

Next-Generation Cyber-Physical Systems: Architectures, Intelligent Technologies, Applications, and Future Research Directions

Dr. Achyuta Nand Mishra, Associate Professor, Department of Electronics and Communication Engineering, Supaul College of Engineering, under the Department of Science, Technology and Technical Education, Bihar, India an.mishra53@gmail.com

Ravi Ranjan, Assistant Professor, Department of Electronics and Communication Engineering, Supaul College of Engineering, under the Department of Science, Technology and Technical Education, Bihar, India ranjan92.ravi@gmail.com

Naveen Kumar, Assistant Professor, Department of Electronics and Communication Engineering, Supaul College of Engineering, under the Department of Science, Technology and Technical Education, Bihar, India naveen6064@gmail.com

Pavan Kumar Shukla, School of Engineering and Technology, CGC University Mohali, Punjab, India pavanfec@gmail.com

Abstract: Cyber-Physical Systems (CPSs) have evolved from tightly coupled sensing, computation, communication, and control platforms into highly distributed intelligent ecosystems in which physical processes, computational resources, communication infrastructures, artificial intelligence, digital twins, and human participants interact continuously. The emergence of edge and fog computing, 5G/6G networks, machine learning, reinforcement learning, federated learning, blockchain, digital twins, autonomous systems, and foundation models is substantially changing the architecture and operational characteristics of next-generation CPSs. Unlike conventional CPS architectures, emerging systems must simultaneously address real-time constraints, heterogeneous data, dynamic environments, distributed intelligence, safety, cybersecurity, privacy, explainability, interoperability, and lifecycle assurance. This review provides a comprehensive analysis of next-generation CPSs from architectural, technological, methodological, application, and research perspectives. A unified taxonomy is developed covering physical, sensing, communication, edge, cloud, intelligence, digital-twin, orchestration, security, and human-interaction layers. The review critically examines centralized, edge-centric, cloud-assisted, distributed, digital-twin-enabled, and AI-native architectures and compares their characteristics with respect to latency, scalability, resilience, computational requirements, interoperability, privacy, and assurance. Mathematical formulations are presented for CPS state estimation, control, resource optimization, anomaly detection, and multi-objective performance assessment. Emerging intelligent methodologies—including deep learning, reinforcement learning, federated learning, graph neural networks, physics-informed learning, generative AI, and foundation models—are analyzed in terms of their applicability and limitations. Applications in smart manufacturing, autonomous transportation, smart grids, healthcare, agriculture, buildings, robotics, and smart cities are discussed, followed by an illustrative industrial CPS case study integrating edge intelligence, digital twins, predictive maintenance, and secure control. Particular attention is given to benchmark datasets, simulation platforms, communication technologies, digital-twin frameworks, and cybersecurity mechanisms. Finally, unresolved challenges concerning real-time AI assurance, semantic interoperability, adversarial robustness, uncertainty quantification, energy efficiency, human autonomy, ethical governance, and verification of learning-enabled controllers are identified. The review concludes by proposing a research agenda toward trustworthy, adaptive, interoperable, and self-evolving CPS ecosystems.

Keywords: Cyber-Physical Systems; Intelligent CPS; Industrial CPS; Edge AI; Digital Twin; Artificial Intelligence; Machine Learning; Reinforcement Learning; Federated Learning; 5G; 6G; Time-Sensitive Networking; CPS Security; Autonomous Systems; Foundation Models; Smart Manufacturing.

1. Introduction

Cyber-Physical Systems represent a fundamental paradigm for integrating computational intelligence with physical processes. A CPS is not simply a network of connected sensors or an Internet of Things deployment. Its defining characteristic is the continuous interaction between computational models and physical processes through sensing, communication, computation, decision-making, and actuation. The physical environment generates observations, computational components interpret these observations, control decisions are generated, and actuators subsequently modify the physical state. This closed-loop relationship distinguishes CPS from conventional information systems, where computation may remain largely independent of physical dynamics. The NIST CPS framework consequently treats CPS as systems involving interacting logical, physical, transducing, and human components whose behavior emerges from the integration of logic and physics [1], [2].

