage of touchless gestural interfaces in terms of safety is that they eliminate potential problems related to hygiene and cleanliness of the interfaces.

Finally, a critical mechanism that can enhance safety is the support for multiple users collaboration in tasks’ control, possibly involving both local and remote users. To support collaborative work in mixed reality, preventing risks of unwanted actions, some authors of this work recently proposed a method to manage concurrency in the editing of augmented/virtual spaces [78] that can be exploited in collaborative industrial applications as well.

Mixed reality technologies, using devices like HoloLens 2 AR and Meta Quest 2, offer innovative solutions for adverse events, mixed human-machine environments, and preventing wrong actions. In adverse event scenarios, these systems can overlay critical information and guide response actions in real-time. For human-machine interactions, they provide immersive visualizations of safety zones and potential hazards around cobots. In preventing wrong actions, AR can highlight correct procedures and warn against potential errors.

The strength of mixed reality lies in its ability to provide context-rich, intuitive guidance directly in the user's field of view. However, adoption may be hindered by the need to wear headsets, which could be cumbersome in some work environments, and can be costly. While highly effective for training and guided operations, their real-time responsiveness to rapidly changing situations may be limited. Privacy concerns are minimal, but data security for the augmented information is crucial. Cost and user adaptation remain significant factors.

G. TECHNOLOGY SELECTION CRITERIA

Some technologies are state-of-the-art with no existing alternatives, so a proper comparison with other solutions is not possible. This is the case, for instance, of the BCI solution. Other more traditional solutions (like visual tracking, HPE, etc...) indeed have been extensively analyzed with the possible alternatives in both methodological alternative solutions and commercial ones.

For visual tracking and HPE, were initially evaluated commercially available marker-based solutions. The issue with these approaches is the high sensitivity to lighting changes, which, unfortunately, due to regulations on natural illumination, is constantly happening during the day.

For the XR devices, the HoloLens 2 by Microsoft is one of the most advanced AR devices on the market, featuring a wide field of view, precise hand tracking, and natural interaction with holograms. These capabilities make it particularly suitable for professional and industrial applications. Widely appreciated in fields such as industrial maintenance, professional training, architectural design, and medicine, HoloLens 2 allows users to visualize and interact with complex data intuitively. Furthermore, it is seamlessly integrated with Microsoft's enterprise tools and platforms. Its design ensures comfort for extended use, with well-balanced weight distribution and ergonomic features that reduce fatigue and enhance productivity.

Regarding the issue of VR headsets, Meta has positioned the Quest 2 as a cost-effective solution, making it one of the most affordable high-quality VR headsets on the market. This price advantage allows businesses to equip their teams with advanced VR technology without a significant financial burden. Additionally, Meta's strategic development ensures that all applications designed for Quest 2 will be easily portable to the upcoming Quest 3. This forward-thinking approach makes the Quest 2 a future-proof investment, as businesses can confidently adopt the platform, knowing that their software and tools will remain compatible with future hardware iterations. The Quest 2 and HoloLens 2 are standalone devices, meaning they do not require a connection to a PC to function. This makes them highly portable and easy to use in various environments without additional hardware.

The wearable device, the Nordic Thingy 52, features the nRF52832 SoC manufactured by Nordic Semiconductors.

Third-party manufacturers utilize this SoC in their wearable devices. Besides, the Thingy 52 offers a versatile, cost-effective solution with extensive developer support and the possibility of developing custom firmware. The interaction through BLE allows the efficient collection of the sensor's data regarding power consumption. On the other hand, other devices like Shimmer 3 or Empatica E4 are more expensive, and they require specific software to collect data, thus without allowing firmware or sensor management software customization.

H. SCALABILITY AND COMPATIBILITY

The technologies developed in the Laboratory were designed to collaborate with other subsystems, even if they previously existed. The laboratory's infrastructure is designed for this very scope. This happens thanks to pub-sub mechanisms (such as those integrated with ROS for robotic systems and Kafka for other software interfaces) that allow for intensive network communication.

The solutions employed in the laboratory are released using container technologies. This means the horizontal scalability is provided by the availability of the computing resources in the Kubernetes cluster, as specified in Section II. The vertical scalability is entrusted to individual technologies and their related maintainers. From the laboratory point of view, a broad 1 Gigabit ethernet connection is implemented, which is solely used for internal data transmission. The data volume transmitted over the network consists of single status messages from the machines, well below the bandwidth's cap.

Finally, it is worth mentioning that many of the technologies were designed to be modular in the sense that they can be inserted into facilities with different configurations. For instance, visual tracking may work on a standard video surveillance system without relying on acquiring new cameras, or the HPE system can be installed using portable PoE devices that compute directly on edge without relying on a powerful server.

I. SUSTAINABILITY AND MAINTENANCE

Since the software components are deployed containerized, updating such systems involves creating a new Docker image that is then pushed to a local docker registry, which is then sent to the Kubernetes cluster described in Section II. Each technology provides input parameters and configurations depending on the tasks they need to solve. For instance, for the visual system, there are the calibration parameters w.r.t. the laboratory's reference system, thresholds for the algorithms, etc. Every technology provides documentation that is shared with the laboratory's operatives.

Regarding obsolescence risks, the technologies employed are well-known or state-of-the-art approaches. The backbone software structures (Kafka, ROS, etc.) are standard for industry-grade facilities, while state-of-the-art technologies are the cutting-edge solutions to address safety and privacy concerns.

V. CONTEXT-AWARE PRIVACY MONITORING

This section discusses how the trade-off between safety and privacy can be achieved in smart manufacturing environments. Section V-A explores the privacy threats that safety technologies can pose to workers and highlight the necessity of implementing a privacy monitoring solution. Then, Section V-B presents the privacy monitoring solution and discusses the challenges related to the implementation of the solution in smart manufacturing environments.

A. ADDRESSING PRIVACY THREATS

The use of safety technologies that monitor workers' activities, while aimed at enhancing workplace safety, raises several privacy threats that can be identified using Daniel Solove's taxonomy of privacy flaws [79]. The taxonomy includes four main categories of privacy threats: personal data collection, processing, dissemination and invasions into people private affairs. The main threat related to data collection is *surveillance* of workers' activities. Continuous monitoring of employees' activities, locations, and behaviors through cameras, wearable devices, and BCI systems can create an environment of constant observation. This pervasive surveillance can lead to a feeling of being watched, eroding trust and potentially creating a stressful work environment.

Moreover, the further processing of the information collected by multi-camera human tracking systems, pose estimation systems, BCI systems, and wearables exposes workers to *aggregation*, *identification*, *secondary use*, and *exclusion* privacy threats. By integrating data collected from these systems (*aggregation*), employers can build extensive profiles of workers, including detailed records of their movements, interactions, and behaviors over time.

