This article has been accepted for publication in Computer. This is the author's version which has not been fully edited and
content may change prior to final publication. Citation information: DOI 10.1109/MC.2026.3671749

Article Type: Description (see below for more detail)

InterwovenXR: A Cyber-Physical-Social Testbed System for Robotics, Digital Twins, and Hyper-Automation

Louis Nisiotis, *University of Central Lancashire, Cyprus, Pyla, 7080, Cyprus*
Marios Raspopoulos, *University of Central Lancashire, Cyprus, Pyla, 7080, Cyprus*

Abstract—Robotics and hyper-automation are moving from specialised tools to core capabilities for designing, operating, and governing complex systems. This article examines the convergence of robotics, digital twins, and AI orchestration within Cyber-Physical-Social Systems (CPSS), where robots act as embodied computing nodes across physical and digital worlds. Digital twins synchronise and coordinate robots, environments, and humans to support simulation, prediction, and decision support. Hyper-automation extends this by enabling adaptive, end-to-end workflows with minimal human intervention. The InterwovenXR testbed illustrates a blended-reality cultural heritage space where robots, digital twins, XR interfaces, and generative AI coordinate visitor experiences and operations. Building on this scenario, the article identifies emerging signals, enabling technologies and phased adoption pathways across Industry 5.0 domains, positioning CPSS as a foundation for scalable, human-centric hyper-automation.

Robotic technologies are advancing rapidly, while the concept of hyper-automation is emerging as a powerful paradigm for automating end-to-end workflows. Advancements in Artificial Intelligence (AI), Machine Learning (ML), embedded computing, vision/language AI models, and edge intelligence enable robots to process large volumes of multimodal data, adapt to environments, learn from their actions and interactions, and adjust how they operate over time [1]. These capabilities are essential for operating in unstructured environments, handling variability, optimizing plans in real time, and personalizing interactions with humans. Hyper automation refers to the coordinated fusion of multiple technologies such as robotic process automation, workflow and business processes, AI and ML, process mining, and low-code platforms to discover, model, execute, monitor, and continuously improve automated processes[1]. Cyber-Physical (CPS) and Cyber-Physical-Social Systems (CPSS) are natural architectural foundations for hyper-automation. CPS tightly couple sensing, computation, and actuation to monitor and control physical processes, while CPSS extend this integration to include human and social behaviours as core system elements. The integration of cyber, physical, and social layers fosters hyper-automation by enabling system perception, analysis, planning, and execution, with humans remaining in the loop for oversight, exception handling and strategic decision-making. This article presents a future-looking scenario of robotic and AI-driven hyper-automation within the *InterwovenXR* CPSS testbed in a cultural heritage space, highlighting emerging signals, mapping enabling technologies, and outlining phased pathways for adoption across Industry 5.0 contexts motivating a new class of human-centric CPSS in which end-to-end workflows can be deployed and executed across real and virtual worlds.

Advancements in Hyper-Automation and Robotics

The rise of automation needs and technologies brought significant industrial interest in Hyper-Automation [2]. This paradigm is commonly described as a business-driven effort to identify and automate as

---

XXXX-XXX © 2021 IEEE
Digital Object Identifier 10.1109/XXX.0000.0000000

[1] https://www.gartner.com/en/information-technology/glossary/hyperautomation

This article has been accepted for publication in Computer. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/MC.2026.3671749

many businesses and IT processes as possible, fusing emerging automation technologies including AI, ML, event-driven architectures, robotics, robotic automation processes, and low-code tools [1]. While automation minimises human intervention across areas such as business processes, IT, or home automation etc., hyper-automation orchestrates multiple tools and platforms to accelerate digitisation, streamline actions, and improve overall operational efficiency. Academic literature often frames it as "*intelligent automation at scale*", where automation expands from basic automation tasks to end-to-end workflows that include discovery, analysis, orchestration, monitoring, and decision support, increasing human capabilities, scaling operations, and reducing risk [2].

One of the key technologies fostering hyper-automation is robotics, extending automated workflows beyond digital processes into the physical world through embodied sensing and actuation. Robotics have moved beyond research labs into industry, education, and society as intelligent interactive actuators that sense, reason, share spaces with humans, execute tasks with precision, and adapt under uncertainty. Modern collaborative robots (cobots) integrate force sensing, vision, and compliance control to operate safely alongside people [3]. Mobile robots have matured, autonomously navigating dynamic environments, capable of semantic mapping, vision, and multi-sensor fusion for inspection, material handling, environmental monitoring, localisation and obstacle awareness [4], [5]. Humanoid robots further extend these capabilities, with AI-driven perception, interaction, and behavioural control for robust bipedal locomotion and operation in unstructured environments [6]. Successful robotics deployments now span educational, healthcare, service, inspection, agriculture, delivery, and social or assistive domains ([7]).

A major shift is the integration between robotics and simulation with the use of digital twins. Digital twins are dynamic digital replicas of physical assets, environments, or processes that are continuously synchronised with their real counterparts via sensor, control, and event data. They provide a shared cyber layer where the state of mirrored actors such as robots, infrastructure, and workflows can be visualised, monitored, analysed, and exercised under hypothetical conditions before changes are applied in the real world [8]. In robotics, digital twins enable teleoperation, task rehearsal, safety validation, predictive simulation, and multi-robot coordination for decision support and automation. Sophisticated software development kits, high-fidelity physics engines, photorealistic rendering, and edge computing now allow robotic skills to be trained, validated, and remotely monitored in real time. These capabilities are encapsulated within the TwinOps paradigm, bringing DevOps-style practices together with model-based engineering and digital twins for CPS, treating models, generated code, and twin instances as priority actors in a continuous lifecycle that spans design, simulation, deployment, monitoring, and diagnosis [9]. Large industries are building full-stack ecosystems that link robotics, digital twins, and physical AI. For example, NVIDIA combines GPU-accelerated edge hardware with Omniverse-based Isaac simulators and robot foundation models such as Isaac GR00T and Cosmos, allowing to model 3D settings such as warehouses and factories, generate synthetic sensor data, and validate autonomous and other types of robots in virtual settings before deployment. Similarly, Siemens integrates industrial digital twins with robotics (e.g. SIMATIC Robot Pick AI), while AWS offers services like RoboMaker and IoT RoboRunner for simulating, deploying, and orchestrating robot fleets within broader cyber-physical infrastructures.