The traditional CPS research landscape has been strongly influenced by embedded control systems, real-time operating systems, wireless sensor networks, industrial control systems, and model-based engineering. These foundations remain important, but contemporary CPS deployments have become significantly more heterogeneous and geographically distributed. A modern autonomous vehicle, for example, may simultaneously incorporate cameras, radar, lidar, inertial sensors, vehicle controllers, high-performance processors, wireless connectivity, cloud services, machine-learning models, maps, digital twins, and human supervision. Similarly, an intelligent manufacturing plant may contain programmable logic controllers, robots, machine-vision systems, industrial gateways, edge servers, cloud analytics, digital representations of machines, enterprise systems, and autonomous optimization services. The complexity of these interactions creates architectural and assurance

problems that cannot be addressed adequately using isolated control, networking, or software-engineering techniques.

The transformation is particularly evident in the increasing movement of intelligence toward the edge. Cloud computing provides substantial computational capacity and facilitates global model management, but sending all CPS data to remote data centers introduces latency, bandwidth consumption, privacy exposure, and connectivity dependencies. Edge AI addresses part of this problem by placing inference and selected learning functions close to data sources [3]. The resulting architecture allows time-critical decisions to remain local while computationally intensive model training, global optimization, and long-term analytics may be performed in cloud or data-center environments. However, edge intelligence introduces new challenges involving heterogeneous hardware, model compression, distributed orchestration, energy constraints, model synchronization, and security.

Digital twins constitute another major architectural development. A digital twin provides a computational representation of a physical asset, process, or system and may be continuously updated using operational data. In CPS, digital twins can support monitoring, simulation, fault diagnosis, predictive maintenance, optimization, and testing. Their significance extends beyond visualization because a sufficiently integrated twin can provide an intermediate computational environment in which alternative actions can be evaluated before being applied to the physical system. Recent systematic research has specifically investigated digital-twin-based testing of CPS and identified the potential of twins to serve as dynamic test oracles and predictive environments [4]. Nevertheless, digital twins introduce difficult questions concerning synchronization fidelity, model uncertainty, semantic interoperability, computational overhead, and validation.

Artificial intelligence is simultaneously changing the computational paradigm of CPS. Earlier CPS controllers generally depended on deterministic mathematical models, whereas contemporary systems increasingly combine mechanistic models with machine learning. Deep neural networks can interpret high-dimensional sensory information; recurrent and transformer-based architectures can process temporal data; graph neural networks can represent interconnected components; reinforcement learning can optimize sequential decisions; and physics-informed learning can incorporate physical constraints into data-driven models. More recently, foundation models and multimodal models have attracted attention as potential general-purpose computational components for CPS-IoT systems [5], [6]. Their potential is substantial, but their stochastic behavior, resource requirements, opaque internal representations, and difficulty in guaranteeing safety make direct integration into safety-critical control loops problematic.

The communication substrate has evolved in parallel. Industrial Ethernet, wireless sensor networks, 5G, private cellular networks, and Time-Sensitive Networking increasingly support distributed CPS applications. TSN is particularly relevant because CPS control requires bounded latency and jitter rather than merely high average throughput. Research has therefore shifted toward deterministic and reliable communication mechanisms capable of supporting converged information-technology and operational-technology networks [7]. Emerging 6G research further introduces integrated sensing and communication, AI-native networking, extreme edge computing, and potentially sub-millisecond communication services.