This aggregated data can reveal sensitive information about workers' habits and routines. Additionally, the ability to uniquely identify (*identification*) and track individuals within the smart factory using wearables enhances the system's capability to link specific actions and behaviors to particular workers. Data collected for safety and operational efficiency might be repurposed (*secondary use*) without workers' consent for other uses, such as performance evaluations, disciplinary actions, or marketing purposes, extending beyond the original intent of the data collection. Furthermore, workers are often unaware of the full extent of data collection and processing related to the use of safety enforcement systems and are typically excluded from the decision-making processes regarding how their data is used, stored, or shared (*exclusion*).

Information dissemination poses additional threats. If the data collected by these systems is not adequately protected, it can be accessed by unauthorized individuals, leading to breaches that expose sensitive data and result in potential misuse, identity theft, or other harms (*breach of confidentiality*). Moreover, data from safety enforcement systems can be shared with third parties, such as insurance

companies or law enforcement, without the workers' consent (*disclosure*). This can lead to unintended consequences for workers.

Invasion of privacy is another major concern with these systems. The use of these safety enforcement systems in the workplace can be seen as an intrusion into the personal space of workers. Even though it is a workplace, workers have a reasonable expectation of some level of privacy, which can be eroded by constant monitoring (*intrusion*). The knowledge that they are being constantly monitored can influence workers' behaviors and decisions, causing them to alter their natural actions to avoid being flagged or scrutinized by their employers (*decisional interference*).

To mitigate these privacy threats, it is essential to implement a privacy monitoring system that tracks information flows from safety enforcement systems to other applications within a smart manufacturing environment. This system should only allow personal data flows in safety-critical situations or when restrictions are enforced to the consumption of the data by these applications.

At the same time, implementing privacy monitoring systems in real-world smart manufacturing settings poses significant challenges. Ensuring compliance with General Data Protection Regulation (GDPR) can be challenging due to several obligations imposed. For example, before deploying such a system, employers have to inform workers about the processing of their personal data in a clear and understandable manner, providing details about the purpose of processing, data retention periods, and the rights of workers with respect to the processing. They also have to obtain workers' consent before the system processes their personal data. The system should also include functionalities to facilitate the enforcement of workers' rights, which include for example the right to access their personal data, the right to rectify inaccurate or incomplete data, and the right to delete data when it is no longer needed for the purposes it was collected for.

Moreover, the privacy monitoring system should implement appropriate technical and organizational measures like encryption and access control to secure personal data against unauthorized access, data breaches, and other security risks.

Another critical issue is related to the integration of the privacy monitoring system with the safety enforcement systems. Due to controls performed by the privacy monitoring system on the information flows, communication latency and processing speeds may increase, negatively impacting the detection of safety-critical situations and the productivity of the smart manufacturing line. Moreover, in order to ensure workforce acceptance, the integration of the privacy monitoring system should be seamless and avoid to disrupt working practices and activities. The acceptance level is also likely to increase if workers are given control on the usage of the data and the advantage of using the system to preserve their right to privacy is communicated to the workers in advance [80].

B. APPLYING CONTEXTUAL INTEGRITY TO MONITOR INFORMATION FLOWS

To allow the use of safety technologies while preserving workers’ privacy, a privacy monitoring system based on the notion of privacy as *contextual integrity* proposed by Nissenbaum [81] was designed. She argues that privacy is not about confidentiality or user control but rather about the appropriateness of information flow within a specific context. Each context has a distinct set of *norms* which govern the flow of information within that context. Norms consist of four key components: *contexts*, *actors*, *attributes*, and *transmission principles*. Contexts represent social contexts (for example, professional or emergency contexts) in which the information is disseminated; actors represent users or software agents involved in the information flow, and their role within the flow e.g, sender, receiver, or information subjects; attributes indicate the type or nature of information, while transmission principles impose conditions on the flow of information from party to party in a given context. For example, a doctor (*sender*) could share sensitive information (*information type*) about a patient (*information subject*) with other doctors (*receivers*) without the explicit consent of the patient (*transmission principle*) in an emergency situation (*context*).

Implementing context-aware privacy monitoring systems presents several technical challenges. One of the first challenges is related to how to model the context in which the information is shared. Another major challenge is accurately identifying the flows of information that involves personal data [82] and the context in which data is being collected and used, as the smart manufacturing environments are complex and dynamic environments where context can rapidly change. This requires advanced algorithms and machine learning models capable of real-time analysis and decision-making. Another challenge is ensuring data quality and consistency across heterogeneous sources, as context-aware privacy monitoring system often integrate data from various safety enforcement systems, sensors, and devices with different formats and standards. The interoperability can be achieved by using a machine-readable format like JSON to share information among the safety enforcement systems, and the context-aware privacy monitoring system. Additionally, the context-aware privacy monitoring system architecture should be robust and scalable to process large volumes of data while maintaining low latency and high reliability. Furthermore, maintaining the security of data when it is shared among the different components of the system and when they are processed is essential to protect the data from potential threats. This requires robust encryption protocols and access control mechanisms to protect the information flows.

Next sections present the context model and the system architecture.

1) MODELING CONTEXT AND PRIVACY NORMS

To ensure privacy, sensitive data should only be released in particular contexts that must be first defined and recognized.




FIGURE 16. The context dimension tree of this project’s use cases.

Contextual features are the different perspectives of the reality in examination, useful to enforce privacy in the data flow. These features are represented through a hierarchical model, the Context Dimension Tree (CDT) [83]. The CDT allows us to specify the points of view of the context with different granularity levels, to be able to restrict and widen access to subsets of sensitive data only in those contexts where safety should take priority over privacy (e.g. adverse events like a medical emergency or a man on the ground, or an environmental event such as a fire emergency or a human-robot collision, and gatherings).

Figure 16 shows a simplified CDT of the ICE lab: black nodes represent the contextual features (i.e., spatio-temporal information, the user, and the situation), whereas white nodes represent the related values. It can be noted that the Emergency situation can be specified with more detail w.r.t. the type of emergency (e.g., Fire emergency, Medical emergency, etc.). Double circle nodes represent parameters whose values are provided by sensors at run-time. Given a CDT \mathcal{C} , a context C is a subtree of \mathcal{C} , where at most one white node is selected for each dimension, e.g., $C_1 = \langle T_Room = Room(\$A) \wedge T_Emerg = FireEmergency \rangle$ is the context representing a Fire Emergency in Room A of ICE lab (assuming for each node the unique name assumption holds, i.e., there is no need to include the path to the node to characterize it). Each context is associated with a condition that allows to determine whether it is active. For example, C_1 is active when a fire sensor in Room A is on. Another, and more general, possible context is $C_2 = \langle Location = ICE LAB \wedge Situation = Emergency \rangle$, which is active when a generic emergency has been declared, for example, by pressing a button on a totem.