The latest advent of generative AI also found its way into robotics. Robotics-focused edge AI platforms now make it possible to deploy Small and Large Language Models (S/LLMs), Vision Language Models (VLMs), and robot foundation models directly on power-efficient embedded hardware, running open-source models locally on devices for private, low-latency natural language assistants, visual monitoring, and high-level task reasoning without cloud reliance. This combination of GPU-accelerated hardware, foundation models, and simulation aligns with the vision of *"edge intelligence"*, shifting computing power to the edge of the network to support real-time responses in physical environments [10], while robots remain connected to digital twins and backend services for coordination at system scale.

Hyper-Automation in Cyber-Physical-Social Systems

CPS integrate computation with the physical world through networks of sensors, actuators, and smart devices to support real-time monitoring, feedback control, optimisation, predictive maintenance and other tasks [11]. They are widely and successfully deployed in industrial domains such as automotive and manufacturing, but typically position humans as external operators rather than integral system components [12]. CPSS extend CPS by treating human and social elements as system components. A CPSS tightly couples physical sensing and actuation, computation, and social layers that support communication, collaboration, and

This article has been accepted for publication in Computer. This is the author's version which has not been fully edited and
content may change prior to final publication. Citation information: DOI 10.1109/MC.2026.3671749

experience sharing [13], [12]. By explicitly modelling human behaviour, preferences, and organisational processes alongside cyber and physical layers, CPSS enable system behaviour to emerge from interactions among people, devices, services, and data-driven intelligence. This integration provides a natural architectural foundation for hyper-automation, supporting processes that monitor complex environments, analyse multimodal data, and trigger coordinated actions under human supervision. CPSS pursue *hyper-connectivity, hyper-intelligence, and hyper-automation*, where learning, self-organisation, and coordination mechanisms enable large-scale automated regulation and control [14]. As such, CPSS act as system-level enablers for automating end-to-end workflows, retaining human influence on the system and its operations.

In robotic-equipped CPSS, robots are key nodes that provide embodied sensing and actuation capabilities for hyper-automated workflows, operating as intelligent physical agents that continuously capture data from their surroundings, reason, plan and execute actions, and exchange state information with their corresponding digital twins and backend services. Each digital twin (DT) maintains a real-time shadow of the robot's kinematics, health status, state, and operating context. Robotic DTs ingest telemetry from onboard sensors and the CPSS infrastructure, enabling visualision, simulation, evaluation and decision making before actions are deployed back to the physical robot or in real time. Through standardised middleware (e.g. ROS) and edge–cloud connectivity, robots can publish their status, receive optimised plans, and coordinate with other cyber, physical and social components, enabling the CPSS to realise hyper-automation across perception, analysis, planning, and execution, while keeping humans in the loop as supervisors and contributors to the system's behaviour.

A Model of Interaction - The Case of *InterwovenXR*

InterwovenXR is a CPSS designed as an experimental testbed prototype that fuses XR, AI, robots and other emerging technologies to create interconnected interactive experiences to blend physical and digital spaces. The current system version is a prototype that links the physical world (i.e. a cultural heritage installation equipped with physical robot workspaces) with virtual spaces infused with S/LLM capabilities and XR interfaces where human in the real world, remote users, and AI agents can meet, communicate, and interact with shared content and with each other (Fig. 1) [15], [16].

Figure 1 depicts the testbed system architecture. The *physical layer* comprises mobile robots, sensing infrastructure and display technologies in the real world; the *cyber layer* comprises virtual environments, robotic digital twins and spaces, and AI services, and the *social layer* captures human presence, interaction, communication, and behavioural influence across co-located and remote participants. In this testbed system, robots and their digital twins operate as intelligent intermediaries between realities. On the physical side, a robot can act as a mobile guide or telepresence actor that navigates the space, responds to visitor requests, evaluates information, and streams audio-visual information. In the digital space, its digital twin is mirroring position and behaviour for remote users to experience the same scene in real time, interact with the real world and explore alternative views of the spatial content. User interactions in XR and sensor data from the physical deployment are continuously fed into the cyber layer, where AI and generative AI services can adapt narratives, personalise interactions, and coordinate tasks between physical robots, their digital twins and virtual agents in the environment. The architecture further includes a *Knowledge Base* layer that stores exhibit data and system states for context aware interaction across components. A *Sensors and Triggers* subsystem monitors user proximity, movement, and interaction events, generating signals that activate the Generative AI module, which formulates context-aware prompts, queries selected S/LLMs, and returns personalised narrative content dispatched to agents, robots, or UI elements. The 'Brain' AI unit acts as the orchestration engine, coordinating behaviours and decisions across virtual agents, robots, and digital twins.

At its current stage of development, the *InterwovenXR* testbed prototype features a live digital twin of a physical mobile robot operating in the real world, navigating and streaming live video and audio feed in the virtual space, allowing remote users to see the real world through the eyes of the robot. The virtual environment is a multi-user space where users interact with each other through their avatars, and communicate using text chat and audio communication for synchronous collaboration and social presence. The virtual space features several heritage exhibits and educational content, and is further enriched with mini-games that guide visitors through cultural heritage material. These interactions are further supported by integrated generative AI services at the edge, where S/LLMs power conversational agents and dynamically populate the environment with context-aware narrative information and guidance.