Security has become equally inseparable from CPS engineering. A cyberattack against a conventional information system may compromise confidentiality or availability, whereas an attack against a CPS may manipulate physical processes. A malicious modification of sensor measurements can therefore lead to inappropriate control decisions, equipment damage, environmental consequences, or safety incidents. Research on CPS intrusion detection has increasingly employed machine learning and deep learning, with datasets such as SWaT and WADI serving as common experimental benchmarks [8], [9]. Nevertheless, high benchmark accuracy does not automatically establish safety in operational CPS environments because attacks may exploit temporal dependencies, physical process knowledge, distribution shifts, or weaknesses in the sensing and actuation pathways.

The next generation of CPS must consequently be viewed as an integrated systems-engineering problem rather than a collection of independent technologies. The key research question is no longer simply how to connect physical assets to computers. Instead, it concerns how heterogeneous computational intelligence can be embedded into physical systems while preserving timing guarantees, safety, security, privacy, explainability, interoperability, resilience, and human control throughout the system lifecycle.

This review addresses that problem by examining the evolution of CPS architectures, developing a technology-oriented taxonomy, analyzing intelligent methods and system architectures, comparing implementation approaches, surveying application domains, presenting an integrated industrial case study, and identifying

unresolved research challenges. The objective is not merely to catalogue technologies but to analyze the architectural relationships among them and identify the conditions under which they can form trustworthy next-generation CPS.

2. Background and Motivation

The conceptual foundations of CPS emerged from the convergence of embedded computing and control engineering. Conventional embedded control systems generally consisted of sensors, controllers, actuators, and deterministic communication links designed for specific physical processes. Their architecture was usually static, their computational functions were predefined, and their behavior could often be analyzed using mathematical models. As networking capabilities expanded, these isolated controllers became interconnected, producing networked control systems. The subsequent proliferation of Internet connectivity, wireless sensing, cloud computing, and IoT technologies created a much broader CPS ecosystem in which physical processes could interact with distributed computational services.

The NIST framework provides an important systems perspective by organizing CPS concerns across functional, human, business, trustworthiness, timing, data, composition, boundary, and lifecycle dimensions [1]. This multidimensional view is significant because CPS failures rarely originate from a single component. A system may possess accurate sensors and a reliable controller but still fail because network delays violate control assumptions. Conversely, a secure communication channel may not protect a CPS if compromised machine-learning models generate unsafe control commands. CPS engineering therefore requires simultaneous reasoning about physical dynamics, software, networks, data, human interaction, and assurance.

The motivation for next-generation CPS arises from the limitations of conventional architectures when exposed to scale, heterogeneity, autonomy, and uncertainty. First, the amount and diversity of CPS data have increased dramatically. Cameras, microphones, vibration sensors, thermal sensors, industrial controllers, environmental sensors, and operational databases produce heterogeneous streams at different frequencies. Processing such data exclusively in centralized cloud environments is inefficient for latency-sensitive applications. Edge and fog architectures provide distributed computational resources closer to the physical process, reducing communication delays and potentially improving resilience to intermittent cloud connectivity [3].

Second, modern CPS applications increasingly require autonomous decision-making. In conventional control systems, engineers specify control laws and operating limits explicitly. AI-enabled CPS instead learn patterns from operational data and may adapt to changing conditions. This capability is attractive for predictive maintenance, autonomous driving, energy optimization, process control, and healthcare monitoring. However, the shift from deterministic logic to learned behavior also changes the verification problem. A conventional controller can often be analyzed through formal stability properties, while a neural network may behave unpredictably outside its training distribution. The resulting gap between empirical performance and guaranteed behavior represents one of the central research challenges of intelligent CPS.

Third, the physical environment is inherently uncertain. Sensor noise, component degradation, disturbances, communication losses, weather conditions, human interventions, and unmodeled physical phenomena can cause deviations from nominal operating conditions. A next-generation CPS must therefore combine model-based and data-driven reasoning. Purely model-based systems can struggle with unknown or highly nonlinear processes, while purely data-driven systems may require large quantities of representative data and may fail to respect physical constraints. Hybrid approaches combining physical models with machine learning are consequently receiving increasing attention.