To represent the personal data that can be disclosed in an active context, the notion of *privacy norms* is introduced. A privacy norm is a first-order logic formula defined as follows:

$$\phi(sndr, rcvr, info, sbj) \rightarrow flow(sndr, rcvr, info, sbj) \quad (1)$$

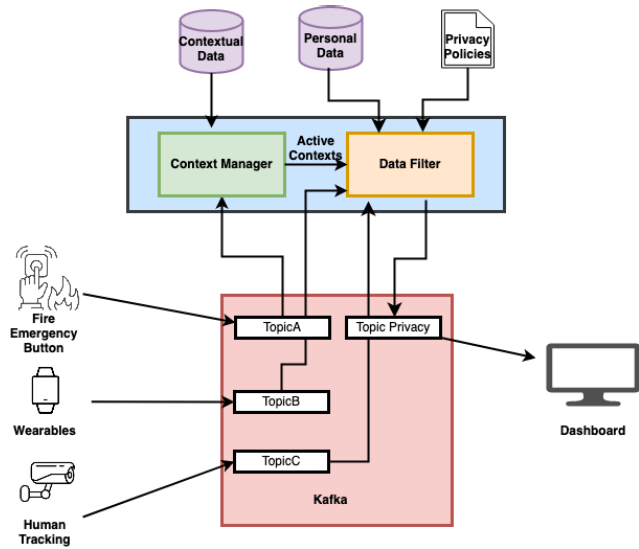


FIGURE 17. Context-based publish-subscribe architecture for privacy management.

where $\phi(sndr, rcvr, info, sbj)$ is a formula representing the transmission principle, and $flowi(sndr, rcvr, info, sbj)$ is a predicate that represents a flow of information. For example, the norm that specifies that the location of a worker should be disclosed to the emergency responder if the worker is in an emergency situation is formalized as follows:

$$\begin{aligned}
 &inEmergencyContext(humantracking, rcvr, worker) \wedge \\
 &publish(humantracking, location) \wedge \\
 &own(worker, location) \wedge hasrole(rcvr, responder) \rightarrow \\
 &flow(humantracking, responder, location, worker) \quad (2)
 \end{aligned}$$

2) SYSTEM COMPONENTS

The components of the context-aware privacy monitoring system are presented in Figure 17. The system has been implemented in Python on top of the data collection architecture based on Kafka. The system consists of three main components, i.e., the Context Manager, the Data Filter, and Dashboard. The Context Manager continuously pulls data from the Kafka topics where the systems like the visual system, human tracking, and sensors deployed in the ICE lab publish collected data. Then, it evaluates which contexts are active among a set of possible contexts representing the environments and the possible situations workers and robots are acting in. Once active contexts have been identified, for each of them, the Data Filter executes a procedure that enforces the privacy norm applicable to that context. The procedure pulls from Kafka topics and an internal database (e.g. a worker’s medical data) that are allowed to be published in that context based on the privacy norm. Then, the Dashboard presents a list of active contexts and the associated information that can be displayed within those contexts. The following example illustrates how the Data Filter selects and visualizes different information.




FIGURE 18. Dashboard - Fire emergency context active.

Example 1. When there is a Regular situation, the Dashboard will only display the number of workers present in the ICE lab. If the context C_1 is active, the Data Filter component accesses the location of the workers published by the human tracking system on a dedicated Kafka topic and checks based on their location whether they are in proximity to the assembly point. If some workers are not near the assembly point, it publishes the location of these workers on another Kafka topic. Then, the Dashboard will pull the workers’ locations from the Kafka topic and display them so that the workers can be rescued, as shown in Figure 18.

VI. CONCLUSION AND FUTURE WORK

In this work, we presented a case study for Safety and Privacy enhanced Industry 4.0, showing a fully operational automated manufacturing production environment, with high-end devices. We employed many different technologies to solve Industry 4.0 problems with state-of-the-art novel approaches (and some of them are also applicable to Industry 5.0). We defined six use cases (presented in Section III), representing the possible threats to the safety and privacy of workers in an Industry 4.0 scenario. We implemented a context-aware privacy enforcement method to address the trade-off between a technology used to increase the safety of a worker, and its impact on the worker’s privacy. Some technologies were re-designed to accommodate for a change in demands (i.e., the face mask detector that is now less critical but can be implemented as a general workers’ PPE detector), while others are still being explored and updated following recent state-of-the-art approaches.

In this sense, the ICE laboratory is just a first step in optimizing Industry 4.0 (and later 5.0), taking into account real-world issues and problems such as privacy and safety, while still researching different aspects of the technologies employed.

Future work involves every technology presented, by trying to improve the existing methods, increasing generalization capabilities, extending the interaction using natural language search [84], and making new promising technologies a reality.

REFERENCES

- [1] Y. Lu, “Industry 4.0: A survey on technologies, applications and open research issues,” *J. Ind. Inf. Integr.*, vol. 6, pp. 1–10, Jun. 2017.
- [2] M. Sanchez, E. Exposito, and J. Aguilar, “Industry 4.0: Survey from a system integration perspective,” *Int. J. Comput. Integr. Manuf.*, vol. 33, nos. 10–11, pp. 1017–1041, Nov. 2020.
- [3] S. Munirathinam, “Industry 4.0: Industrial Internet of Things (IIOT),” in *Advances in Computers*, vol. 117. Amsterdam, The Netherlands: Elsevier, 2020, pp. 129–164.