This article has been accepted for publication in Computer. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/MC.2026.3671749

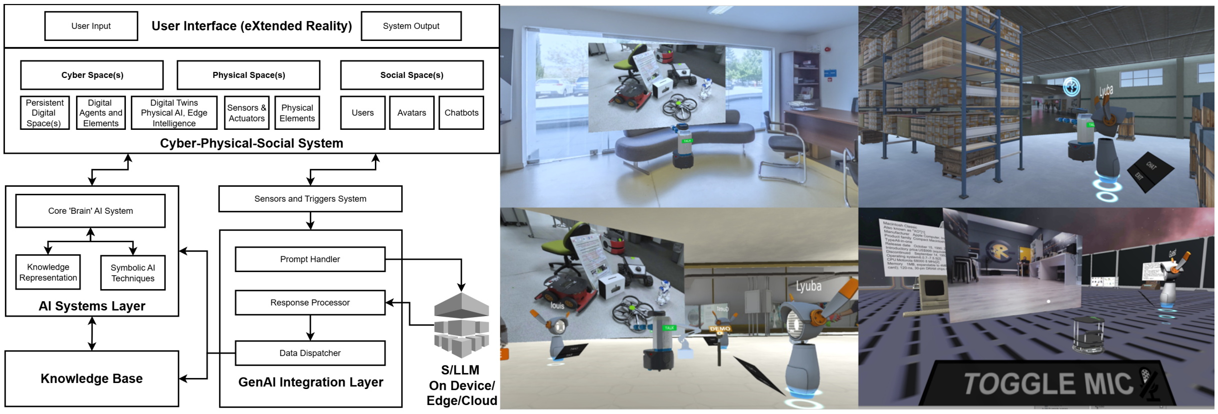

FIGURE 1. Overview of the *InterwovenXR* CPSS testbed system architecture (left) and example prototype deployments (right)

To evaluate the efficacy and user experience of InterwovenXR, a series of user experience evaluations have been conducted, revealing that the system is offering an immersive, engaging, and well-accepted mixed-reality experience [17], [15], [16]. Across these studies, participants responded positively to the synchronised behaviour of the robot and its digital twin, and described interactions with the AI guide and virtual environment as natural and meaningful. Engagement evaluation results further showed that users found the environment rewarding, visually appealing, and easy to navigate, reinforcing the value of integrating robotics, XR and AI as complementary interaction channels. Expert evaluations of the integrated LLM and SLM components for conversational intelligence demonstrated improved quality and relevance of generated content when fine-tuned models were used, although occasional hallucinations, synchronisation issues, and response-quality problems remain priorities for optimisation. Overall, these early findings do not yet constitute full system validation, but demonstrate that *InterwovenXR* can be used for studying real-time human–robot–AI interaction across interconnected physical/virtual spaces, indicating the need for further empirical work before full system validation. These publications focused on system design, user experience evaluation, and AI integration. Building on this foundation, the present article describes the applicability of *InterwovenXR* as a robotics-driven hyper-automation testbed, presenting a unified CPSS architecture that fuses robots, digital twins, edge intelligence, and generative AI, and introducing a CPSS hyper-automation model that links these enabling technologies to Industry 5.0 domains. It further outlines staged pathways toward closed-loop, human-centric hyper-automation, generalising the *InterwovenXR* paradigm as a scalable blueprint applicable to a plethora of domains.

InterwovenXR can be seen as a living lab for robotic and AI-driven hyper-automation, where sensing, perception, user modelling, and robotic guidance form end-to-end workflows that are partly automated and facilitated by human choice. It illustrates how CPSS can be engineered to support intelligent human–robot–AI interaction, and how they can evolve towards hyper-automation that still preserves human agency and control. Within this context, a future desired scenario is described below, where *InterwovenXR* operates as a futuristic hyper-automated blended reality exhibition space aligned with the Industry 5.0 paradigm, combining human-centred experience design with efficient operation orchestrated through robotic, AI, data-driven operations, influenced by humans and their behaviour.

Future Desired Scenario

In a forward-looking deployment, *InterwovenXR* operates as a CPSS that hyper-automates a heritage venue across blended physical and virtual spaces. The exhibition gallery, robot work zones, and visitor services are digitised and mirrored through their digital twins, maintained through multi-sensor fusion from robot-mounted perception and infrastructure sensing. On-site robots act as embodied AI intermediaries that deliver navigation, curation, and assistance, shadowed by their digital twin, where the cyber layer monitors higher-level situational understanding for workflow orchestration.

Interactions in the Physical Space. As visitors arrive, they are greeted by robots that support navigation, curation, and interaction. Robots use edge perception for obstacle avoidance, human detection, proxemics,

This article has been accepted for publication in Computer. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/MC.2026.3671749

visual perception, and natural language understanding via VLMs and S/LLMs, contributing higher-level environment understanding to the cyber layer. The AI orchestration layer converts perception and interaction into system behaviour, such as the S/LLM subsystem receiving signals from the twin and live interaction between users and robots, for producing context to support the system knowledge and enhancing robots' contextual understanding.

A school group arrives, guided by a mobile robot through the gallery and key artefacts. The robot welcomes and leads them through the gallery and key artefacts. The system estimates congestion and generates plans such as rerouting to adjacent spaces, adjusting narrative pacing, or allocating a second robot to split the group. Robots help maintain group cohesion, support telepresence for remote participants, and adapt to context. In quiet zones, interaction switches to captions on robot displays or personal devices, and accessibility modes select routes that avoid stairs, adjust pace, and reduce reliance on speech.