Digital twins provide a particularly useful mechanism for bridging these approaches. A physical process can be mirrored by a dynamic digital representation that receives real-time observations, estimates hidden states, simulates possible future trajectories, and evaluates alternative actions. Digital twins can also support testing without exposing physical equipment to unsafe experiments. However, systematic research indicates that digital-twin-based CPS testing remains an emerging field, with unresolved issues surrounding test oracles, modeling approaches, predictive accuracy, and automated test generation [4].

The communication environment is another important motivation. CPS control depends not merely on whether a message arrives but when it arrives. Variable latency and jitter can destabilize a feedback loop even when packet loss is negligible. TSN research has therefore focused on deterministic communication, time synchronization, traffic scheduling, reliability, and coexistence between information and operational technology [7]. 5G private

networks and emerging 6G technologies add flexible wireless connectivity, network slicing, edge integration, and advanced quality-of-service mechanisms. Nevertheless, wireless communication introduces additional uncertainty due to interference, mobility, channel variation, and security threats.

Cybersecurity presents a particularly strong motivation for architectural transformation. Traditional IT security mechanisms often focus on confidentiality and data integrity, whereas CPS additionally requires physical safety and process integrity. The attack surface includes sensors, actuators, controllers, communication networks, edge devices, cloud services, machine-learning models, digital twins, and human interfaces. Intrusion-response research emphasizes that CPS defense must preserve stable and safe operation even when attacks occur [8]. Consequently, resilience—not merely prevention—must become a primary design objective.

Finally, CPS is increasingly becoming human-centered. Industry 5.0 emphasizes collaboration between humans and intelligent machines, while healthcare, transportation, and smart-city systems directly influence human decisions and well-being. The resulting systems must therefore account for human intent, trust, cognitive workload, privacy, accessibility, and accountability. These considerations motivate the transition from purely autonomous CPS toward adaptive human–cyber–physical ecosystems.

The combined effect of these trends establishes the need for a new CPS paradigm characterized by distributed intelligence, deterministic communication, continuous digital representation, adaptive control, cybersecurity-by-design, and lifecycle assurance. The subsequent taxonomy organizes these capabilities into a coherent framework.

3. Taxonomy / Classification

A useful taxonomy of next-generation CPS should distinguish between architectural layers, computational intelligence, communication mechanisms, physical domains, autonomy levels, and assurance requirements. Existing classifications frequently focus on individual technologies, such as IoT, edge computing, or industrial CPS. Such classifications become insufficient when a single deployment simultaneously incorporates all of these components. A next-generation CPS should instead be classified according to how sensing, communication, computation, intelligence, physical actuation, digital representation, and human participation are integrated.

The first classification dimension concerns the physical scale of the system. Micro-CPSs operate at the device or embedded-system level and include intelligent sensors, wearable devices, actuators, and small robots. Device-level systems generally have severe computational and energy constraints. Meso-CPSs integrate multiple devices into a machine, production cell, vehicle, building, or local energy system. Macro-CPSs encompass large geographical infrastructures such as transportation networks, smart grids, water systems, and smart cities. System-of-systems CPSs connect multiple independently managed CPSs. The distinction is important because increasing scale introduces greater heterogeneity, organizational boundaries, and interoperability requirements.

The second dimension concerns intelligence placement. Centralized CPSs send most data to a central computational platform. Edge-centric CPSs execute time-critical analytics locally. Fog or continuum architectures distribute computation across devices, gateways, edge servers, regional clouds, and central clouds. Fully distributed CPSs allow multiple agents to make coordinated decisions without a single computational center. The appropriate configuration depends on latency, bandwidth, privacy, computational capacity, and resilience requirements.

The third dimension concerns the relationship between models and data. Model-driven CPSs primarily use differential equations, state-space models, finite-state machines, optimization models, and control theory. Data-driven CPSs use machine learning to infer system behavior. Hybrid CPSs combine physical models with learned components. The latter category is increasingly important because physical models provide interpretability and constraints, while machine learning can approximate unknown nonlinearities.