- [4] M. Soori, B. Arezoo, and R. Dastres, "Internet of Things for smart factories in industry 4.0, a review," *Internet Things Cyber-Phys. Syst.*, vol. 3, pp. 192–204, Jan. 2023.
- [5] C. Cimini, F. Pirola, R. Pinto, and S. Cavalieri, "A human-in-the-loop manufacturing control architecture for the next generation of production systems," *J. Manuf. Syst.*, vol. 54, pp. 258–271, Jan. 2020.
- [6] C. J. Turner, R. Ma, J. Chen, and J. Oyekan, "Human in the loop: Industry 4.0 technologies and scenarios for worker mediation of automated manufacturing," *IEEE Access*, vol. 9, pp. 103950–103966, 2021.
- [7] Y. Li and F. W. Guldenmund, "Safety management systems: A broad overview of the literature," *Saf. Sci.*, vol. 103, pp. 94–123, Mar. 2018. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0925753517309463>
- [8] A. R. Paiva and A. Tewari, "Methodology for testing and evaluation of safety analytics approaches," *Saf. Sci.*, vol. 152, Aug. 2022, Art. no. 105737. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0925753522000777>
- [9] L. Zanella, B. Liberatori, W. Menapace, F. Poesi, Y. Wang, and E. Ricci, "Delving into CLIP latent space for video anomaly recognition," *Comput. Vis. Image Understand.*, vol. 249, Dec. 2024, Art. no. 104163.
- [10] S. Gaiardelli, S. Spellini, M. Panato, M. Lora, and F. Fummi, "A software architecture to control service-oriented manufacturing systems," in *Proc. Design, Autom. Test Eur. Conf. Exhib. (DATE)*, Mar. 2022, pp. 40–43.
- [11] F. Tao, B. Xiao, Q. Qi, J. Cheng, and P. Ji, "Digital twin modeling," *J. Manuf. Syst.*, vol. 64, pp. 372–389, Jul. 2022. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0278612522001108>
- [12] L. Capogrosso, F. Cunico, D. S. Cheng, F. Fummi, and M. Cristani, "A machine learning-oriented survey on tiny machine learning," *IEEE Access*, vol. 12, pp. 23406–23426, 2024.
- [13] F. Cunico, L. Capogrosso, F. Setti, D. Carra, F. Fummi, and M. Cristani, "I-SPLIT: Deep network interpretability for split computing," in *Proc. 26th Int. Conf. Pattern Recognit. (ICPR)*, Aug. 2022, pp. 2575–2581.
- [14] L. Capogrosso, F. Cunico, M. Lora, M. Cristani, F. Fummi, and D. Quaglia, "Split-Et-impera: A framework for the design of distributed deep learning applications," in *Proc. 26th Int. Symp. Design Diag. Electron. Circuits Syst. (DDECS)*, May 2023, pp. 39–44.
- [15] L. Capogrosso, E. Fraccaroli, S. Chakraborty, F. Fummi, and M. Cristani, "MTL-split: Multi-task learning for edge devices using split computing," 2024, *arXiv:2407.05982*.
- [16] L. Capogrosso, E. Fraccaroli, G. Petrozziello, F. Setti, S. Chakraborty, F. Fummi, and M. Cristani, "Enhancing split computing and early exit applications through predefined sparsity," in *Proc. Forum Specification Design Lang. (FDL)*, Sep. 2024, pp. 1–8.
- [17] M. Boldo, D. Carra, D. Quaglia, and N. Bombieri, "A dynamic and collaborative deep inference framework for human motion analysis in telemedicine," in *Proc. IEEE Int. Conf. Edge Comput. Commun. (EDGE)*, vol. 225, Jul. 2023, pp. 221–226.
- [18] F. Cunico, F. Girella, A. Avogaro, M. Emporio, A. Giachetti, and M. Cristani, "OO-dMVM: A deep multi-view multi-task classification framework for real-time 3D hand gesture classification and segmentation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. Workshops (CVPRW)*, Jun. 2023, pp. 2745–2754.
- [19] M. Emporio, A. Caputo, A. Giachetti, M. Cristani, G. Borghi, A. D'Eusano, M.-Q. Le, H.-D. Nguyen, M.-T. Tran, and F. Ambellan, "SHREC 2022 track on online detection of heterogeneous gestures," *Comput. Graph.*, vol. 107, pp. 241–251, Oct. 2022.
- [20] P. Forouzandeh, K. O'Dowd, and S. C. Pillai, "Face masks and respirators in the fight against the COVID-19 pandemic: An overview of the standards and testing methods," *Saf. Sci.*, vol. 133, Jan. 2021, Art. no. 104995.
- [21] E. K. J. Dzisi and O. A. Dei, "Adherence to social distancing and wearing of masks within public transportation during the COVID 19 pandemic," *Transp. Res. Interdiscipl. Perspect.*, vol. 7, Sep. 2020, Art. no. 100191.
- [22] A. Sampieri, G. M. D. di Melendugno, A. Avogaro, F. Cunico, F. Setti, G. Skenderi, M. Cristani, and F. Galasso, "Pose forecasting in industrial human-robot collaboration," in *Proc. Eur. Conf. Comput. Vis. Berlin, Germany: Springer*, 2022, pp. 51–69.
- [23] M. Boldo, N. Bombieri, S. Centomo, M. De Marchi, F. Demrozi, G. Pravadelli, D. Quaglia, and C. Turetta, "Integrating wearable and camera based monitoring in the digital twin for safety assessment in the Industry 4.0 era," in *Proc. Int. Symp. Leveraging Appl. Formal Methods*. Cham, Switzerland: Springer, 2022, pp. 184–194.
- [24] M. Boldo, N. Bombieri, M. De Marchi, L. Geretti, S. Germiniani, and G. Pravadelli, "Risk assessment and prediction in human-robot interaction through assertion mining and pose estimation," in *Proc. IEEE 23rd Latin Amer. Test Symp. (LATS)*, Sep. 2022, pp. 1–5.