Visitors' actions influence system behaviour in real time, affecting robot interactions, digital twin monitoring, AI decision-making and generative AI-driven narratives. Visitors ask robots open questions, or to pause, repeat, switch language, request staff help etc., and these actions directly update planning and content generation. Actions within policy execute automatically, while others require staff approval through an XR control interface. When anomalies occur, congestion spikes, or environmental alerts, the system generates intervention options, simulates outcomes, and either executes within policy or escalates to staff with an explanation and recommended actions. At shift-based intervals, the system produces summaries and log keeping support optimisation of simulation models, intervention policies, and interaction design. Outside opening hours, the system runs calibration and integrity checks such as mobile robots location validation, map alignment, and safety perimeters, and infrastructure sensing confirms spatial anchors against the twin, flagging errors or gaps for correction.

The Digital Space While the visit unfolds in the physical gallery, a virtual world mirrors the venue through digital twins. Remote users connect to a high-fidelity reconstruction of rooms, exhibits, and interpretive media, interacting with these and with one another via avatars. Robot and sensor data from the physical space update the virtual environment so remote participants can see robot positions, group location, active exhibits, and how the space is evolving, such as crowding patterns. Remote visitors can join the tour as telepresence participants, follow a companion view aligned with the robots' narrative and group position, switch languages, request captions, and submit questions routed through the same interaction layer as on-site queries. The system also supports asynchronous engagement, where visitors can explore the virtual gallery independently and interact with exhibits and contextual stories generated from curated knowledge. When the venue is closed, physical robots operate autonomously while their twins remain accessible in cyberspace, acting as embodied cameras and narrators that stream after-hours walkthroughs and support on-demand visits, extending access and maintaining readiness for reopening.

Signals

The future desired scenario described above envisions how robots, digital twins, and AI-driven orchestration can deliver hyper-automation across physical and virtual spaces. Although framed around a cultural heritage venue, the underlying principles generalise across industrial verticals. Several technological, industrial, research and policy developments indicate that the integration of robotics, hyper-automation, and CPSS is already underway. One key signal for hyper-automated Industry 5.0 is the increasing deployment of industrial robots and the growth in research and applications. World Robotics statistics reports more than 4 million industrial robots in operation worldwide in 2024, with annual installations exceeding 500,000 units[2], and global robot demand in factories in 2025 has doubled over the past 10 Years[3]. Digital twin technologies are becoming integrated with robotics and industrial processes, ingesting robot telemetry and infrastructure sensing to support monitoring, simulation-based decision support and optimisation of automated workflows [18].

Another signal towards hyper-automated CPSS is the maturation of connectivity, sensing, and localisation technologies that enable persistent spatial awareness across physical and digital domains. Indoor localisation, long a limiting factor for autonomous systems in human-centric environments and particularly indoors, has advanced through UWB ranging, Wi-Fi, mmWave radar, and vision–inertial odometry, with growing deployment in factories, hospitals, museums, and logistics hubs. Within Integrated Communications

---

[2] https://ifr.org/ifr-press-releases/news/record-of-4-million-robots-working-in-factories-worldwide

[3] https://ifr.org/ifr-press-releases/news/global-robot-demand-in-factories-doubles-over-10-years

This article has been accepted for publication in Computer. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/MC.2026.3671749

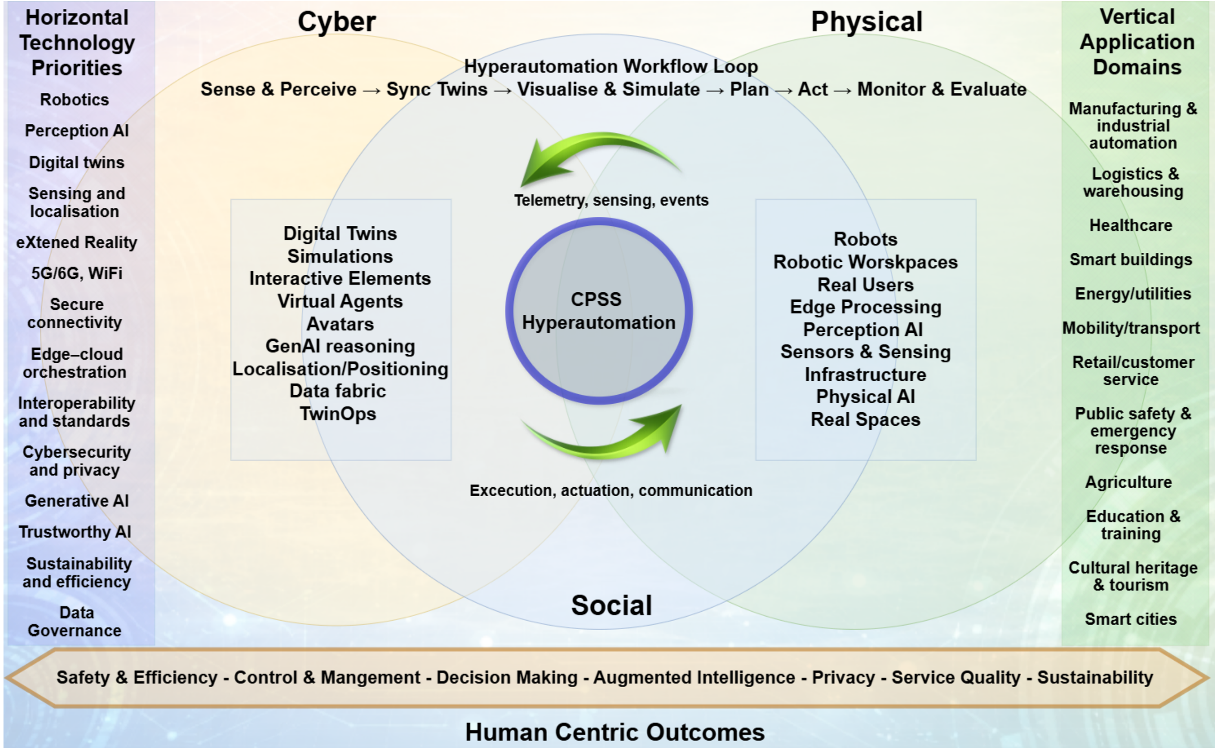

FIGURE 2. CPSS hyper-automation model connecting horizontal technology priorities (left) to vertical application domains (right) through integrated cyber, physical, and social capabilities, targeting human-centric outcomes.

and Sensing (ISAC) [19], 5G and emerging 6G integrate positioning, sensing, and communication into a unified fabric, enabling sub-meter accuracy, ultra-low latency control, and high-bandwidth sensor fusion. These gains are reinforced by edge computing, which runs localisation, mapping, and perception close to the data source for lower latency and greater resilience. Together, they signal a shift from isolated localisation capabilities to position-aware CPSS, where spatial context becomes a shared resource for digital twins, orchestration, and human–robot interaction.

At the same time, hyper-automation appears in industry roadmaps as an operating model that combines workflow automation with AI-driven prediction and optimisation to automate processes end-to-end, with major vendors and consultancies outlining staged transitions toward more adaptive, self-optimising operations (i.e. [4] and [5]).

The notion of industrial Metaverse and AI factories is gaining traction in research and practice. Platforms such as NVIDIA Omniverse and Isaac Sim support high-fidelity simulation, synthetic data generation and validation of AI-driven robots within digital twins of factories, warehouses and other environments. Robotics workflows are moving toward an end-to-end stack where foundation models are trained in simulators and then deployed on robots for real-time control. Edge AI platforms like NVIDIA Jetson run S/LLMs, VLMs, and robot foundation models directly on embedded GPUs, enabling private, low-latency perception and reasoning on-board robots rather than in the cloud[6]. VLA models such as RT-2 and robot foundation model stacks like NVIDIA's GR00T line show that robots able to interpret open questions, context, and intent are now a mainstream research and industry direction[7].

Industrial analyses forecast factories orchestrated through AI-governed digital twin control loops, where most production decisions are modelled in the virtual

---

[4] https://www.deloitte.com/us/en/insights/topics/talent/intelligent-automation-2022-survey-results.html

[5] https://www.ibm.com/roadmaps/automation/

[6] https://developer.nvidia.com/blog/getting-started-with-edge-ai-on-nvidia-jetson-llms-vlms-and-foundation-models-for-robotics/

[7] https://deepmind.google/blog/rt-2-new-model-translates-vision-and-language-into-action/

This article has been accepted for publication in Computer. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/MC.2026.3671749

environment before execution in the physical facility[8]. Policy and funding initiatives increasingly target the intersection of AI, robotics, virtual worlds, and digital twins. For example, European strategic documents on AI, data and robotics support networks of excellence and partnerships that promote trustworthy AI, physical AI and robotics integrated with virtual environments. Several funding calls explicitly encourage generative AI in robotics and industrial automation, expecting measurable productivity gains and requiring demonstrators in realistic contexts. Standardisation efforts, including collaborative robot safety requirements (ISO 10218 family, ISO/TS 15066)[9] and digital twin frameworks such as ISO 23247[10], move practice from ad hoc implementations toward structured lifecycle management and clearer interfaces for manufacturing digital twins.

Early CPSS deployments already appear in smart factories and infrastructure management, where robotic digital twins and automation platforms explore what-if scenarios and coordinate human and robotic responses to disruptions [14] These cases demonstrate the feasibility of treating robots, digital twins, and human operators as interconnected agents in socio-technical systems, signalling that the envisioned intelligent CPSS-oriented future is emerging.

Required Technologies

Realising this vision of fused robotics, hyper-automation and human–AI interaction in intelligent CPSS depends on advances across several technological layers.

Digital twins and simulation platforms. A core enabler of hyper-automation in robotics-driven CPSS is a digital twin infrastructure that represents robots, environments, and human actors with sufficient fidelity for visualisation, simulation, prediction, and control. Beyond static 3D models, hyper-automation requires live data-driven twins that fuse heterogeneous sensing (robot perception, infrastructure IoT, positioning) into continuously updated world states. This calls for standardised and interoperable 3D scene representations, physics-consistent simulation engines for robot, object, and crowd dynamics, and robust synchronisation between physical and virtual entities. To operationalise automation, twin platforms must expose state and events through APIs for analytics, policy orchestration, and human-in-the-loop supervision, so actions can be triggered from predicted conditions.

Sensing, localisation, and connectivity. Hyper-automated CPSS rely on dense, heterogeneous sensing to perceive the environment and maintain continuous spatial alignment between physical entities, humans, and their digital twins. Robust localisation is needed in dynamic, cluttered, human-populated settings. Core technologies of multi-modal sensing on robots, infrastructure-based sensing, and wearable or personal devices that contribute human-centric spatial signals are essential to achieve reliable positioning under occlusion, interference, and environmental change. Positioning accuracy and confidence directly affect digital twin fidelity, so synchronisation mechanisms must incorporate uncertainty modelling, drift detection, and periodic recalibration. Low-latency and resilient connectivity enabled by edge computing, time-sensitive networking, and emerging 5G/6G infrastructures supports the continuous exchange of spatial state, sensor data, and control signals across the CPSS stack.