- [25] A. Avogaro, F. Cunico, B. Rosenhahn, and F. Setti, "Markerless human pose estimation for biomedical applications: A survey," *Frontiers Comput. Sci.*, vol. 5, Jul. 2023, Art. no. 1153160.
- [26] S. Bragança, E. Costa, I. Castellucci, and P. M. Arezes, "A brief overview of the use of collaborative robots in Industry 4.0: Human role and safety," in *Occupational and Environmental Safety and Health*. Cham, Switzerland: Springer, 2019, pp. 641–650.
- [27] R. Galin and R. Meshcheryakov, "Automation and robotics in the context of industry 4.0: The shift to collaborative robots," *IOP Conf. Ser., Mater. Sci. Eng.*, vol. 537, no. 3, May 2019, Art. no. 032073.
- [28] F. Sherwani, M. M. Asad, and B. S. K. K. Ibrahim, "Collaborative robots and industrial revolution 4.0 (IR 4.0)," in *Proc. Int. Conf. Emerg. Trends Smart Technol. (ICETST)*, Mar. 2020, pp. 1–5.
- [29] A. Borboni, K. V. V. Reddy, I. Elamvazuthi, M. S. Al-Quraishi, E. Natarajan, and S. S. A. Ali, "The expanding role of artificial intelligence in collaborative robots for industrial applications: A systematic review of recent works," *Machines*, vol. 11, no. 1, p. 111, Jan. 2023.
- [30] B. Yalçinkaya, M. S. Couceiro, S. P. Soares, and A. Valente, "Human-aware collaborative robots in the wild: Coping with uncertainty in activity recognition," *Sensors*, vol. 23, no. 7, p. 3388, Mar. 2023.
- [31] *Robots and Robotic Devices—Safety Requirements for Industrial Robots—Part 2: Robot Systems and Integration*, Int. Org. Standardization, Geneva, Switzerland, Jul. 2011.
- [32] A. Toiari, F. Cunico, F. Taioli, A. Caputo, G. Menegaz, A. Giachetti, G. M. Farinella, and M. Cristani, "Scene-path: Capturing the visual selective attention of people towards scene elements," in *Proc. Int. Conf. Image Anal. Process*. Cham, Switzerland: Springer, 2023, pp. 352–363.
- [33] S. Germiniani and G. Pravadelli, "HARM: A hint-based assertion miner," *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, vol. 41, no. 11, pp. 4277–4288, Nov. 2022.
- [34] E. Cinquetti, I. Siviero, S. F. Storti, and G. Menegaz, "Enhancing safety in Industry 4.0: The use of passive brain computer interfaces for vigilance monitoring," in *Proc. 8th Nat. Congr. Bioeng.*, 2023, pp. 774–777.
- [35] F. Demrozi, C. Turetta, P. H. Kindt, F. Chiarani, R. Bacchin, N. Valè, F. Pascucci, P. Cesari, N. Smania, S. Tamburini, and G. Pravadelli, "A low-cost wireless body area network for human activity recognition in healthy life and medical applications," *IEEE Trans. Emerg. Topics Comput.*, vol. 11, no. 4, pp. 839–850, 2023.
- [36] M. De Marchi, C. Turetta, G. Pravadelli, and N. Bombieri, "Real-time multi-person identification and tracking via HPE and IMU data fusion," in *Proc. Design, Autom. Test Eur. Conf. Exhib. (DATE)*, vol. 20, Mar. 2024, pp. 1–6.
- [37] I. Siviero, L. Brusini, G. Menegaz, and S. F. Storti, "Motor-imagery EEG signal decoding using multichannel-empirical wavelet transform for brain computer interfaces," in *Proc. IEEE-EMBS Int. Conf. Biomed. Health Informat. (BHI)*, vol. 16, Sep. 2022, pp. 1–4.
- [38] I. Siviero, G. Menegaz, and S. F. Storti, "Functional connectivity and feature fusion enhance multiclass motor-imagery brain-computer interface performance," *Sensors*, vol. 23, no. 17, p. 7520, Aug. 2023.
- [39] L. Brusini, F. Stival, F. Setti, E. Menegatti, G. Menegaz, and S. F. Storti, "A systematic review on motor-imagery brain-connectivity-based computer interfaces," *IEEE Trans. Human-Machine Syst.*, vol. 51, no. 6, pp. 725–733, Dec. 2021.
- [40] F. Cunico and M. Cristani, "Multi-camera industrial open-set person identification and tracking," 2024, *arXiv:2409.03879*.
- [41] L. Geretti, S. Centomo, M. Boldo, E. Martini, N. Bombieri, D. Quaglia, and T. Villa, "Process-driven collision prediction in human-robot work environments," in *Proc. IEEE 27th Int. Conf. Emerg. Technol. Factory Autom. (ETFA)*, Sep. 2022, pp. 1–8.
- [42] F. Cunico, A. Toiari, and M. Cristani, "A masked face classification benchmark on low-resolution surveillance images," in *Proc. Int. Conf. Pattern Recognit*. Cham, Switzerland: Springer, 2023, pp. 49–63.
- [43] F. Cunico, L. Capogrosso, A. Castellini, F. Setti, P. Pluchino, F. Zordan, V. Santus, A. Spagnolli, S. Cordibella, G. Gennari, M. Borgo, A. Sozza, S. Troiano, R. Flor, A. Zanella, A. Farinelli, L. Gamberini, and M. Cristani, "The post-pandemic effects on IoT for safety: The safe place project," in *Proc. Design, Autom. Test Eur. Conf. Exhib. (DATE)*, Apr. 2023, pp. 1–4.
- [44] Q. A. Ahmed and Z. A. Memish, "The cancellation of mass gatherings (MGs)? Decision making in the time of COVID-19," *Travel Med. Infectious Disease*, vol. 34, Mar. 2020, Art. no. 101631.