Intelligence, planning and hyper-automation. AI is required at multiple levels, from low-level robot control and perception to high-level task planning and process optimisation. This includes learned perception models, motion and path planners for dynamic human environments, multi-robot coordination, reasoning, and decision-support systems that operate over twin data. Hyper-automation demands orchestration engines that integrate event detection, process mining, optimisation and workflow execution, with generative AI that maps high-level human goals into executable plans across robots, services and human roles.

Human–AI–robot interaction technologies. Interaction with intelligent CPSS through XR interfaces, conversational agents and adaptive visual analytics is tightly integrated with the digital twins. Required technologies include user interfaces, explanation facilities for AI and robotic decisions, mechanisms for sharing and managing control between humans and automation, and instrumentation for measuring KPIs and success, ensuring hyper-automation augments rather than replaces human expertise.

Infrastructure, interoperability and governance. Realising such systems demand secure, scalable and interoperable infrastructure, encompassing optimised architectures for distributing computation, common data models and standards for integrating robots, twins and

---

[8] https://www.accenture.com/us-en/insights/industrial/future-of-manufacturing

[9] https://www.iso.org/standard/62996.html

[10] https://www.iso.org/standard/75066.html

This article has been accepted for publication in Computer. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/MC.2026.3671749

enterprise systems, and security, privacy and safety mechanisms for safety-critical and public environments. Monitoring, auditing, and policy-enforcement tools are needed to keep AI-driven and robot-assisted processes aligned with regulatory, ethical, and organisational requirements.

Paths to Realisation

Moving from today's deployments to envisaged hyper-automated CPSS requires staged progress that manages risk while increasing autonomy, scale, and organisational adoption. The following paths outline practical steps applicable to Industry 5.0 verticals.

Establish CPSS-ready sensing and connectivity. Organisations should identify and prioritise processes where CPSS can deliver measurable value, mapping key workflows, safety, and data dependencies across people, robots, and infrastructure, and then selecting a small set of high-impact processes for initial instrumentation. To avoid over-instrumentation, sensing should follow a *"minimum viable observability"* principle, introducing measurements only when they directly support key metrics. Priority should focus on safety and compliance activities, high-value tasks and frequent issues, coordination-heavy tasks across teams or systems, and strong reuse potential across analytics and digital twin services. Deployment can start in sandbox zones, expand to controlled robot workspaces, and then scale to high-value or safety-critical spaces as evidence accumulates. Data quality checks, timestamping, and metadata governance should ensure trustworthy reuse of sensor streams. Early criteria should emphasise accuracy, connectivity, repeatability for analytics and simulation, and clear evidence that each sensing stream improves decisions relative to costs and complexity. Initial deployments should prioritise baseline spatial observability, enabling coarse localisation of robots, assets, and zones through infrastructure-supported technologies such as UWB, mmWave or vision-based mapping, targeting repeatability, coverage, and uncertainty estimation rather than absolute accuracy.

Develop twin-centric monitoring, simulation, and decision support. The next step is to introduce digital twins that mirror selected equipment, rooms, robots, and workflows, and expose their state to the CPSS for observation and analysis. Early deployments should emphasise visual monitoring and decision-support dashboards that keep humans in control while generating value through measurable gains in situational awareness, safety, and throughput. This phase should extend twins with simulation, prediction, and what-if analysis of robot behaviour, with outputs fed back for evaluation. Establishing such consistency aligns with the emerging concept of Digital Thread (DTh) approaches, which emphasise standards-based continuity between physical systems, digital twins, and organisational workflows [20]. In the meantime, organisations should mature synchronisation mechanisms, validate twin accuracy against the physical space through calibration checks, and establish standard interfaces to later support planning and automation. Evaluation should be embedded from the outset, considering accuracy and latency, decision time, incident rates, downtime, throughput, and user evaluations of perceived workload and usability. At this stage, digital twins should explicitly visualise localisation confidence, drift, and coverage gaps, allowing operators to understand where spatial uncertainty may affect decision-making, where positioning data becomes a diagnostic signal rather than just an input.

Integrate robots as digital twins. With twin-based monitoring and decision support mechanisms established, robots can be introduced as embodied agents in the workspace and connected to their digital twins. Early deployments should target well-bounded tasks such as inspection, material handling, or guided tours, where robot behaviour is simulated and validated in the twin before execution, e.g. path simulation and collision-risk checks, multi-robot coordination in shared zones, remote supervision or teleoperation. This stage should align with organisational readiness, including safety culture, clear procedures, maintenance and incident-response processes. Co-design with frontline staff and user experience studies should be used to identify concerns and improve usability and acceptance, particularly for collaborative and visitor-facing robots. Evaluation criteria include task success rate, intervention frequency, navigation and localisation error, downtime and maintenance time, safety events and near-misses, operator workload and user trust, usability, and acceptance. Introducing robots as digital twins requires validated localisation tests, including map alignment, anchor calibration, and continuous consistency checks between physical trajectories and twin states. Multi-robot environments particularly benefit from shared positioning references to support coordination and collision avoidance.

Introduce hyper-automation loops. As confidence in twin-centric operation grows, orchestration engines can progressively connect event detection, analytics, and robotic actions into closed control loops, shifting from isolated task automation to end-to-end hyper-automated workflows, and retaining human oversight for goals, policy constraints, exception handling,

This article has been accepted for publication in Computer. This is the author's version which has not been fully edited and
content may change prior to final publication. Citation information: DOI 10.1109/MC.2026.3671749

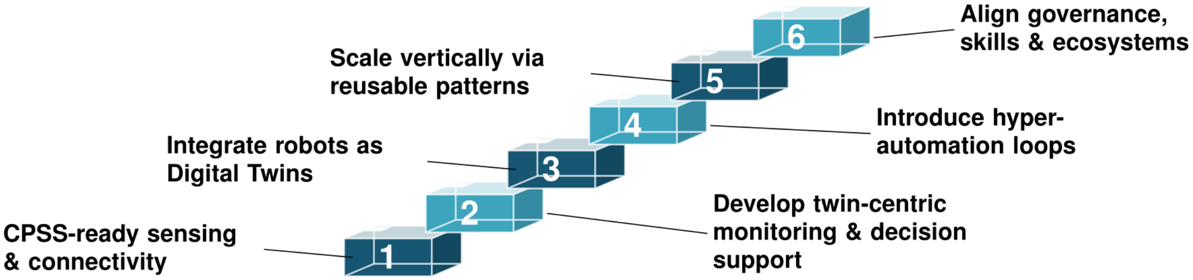

FIGURE 3. Staged paths to realising hyper-automated CPSS, progressing from bottom-left (Stage 1) to top-right (Stage 6).

and continuous improvement. Given the risk of poor or miscalibrated decisions, guardrails should define what is fully automated, what needs human approval, and which conditions trigger safe-stops. Failure modes should be anticipated through risk analysis, with escalation paths such as confidence thresholds, redundancy checks, human-in-the-loop confirmation for high-impact actions, and possible revert to manual control. Exceptions should be treated as priority system events, with logging, explanation traces, and structured handover to operators to support rapid diagnosis and learning. Continuous monitoring and evaluation are required to detect drift, emerging risks, automation issues, such as false alarm and miss rates, time-to-detect and time-to-intervene, autonomy level achieved, exception frequency, safety incidents and near-misses, operational throughput, and operator workload and trust. Localisation and positioning systems must be relevant to organisational policies, defining when spatial uncertainty triggers degraded autonomy, human approval, safe-stop behaviours and privacy concerns. Hyper-automation should operate on confidence-weighted spatial data, not raw estimates.

Scale vertically through reusable patterns. Design and implementation patterns for sensing, twinning, orchestration, and human interaction can be packaged as reusable blueprints and reference architectures, then adapted to additional lines, facilities, and domains such as logistics hubs, campuses, or cultural venues using interoperable data models and standard interfaces. Modularity and configurability should be prioritised so organisations can extend the CPSS stack without redesigning core components. Scaling must recognise that robot capabilities, safety standards, and human interaction patterns vary across verticals, leading to different requirements for autonomy, certification, supervision, and user experience. Reusable patterns should therefore separate services from domain-specific policies and interaction designs, enabling consistent deployment while respecting sector constraints and local practices.

Align governance, skills, and ecosystems. Across all stages, organisations should invest in workforce skills, change management, and governance mechanisms that define responsibilities, risk controls, and evaluation criteria for AI-driven and robot-assisted processes. This includes addressing organisational resistance and the operational complexity of multi-stakeholder governance, spanning management, safety officers, IT, frontline staff, and, where relevant, public-facing stakeholders. Clear boundaries of robot responsibility and human authority, together with procedures for incidents, audit trails, and explainability of AI decisions, are needed to support accountability and trust.

Progress between stages should be evidence-based, using staged gates with indicators that combine technical performance, safety validation, operational impact, and socio-technical evidence. Engagement with standardisation bodies, industry alliances, and research testbeds can accelerate learning, improve interoperability, and help ensure deployments remain human-centric, trustworthy, and aligned with regulatory and societal expectations.

Inhibitors, Risks, and Mitigation Considerations. While these paths outline a progressive trajectory, it is equally important to acknowledge inhibitors that could slow adoption, particularly those arising from human–robot interaction. Organisations often encounter concerns around safety, trust, unpredictability of robot behaviour, and the cognitive burden placed on operators overseeing autonomous systems. Misalignment between human expectations and robot actions, unclear handover points, and limited transparency in decision-making can lead to hesitancy and resistance, among other issues. Mitigation requires intentional socio-technical design such as co-design with frontline users, clear communication of robot intent, explainable AI, graduated autonomy with override options, and evaluation of failure modes before scaling. Address-

This article has been accepted for publication in Computer. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/MC.2026.3671749

ing these factors early helps keep hyper-automation human-centric and supports organisational readiness. Beyond technical pathways, real-world deployment of CPSS must handle broader open challenges and risks that shape long-term viability. Data privacy is a core concern as sensing, localisation, and behavioural analytics increase the volume and sensitivity of collected information. AI explainability is critical, since opaque decisions can undermine trust and hinder incident investigation. Safety assurance adds further complexity, requiring continuous verification of autonomy under uncertainty, robust fail-safe mechanisms, and processes to evaluate drift and unintended behaviour. Recognising and managing these socio-technical risks is essential for trustworthy, transparent, human-centred CPSS.

Conclusions

Robotics and AI-driven hyper-automation in CPSS can be achieved by engineering interconnected systems that fuse sensing, digital twins, AI reasoning, and human influence rather than deploying more autonomous machines in isolation. The future desired scenario of *InterwovenXR* envisions how twin-centric monitoring and simulation orchestration can turn robots into safe service actors across blended reality environments. Progress toward this vision requires accurate localisation and synchronisation, interoperable twin platforms, high-performance edge intelligence, and interaction technologies that preserve human interaction and influence. The staged pathways outlined in this article provide a practical route from minimum viable observability to reusable deployment patterns. The success of hyper-automated CPSS should be measured by throughput and cost reduction, together with safety, resilience, usability, and human trust, with future work prioritising rigorous evaluation protocols, standard interfaces, and socio-technical design that align autonomy with human values.