- [45] A. Ahammer, M. Halla, and M. Lackner, "Mass gatherings contributed to early COVID-19 mortality: Evidence from U.S. sports," *Contemp. Econ. Policy*, vol. 41, no. 3, pp. 471–488, Jul. 2023.
- [46] C. Martella, J. Li, C. Conrado, and A. Vermeeren, "On current crowd management practices and the need for increased situation awareness, prediction, and intervention," *Saf. Sci.*, vol. 91, pp. 381–393, Jan. 2017.
- [47] F. Becattini, G. Becchi, A. Ferracani, A. D. Bimbo, L. L. Presti, G. Mazzola, M. L. Cascia, F. Cunico, A. Toiari, M. Cristani, A. Greco, A. Saggese, and M. Vento, "I-MALL an effective framework for personalized visits. Improving the customer experience in stores," in *Proc. 1st Workshop Multimedia Comput. towards Fashion Recommendation*, vol. 28, Oct. 2022, pp. 11–19.
- [48] Y. Zhang, P. Sun, Y. Jiang, D. Yu, F. Weng, Z. Yuan, P. Luo, W. Liu, and X. Wang, "ByteTrack: Multi-object tracking by associating every detection box," in *Proc. Eur. Conf. Comput. Vis. (ECCV)*. Cham, Switzerland: Springer, 2022, pp. 1–21.
- [49] M. Wiczorek, B. Rychalska, and J. Dabrowski, "On the unreasonable effectiveness of centroids in image retrieval," in *Proc. Int. Conf. Neural Inf. Process.* Cham, Switzerland: Springer, 2021, pp. 212–223.
- [50] L. Zheng, L. Shen, L. Tian, S. Wang, J. Wang, and Q. Tian, "Scalable person re-identification: A benchmark," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Dec. 2015, pp. 1116–1124.
- [51] Q. Leng, M. Ye, and Q. Tian, "A survey of open-world person re-identification," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 30, no. 4, pp. 1092–1108, Apr. 2020.
- [52] *Statement on the Fifteenth Meeting of the IHR (2005) Emergency Committee on the COVID-19 Pandemic, 5 May 2023*, World Health Org., Geneva, Switzerland, 2021.
- [53] M. Capuzzo, A. Zanella, M. Zuccotto, F. Cunico, M. Cristani, A. Castellini, A. Farinelli, and L. Gamberini, "IoT systems for healthy and safe life environments," in *Proc. IEEE 7th Forum Res. Technol. Soc. Ind. Innov. (RTSI)*, Aug. 2022, pp. 31–37.
- [54] A. K. Inkulu, M. V. A. R. Bahubalendruni, and A. Dara, "Challenges and opportunities in human robot collaboration context of Industry 4.0—A state of the art review," *Ind. Robot, Int. J. Robot. Res. Appl.*, vol. 49, no. 2, pp. 226–239, Feb. 2022.
- [55] K. A. Demir, G. Döven, and B. Sezen, "Industry 5.0 and human-robot co-working," *Procedia Comput. Sci.*, vol. 158, pp. 688–695, Jan. 2019.
- [56] J. Alves, T. M. Lima, and P. D. Gaspar, "Is Industry 5.0 a human-centred approach? A systematic review," *Processes*, vol. 11, no. 1, p. 193, Jan. 2023.
- [57] M. Boldo, M. De Marchi, E. Martini, S. Aldegheri, D. Quaglia, F. Fummi, and N. Bombieri, "Real-time multi-camera 3D human pose estimation at the edge for industrial applications," *Expert Syst. Appl.*, vol. 252, Oct. 2024, Art. no. 124089.
- [58] A. Avogaro, A. Toiari, F. Cunico, X. Xu, H. Dafas, A. Vinciarelli, E. Li, and M. Cristani, "Exploring 3D human pose estimation and forecasting from the robot's perspective: The HARPER dataset," 2024, *arXiv:2403.14447*.
- [59] Z. Cao, T. Simon, S.-E. Wei, and Y. Sheikh, "Realtime multi-person 2D pose estimation using part affinity fields," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 1302–1310.
- [60] C. Buizza and Y. Demiris, "Rotational adjoint methods for learning-free 3D human pose estimation from IMU data," in *Proc. 25th Int. Conf. Pattern Recognit. (ICPR)*, Jan. 2021, pp. 7868–7875.
- [61] D. L. Mills, "Network time protocol (NTP)," NTP, Tech. Rep., 1985.
- [62] W. Kabsch, "A solution for the best rotation to relate two sets of vectors," *Acta Crystallographica A*, vol. 32, no. 5, pp. 922–923, Sep. 1976.
- [63] C. Maureen, B. Laurent, and L. Fabien, *Brain-Computer Interfaces I: Methods and Perspectives*. Hoboken, NJ, USA: Wiley, 2016.
- [64] M. Sebastiani, G. Di Flumeri, P. Aricò, N. Sciaraffa, F. Babiloni, and G. Borghini, "Neurophysiological vigilance characterisation and assessment: Laboratory and realistic validations involving professional air traffic controllers," *Brain Sci.*, vol. 10, no. 1, p. 48, Jan. 2020.
- [65] A. B. Geva and D. H. Kerem, "Forecasting generalized epileptic seizures from the EEG signal by wavelet analysis and dynamic unsupervised fuzzy clustering," *IEEE Trans. Biomed. Eng.*, vol. 45, no. 10, pp. 1205–1216, Jul. 1998.
- [66] W. W. Wei, *Multivariate Time Series Analysis and Applications*. Hoboken, NJ, USA: Wiley, 2018.
- [67] R. C. Staudemeyer and E. Rothstein Morris, "Understanding LSTM—A tutorial into long short-term memory recurrent neural networks," 2019, *arXiv:1909.09586*.
- [68] R. Wan, S. Mei, J. Wang, M. Liu, and F. Yang, "Multivariate temporal convolutional network: A deep neural networks approach for multivariate time series forecasting," *Electronics*, vol. 8, no. 8, p. 876, Aug. 2019.
- [69] K. Li, J. Daniels, C. Liu, P. Herrero, and P. Georgiou, "Convolutional recurrent neural networks for glucose prediction," *IEEE J. Biomed. Health Informat.*, vol. 24, no. 2, pp. 603–613, Feb. 2020.
- [70] R. E. Stirling, M. I. Maturana, P. J. Karoly, E. S. Nurse, K. McCutcheon, D. B. Grayden, S. G. Ringo, J. M. Heasman, R. J. Hoare, A. Lai, W. D'Souza, U. Seneviratne, L. Seiderer, K. J. McLean, K. J. Bulluss, M. Murphy, B. H. Brinkmann, M. P. Richardson, D. R. Freestone, and M. J. Cook, "Seizure forecasting using a novel sub-scalp ultra-long term EEG monitoring system," *Frontiers Neurol.*, vol. 12, Aug. 2021, Art. no. 713794.
- [71] A. Craik, Y. He, and J. L. Contreras-Vidal, "Deep learning for electroencephalogram (EEG) classification tasks: A review," *J. Neural Eng.*, vol. 16, no. 3, Jun. 2019, Art. no. 031001.
- [72] I. Mehmood, H. Li, Y. Qarout, W. Umer, S. Anwer, H. Wu, M. Hussain, and M. F. Antwi-Afari, "Deep learning-based construction equipment operators' mental fatigue classification using wearable EEG sensor data," *Adv. Eng. Informat.*, vol. 56, Apr. 2023, Art. no. 101978.
- [73] M. U. Iqbal, B. Srinivasan, and R. Srinivasan, "Multi-class classification of control room operators' cognitive workload using the fusion of eye-tracking and electroencephalography," *Comput. Chem. Eng.*, vol. 181, Feb. 2024, Art. no. 108526.
- [74] J. Y. Lee and S. H. Seo, "EEG-based emotion recognition using deep learning model for workers safety," *Nanotechnol. Perceptions*, vol. 20, pp. 332–345, 2024.
- [75] Y. Yin, P. Zheng, C. Li, and L. Wang, "A state-of-the-art survey on augmented reality-assisted digital twin for futuristic human-centric industry transformation," *Robot. Comput.-Integr. Manuf.*, vol. 81, Jun. 2023, Art. no. 102515.
- [76] L. Shi, Y. Zhang, J. Cheng, and H. Lu, "Decoupled spatial-temporal attention network for skeleton-based action-gesture recognition," in *Proc. Asian Conf. Comput. Vis.*, 2020, pp. 1–16.
- [77] F. Yang, Y. Wu, S. Sakti, and S. Nakamura, "Make skeleton-based action recognition model smaller, faster and better," in *Proc. ACM Multimedia Asia*, Dec. 2019, pp. 1–6.
- [78] D. Pintani, A. Caputo, D. Mendes, and A. Giachetti, "CIDER: Collaborative interior design in extended reality," in *Proc. 15th Biannual Conf. Italian SIGCHI Chapter*, Sep. 2023, pp. 1–11.
- [79] D. J. Solove, "A taxonomy of privacy," *Univ. PA Law Rev.*, vol. 154, no. 3, pp. 477–560, Jan. 2006.
- [80] M. Abraham, C. Niessen, C. Schnabel, K. Lorek, V. Grimm, K. Möslein, and M. Wrede, "Electronic monitoring at work: The role of attitudes, functions, and perceived control for the acceptance of tracking technologies," *Hum. Resource Manage. J.*, vol. 29, no. 4, pp. 657–675, Nov. 2019.
- [81] H. Nissenbaum, "Privacy as contextual integrity," *Washington Law Rev.*, vol. 79, no. 1, pp. 119–157, 2004.
- [82] E. Papadogiannaki and S. Ioannidis, "A survey on encrypted network traffic analysis applications, techniques, and countermeasures," *ACM Comput. Surv.*, vol. 54, no. 6, pp. 1–35, Jul. 2021.
- [83] C. Bolchini, E. Quintarelli, and L. Tanca, "CARVE: Context-aware automatic view definition over relational databases," *Inf. Syst.*, vol. 38, no. 1, pp. 45–67, Mar. 2013.
- [84] F. Taioli, F. Cunico, F. Girella, R. Bologna, A. Farinelli, and M. Cristani, "Language-enhanced RNR-map: Querying renderable neural radiance field maps with natural language," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. Workshops (ICCVW)*, vol. 33, Oct. 2023, pp. 4671–4676.



FEDERICO CUNICO (Graduate Student Member, IEEE) received the master's degree in computer science and engineering from the University of Verona, in 2019, where he is currently pursuing the Ph.D. degree in computer science, under the supervision of Prof. Marco Cristani. His main research interests include deep learning and computer vision, with a focus on people analysis.



STEFANO ALDEGHERI received the Ph.D. degree in computer science from the University of Verona, Verona, Italy, in 2020. Currently, he is a Senior Software Developer and a Researcher with the Department of Engineering for Innovation Medicine, University of Verona. His research interests include the development and adaptation of computer vision applications for embedded systems, with a particular emphasis on automatic code synthesis, system-level optimization, and multi/many-core acceleration.



ARIEL CAPUTO received the Ph.D. degree in computer science from the University of Verona, Verona, Italy, in 2019. She is currently a Research Assistant with the Department of Engineering for Innovation Medicine. Her research interests include variety of HCI topics, including interaction design in mixed-reality systems and computer graphics.



ANDREA AVOGARO (Associate Member, IEEE) received the bachelor's and master's degrees in computer science from the University of Verona, in 2019 and 2021, respectively. He is currently pursuing the Ph.D. degree in computer science supervised by Prof. Marco Cristani. He specializes in human pose forecasting and fashion sales forecasting. His work involves the study of deep learning for predicting human body poses and enhancing sales predictions in the fashion industry.



DAMIANO CARRA received the Laurea degree in telecommunication engineering from the Politecnico di Milano and the Ph.D. degree in computer science from the University of Trento. He is currently an Associate Professor with the Computer Science Department, University of Verona. His research interest includes modeling and performance evaluation of large scale distributed systems.



MICHELE BOLDO (Student Member, IEEE) is currently pursuing the Ph.D. degree with the Department of Computer Science, University of Verona, Italy. His research interests include the development of efficient human pose estimation model for remote monitoring and motion analysis and optimization of embedded software for edge computing.



STEFANO CENTOMO received the Ph.D. degree in computer science from the University of Verona, in 2021. He is currently a Postdoctoral Researcher with the Department of Computer Science, University of Verona. His main research interest includes electronic design automation methodologies for efficient simulation of cyber-physical production systems in the context of smart factory (industry 4.0).



NICOLA BOMBIERI (Member, IEEE) is currently a Full Professor with the Department of Engineering for Innovation Medicine, University of Verona, Italy. He develops embedded and efficient software applications, with a particular emphasis on addressing multi- and extra-functional constraints, such as performance, power, and energy efficiency. His research interests include parallel and heterogeneous architectures, artificial intelligence at the edge, and parallel programming languages.



DONG SEON CHENG (Member, IEEE) received the Ph.D. degree in computer science from the University of Verona, Italy, in 2008. Currently, he is a Research Associate with the University of Verona. From 2012 to 2017, he was an Assistant Professor with the Department of CSE, HUFS, South Korea. His main research interests include machine learning for computer vision and pattern recognition.



LUIGI CAPOGROSSO (Student Member, IEEE) received the B.Sc. degree in computer science and the M.Sc. degree in computer engineering for robotics and smart industry from the University of Verona, in 2019 and 2021, respectively. He is currently pursuing the joint Ph.D. degree with the National Program in Artificial Intelligence, University of Verona, in collaboration with the Intelligo ML Laboratories, Polytechnic of Turin, under the supervision of Prof. Marco Cristani and Prof. Franco Fummi. His research interests include efficient deep learning and representation learning.



ETTORE CINQUETTI received the Graduate degree (Hons.) in computer science and engineering from the University of Verona, where he is currently pursuing the Ph.D. degree. He is studying brain-computer interfaces applications to improve safety conditions in critical areas, with a focus on industrial settings.



MARCO CRISTANI (Member, IEEE) is currently a Full Professor with the Computer Science Department, University of Verona; an Associate Member of the National Research Council (CNR); and an External Collaborator of Italian Institute of Technology (IIT). He is also the Managing Director of the Computer Science Park, a Technology Transfer Center, University of Verona. His main research interests include statistical pattern recognition and computer vision, mainly in deep learning and generative modeling, with application to social signal processing and fashion modeling. On these topics, he has published more than 180 papers, including two edited volumes, 46 international journal articles, 126 conference papers, and 13 book chapters. He has organized 11 international workshops, co-founded a spin-off company, Humatics, dealing with e-commerce for fashion. He is or has been a PI of several national and international projects, including PRIN and H2020 projects. He is a member of ACM and IAPR. He is a member of the Editorial Board of the *Pattern Recognition* and *Pattern Recognition Letters*.



MIRCO DE MARCHI is currently pursuing the Ph.D. degree with the Department of Computer Science, University of Verona, Italy. His research interests include development of efficient human pose estimation model for remote monitoring and motion analysis and optimization of embedded software for edge computing.



FLORENC DEMROZI (Member, IEEE) received the B.Sc. and M.Sc. degrees in computer science and engineering and the Ph.D. degree in computer science from the University of Verona, Italy, in 2014, 2016, and 2020, respectively. Currently, he is an Associate Professor of biomedical engineering with the Department of Electrical Engineering and the Department of Computer Science, University of Stavanger, Norway. His research interests include the Artificial Intelligence of Things (AIoT), human activity recognition (HAR), active and assisted living (AAL), and the Internet of Medical Things (IoMT).



MARCO EMPORIO (Student Member, IEEE) received the Graduate degree in computer engineering from the University of Verona, in 2021, where he is currently pursuing the Ph.D. degree with the Department of Computer Science. His research interest includes human-computer interaction in AR and VR.



FRANCO FUMMI (Member, IEEE) received the Laurea degree in electronic engineering from the Polytechnic of Milan, in 1990, and the Ph.D. degree in electronic and communication engineering from the Polytechnic of Milan, in 1994. In 1993, he was a Research Assistant with the Department of Computer Science, University of Victoria. In 1996, he received the position of an Assistant Professor in computer science with the Dipartimento di Elettronica e Informazione, Polytechnic of Milan, until October 1998. In July 1998, he received the position of an Associate Professor in computer architecture with the Computer Science Department, University of Verona, where he has been a Full Professor of computer architecture, since March 2001, before with the Computer Science Department and then the Department of Engineering for Innovation Medicine. He is also leading the Cyber-Physical and IoT Systems Design (CISD) Group, University of Verona, currently composed of more than 20 people and working on hardware description languages and electronic design automation methodologies for modeling, verification, testing, and optimization of cyber-physical systems. Since 2018, he has been the Project Manager of the Industrial Computer Engineering (ICE) Laboratory, University of Verona: a facility serving as a Technological Demonstrator and Research Laboratory, functional to rethink industrial processes, such as additive and subtractive manufacturing, quality control, assembly, and parts storage. He is also a Co-Founder of two spin-off companies: EDALab, focused on networked embedded systems design and the automation control software company FACTORYAL.