REFERENCES

1. M. Soori, B. Arezoo, and R. Dastres, “Artificial intelligence, machine learning and deep learning in advanced robotics, a review,” *Cognitive Robotics*, vol. 3, pp. 54–70, 2023.
2. A. Haleem, M. Javaid, R. P. Singh, S. Rab, and R. Suman, “Hyperautomation for the enhancement of automation in industries,” *Sensors International*, vol. 2, p. 100124, 2021.
3. F. Semeraro, A. Griffiths, and A. Cangelosi, “Human–robot collaboration and machine learning: A systematic review of recent research,” *Robotics and Computer-Integrated Manufacturing*, vol. 79, p. 102432, 2023.
4. F. Rubio, F. Valero, and C. Llopis-Albert, “A review of mobile robots: Concepts, methods, theoretical framework, and applications,” *International Journal of Advanced Robotic Systems*, vol. 16, no. 2, p. 1729881419839596, 2019.
5. Y. Liu, S. Wang, Y. Xie, T. Xiong, and M. Wu, “A review of sensing technologies for indoor autonomous mobile robots,” *Sensors*, vol. 24, no. 4, 2024.
6. Y. Tong, H. Liu, and Z. Zhang, “Advancements in humanoid robots: A comprehensive review and future prospects,” *IEEE/CAA Journal of Automatica Sinica*, vol. 11, no. 2, pp. 301–328, 2024.
7. J. T. Licardo, M. Domjan, and T. Orehovački, “Intelligent robotics—a systematic review of emerging technologies and trends,” *Electronics*, vol. 13, no. 3, 2024.
8. N. Napp and L. Nisiotis, “Digital reality technology, challenges, and human factors,” *IEEE Technology and Engineering Management Society Body of Knowledge (TEMSBOK)*, pp. 331–349, 2023.
9. J. Hugues, A. Hristosov, J. J. Hudak, and J. Yankel, “Twinops-devops meets model-based engineering and digital twins for the engineering of cps,” in *Proceedings of the 23rd ACM/IEEE International Conference on Model Driven Engineering Languages and Systems: Companion Proceedings*, pp. 1–5, 2020.
10. Z. Zhou, X. Chen, E. Li, L. Zeng, K. Luo, and J. Zhang, “Edge intelligence: Paving the last mile of artificial intelligence with edge computing,” *Proceedings of the IEEE*, vol. 107, no. 8, pp. 1738–1762, 2019.
11. Q. Chang, R. Gao, Y. Lei, L. Wang, C. Wu, *et al.*, “Cyber-physical systems in manufacturing and service systems,” *Mathematical Problems in Engineering*, vol. 2015, pp. 1–2, 2015.
12. B. A. Yilma, H. Panetto, and Y. Naudet, “Systemic formalisation of cyber-physical-social system (CPSS): A systematic literature review,” *Computers in Industry*, vol. 129, p. 103458, 2021.
13. L. Monostori, “Cyber-physical production systems: Roots, expectations and R&D challenges,” *Procedia CIRP*, vol. 17, pp. 9–13, 2014.
14. Y. Zhou, F. R. Yu, J. Chen, and Y. Kuo, “Cyber-physical-social systems: A state-of-the-art survey, challenges and opportunities,” *IEEE Communications Surveys Tutorials*, vol. 22, no. 1, pp. 389–425, 2020.
15. L. Nisiotis, C. Nikolaou, N. Markov, and A. Hadjiliasi, “Developing a cyber-physical-social metaverse system for interactive cultural heritage experiences,” in *2025 IEEE 49th Annual Computers, Software, and*

This article has been accepted for publication in Computer. This is the author's version which has not been fully edited and
content may change prior to final publication. Citation information: DOI 10.1109/MC.2026.3671749

Applications Conference (COMPSAC), pp. 654–663, IEEE, 2025.
16. L. Nisiotis, N. Markov, C. Nikolaou, A. Hadjiliasi, and M. Raspopoulos, “Enhancing digital heritage experiences: Evaluating fine-tuned llm integration within a cyber-physical-social virtual museum system,” in *2nd IEEE International Conference on Emerging Metaverse (ISEMV), Honolulu, HI*, 2025.
17. L. Nisiotis and L. Alboul, “Evaluating presence and technology acceptance of an intelligent reality virtual museum prototype,” in *2022 IEEE 2nd International Conference on Intelligent Reality (ICIR)*, pp. 1–6, IEEE, 2022.
18. W. Kritzinger, M. Karner, G. Traar, J. Henjes, and W. Sihn, “Digital twin in manufacturing: A categorical literature review and classification,” in *IFAC-PapersOnLine*, vol. 51, pp. 1016–1022, 2018.
19. Y. Cui, F. Liu, X. Jing, and J. Mu, “Integrating sensing and communications for ubiquitous IoT: Applications, trends, and challenges,” *IEEE Network*, vol. 35, no. 5, pp. 158–167, 2021.
20. E. Rodriguez and A. J. Alvares, “Toward a digital ecosystem for additive manufacturing driven by standards-based digital thread and digital twins,” *IEEE Access*, vol. 13, pp. 207776–207807, 2025.

Louis Nisiotis is an Assistant Professor at UCLan Cyprus, in Pyla, 7080, Cyprus. His research interests include eXtended Reality, Virtual Worlds, Digital Twins, and Entertainment Computing. Dr Nisiotis received his PhD in Computing from Sheffield Hallam University, UK. He is a member of IEEE, and IEEE Computer Society. Contact him at LNisiotis@uclan.ac.uk.

Marios Raspopoulos is an Associate Professor at UCLan Cyprus, in Pyla, 7080, Cyprus. His research interests include wireless communication, radio channel modelling, Internet of Things, and Indoor Localisation, Positioning and Navigation. Dr Raspopoulos received his PhD in Telecommunication Engineering from the University of Surrey, UK. He is a member of IEEE, and IEEE Communications Society (ComSoc). Contact him at MRaspopoulos@uclan.ac.uk.