LUCA GERETTI received the Electrical Engineering degree and the Ph.D. degree in computer engineering from the University of Udine, Italy, in 2019. Then he moved to the University of Verona, Italy, working on formal verification with particular application to hybrid systems. Since 2020, he has been the Head of the nonlinear systems track of the international ARCH competition for formal verification tools. Currently, he is an Assistant Professor at the University of Verona. His research focus being on mobile computing and neural networks.



SAMUELE GERMINIANI (Member, IEEE) received the Ph.D. degree in computer science from the University of Verona, Italy, in 2023. He is currently a Postdoctoral Research Fellow with the Department of Engineering for Innovation Medicine, University of Verona. His main research interests include semi-formal verification and embedded security of cyber-physical systems.



ANDREA GIACHETTI is currently a Professor of information processing systems with the Department of Engineering for Innovation Medicine, University of Verona. His research interests include computer graphics, computer vision, and human-computer interaction. In these fields, he has published more than 150 articles in international journals and conferences and has been a coordinator or a key personnel of several national and European projects.



FEDERICO GIRELLA received the master's degree in computer science and engineering from the University of Verona, Verona, Italy, in 2022, where he is currently pursuing the Ph.D. degree in artificial intelligence. His main research interest includes vision and language for AI using deep neural networks.



MARCO PANATO received the B.S. and M.E. degrees in computer science and engineering from the University of Verona, Italy, in 2016 and 2018, respectively. He is currently a Software Engineer and a Co-Founder of Factorial S.r.l., San Giovanni Lupatoto, Verona, Italy, a spinoff of the University of Verona. Previously, he was a Research Assistant with the University of Verona.



ENRICO MARTINI (Member, IEEE) is currently pursuing the Ph.D. degree with the Department of Computer Science, University of Verona, Italy. His research interests include the development of efficient human pose estimation model for remote monitoring and motion analysis and optimization of embedded software for edge computing.



GRAZIANO PRAVADELLI (Senior Member, IEEE) received the Ph.D. degree in computer science. He is currently the IFIP 10.5 WG Chair and a Full Professor in information processing systems with the Department of Engineering for Innovation Medicine, University of Verona, Italy, where he leads the IoT4care Research Group. In 2007, he co-founded EDALab s.r.l., an SME working on the design of IoT-based monitoring systems. His main research interests include system-level modeling, simulation and semi-formal verification of embedded systems, and their application to develop IoT-based virtual coaching platforms for people with special needs. In the previous contexts, he collaborated in several national and European projects and has published about 150 papers in international conferences and journals.



GLORIA MENEGAZ (Senior Member, IEEE) received the M.Sc. degree in electronic engineering from the Politecnico di Milano and the Ph.D. degree in applied sciences from the Swiss Federal Institute of Technology (EPFL). She is currently a Professor in bioengineering with the Department of Engineering for Innovation Medicine, University of Verona, Italy, where she founded and leads the BrainAVLab. She received the Rita Levi Montalcini Grant. Her research interests include multimodal and multidimensional biomedical signal processing, with a focus on neuroimaging and explainable artificial intelligence. She has authored more than 130 publications in international peer-reviewed journals and conferences.



ELISA QUINTARELLI received the Ph.D. degree in computer engineering from the Politecnico di Milano, Italy, in 2002. Currently, she is an Associate Professor with the Computer Science Department, University of Verona, Italy. She is the co-author of more than 90 articles on databases, data management, and analysis.



NIEK MUIJS received the Master of Science degree in computer science and in engineering from the University of Technology, Eindhoven, The Netherlands. In his final thesis, he worked on the design and development of a context-based monitoring solution for the industry 4.0 factories that ensures privacy.



ILARIA SIVIERO (Member, IEEE) received the M.S. degree in bioengineering from the University of Padova, in 2021. She is currently pursuing the Ph.D. degree in computer science with the University of Verona. Her main research interests include brain-computer interface, brain connectivity, and biomedical signal processing.



FEDERICA PACI received the Ph.D. degree in information technology from the University of Milan, Italy, in 2008. Currently, she is an Associate Professor with the Computer Science Department, University of Verona, Italy. She is the co-author of more than 80 articles on data security and privacy.




SILVIA FRANCESCA STORTI (Member, IEEE) received the master's degree in electronic engineering from the University of Padova and the Ph.D. degree in neuroscience from the University of Verona, Italy. She is currently an Associate Professor in bioengineering with the Department of Engineering for Innovation Medicine, University of Verona. She is also a member of the BrainAVLab, and focuses her research on neuroengineering, including multimodal functional neuroimaging integration, brain functional connectivity inference, network analysis, and brain-computer interfaces. She has co-authored 62 publications in international peer-reviewed journals, with an H-index of 22 and 1,318 article citations according to Scopus.



CARLO TADIELLO received the B.S. degree in computer science and the M.E. degree in engineering from the University of Verona, Italy, in 2015 and 2018, respectively. He is currently a Software Engineer and the Co-Founder of Factorial S.r.l., San Giovanni Lupatoto, Verona, Italy, a spinoff of the University of Verona. Previously, he was a Research Assistant with the University of Verona.



NICOLA ZANNONE received the master's degree in computer science from the University of Verona, in 2003, and the Ph.D. degree from the University of Trento, in 2007. He is currently an Associate Professor with the Department of Mathematics and Computer Science, Eindhoven University of Technology.




CRISTIAN TURETTA (Graduate Student Member, IEEE) received the B.S. and M.E. degrees in computer science and engineering from the University of Verona, Italy, in 2017 and 2020, respectively, where he is currently pursuing the Ph.D. degree with the Department of Computer Science. He is a member of the IoT4care Research Group, working on ambient assisted living and the Internet of Things with the University of Verona.



TIZIANO VILLA (Member, IEEE) received the Ph.D. degree in EECS from the University of California at Berkeley, in 1995. Since 2006, he has been a Professor with the Department of Computer Science (DI), University of Verona, Italy. He has co-authored three books: *Synthesis of Finite State Machines: Functional Optimization* (Kluwer/Springer, 1997 and Reprint 2010), *Synthesis of Finite State Machines: Logic Optimization* (Kluwer/Springer, 1997 and Reprint 2012),

and *The Unknown Component Problem: Theory and Applications* (Springer, 2012); and co-edited the book *Coordination Control of Distributed Systems* (Springer, 2015). His research interests include formal methods for electronic design automation, including logic synthesis, formal verification, models of computation, discrete-event dynamic systems, and cyber-physical systems.



DAVIDE QUAGLIA received the Ph.D. degree in computer engineering from the Politecnico di Torino, Italy, in 2003. Currently, he is an Associate Professor with the Computer Science Department, University of Verona, Italy. He is the co-author of about 90 articles on networked embedded systems, networked control systems, and cyber-physical systems. He is the Coordinator of the project on which this article is based.

...

Open Access funding provided by 'Università degli Studi di Verona' within the CRUI CARE Agreement