The fourth dimension concerns autonomy. At the lowest level, CPS merely observes physical processes. At higher levels it provides recommendations, executes supervised actions, performs autonomous control, and ultimately coordinates multiple autonomous agents. Autonomy should not be treated as a binary property because different functions within a CPS may have different autonomy levels.

The fifth dimension concerns digital representation. Conventional CPS may maintain a limited system model. Digital-twin-enabled CPS maintain continuously updated digital representations. Multi-twin systems represent

multiple interacting assets. At larger scales, federated or system-level digital twins can represent production lines, cities, or transportation systems.

The sixth dimension concerns assurance. Non-critical CPS may tolerate occasional failures, while safety-critical CPS such as aircraft, medical devices, nuclear systems, and autonomous vehicles require strict assurance. Security-critical CPS such as energy and water infrastructure require protection against malicious manipulation. Mission-critical CPS require availability and resilience under severe disruptions. Next-generation systems increasingly combine all three.

Table 1. Taxonomy of Next-Generation CPS

Classification dimension	Category	Principal characteristics	Typical applications
Physical scale	Micro-CPS	Device-level sensing/control	Wearables, smart sensors
	Meso-CPS	Machine/system-level coordination	Robots, vehicles, production cells
	Macro-CPS	Geographically distributed	Smart grids, transportation
	System-of-systems CPS	Interdependent CPS networks	Smart cities, autonomous mobility
Intelligence placement	Centralized	Cloud/data-center analytics	Batch optimization
	Edge-centric	Local inference/control	Industrial automation
	Fog/continuum	Multi-tier computing	Smart manufacturing
	Distributed	Collaborative agents	Multi-robot systems
Modeling	Model-driven	Physics/control models	Process control
	Data-driven	ML-based models	Fault diagnosis
	Hybrid	Physics + ML	Predictive control
Digital representation	Conventional	Static/limited models	Embedded CPS
	Digital twin	Dynamic synchronized model	Smart manufacturing
	Multi-twin	Interconnected twins	Factory/city systems
Autonomy	Monitoring	Observation only	Environmental sensing
	Advisory	Decision support	Healthcare
	Supervised autonomy	Human-authorized actions	Robotics
	Autonomous	Independent decisions	Autonomous vehicles
Assurance	Safety-critical	Strong safety constraints	Aviation, healthcare
	Security-critical	Adversarial resilience	Power, water
	Mission-critical	Availability/resilience	Emergency systems

The taxonomy also reveals an important architectural trend: categories are not mutually exclusive. A modern autonomous vehicle can simultaneously be a meso-scale CPS, edge-centric, hybrid model/data-driven, digital-twin-enabled, autonomous, and safety-critical. Consequently, research should avoid treating CPS classifications as isolated categories.

A complementary technology taxonomy is presented in Table 2.

Table 2. Technology Taxonomy for Next-Generation CPS

Technology	CPS function	Primary advantage	Principal limitation
IoT sensing	Perception	Large-scale observability	Data heterogeneity
Edge computing	Local computation	Low latency	Resource constraints


Fog computing	Distributed orchestration	Flexible resource placement	Management complexity
5G/private 5G	Wireless connectivity	Low latency/QoS	Deployment and security complexity
TSN	Deterministic networking	Bounded latency/jitter	Configuration complexity
Digital twins	Modeling/simulation	Prediction and testing	Synchronization fidelity
Deep learning	Perception/analytics	High nonlinear modeling capability	Data and explainability
Reinforcement learning	Sequential control	Adaptive optimization	Safety and sample efficiency
Federated learning	Distributed learning	Data locality/privacy	Communication and non-IID data
GNNs	Relational reasoning	Graph-aware modeling	Scalability
Blockchain	Trust/coordination	Distributed integrity	Latency/energy overhead
Foundation models	General-purpose intelligence	Multimodal/generalizable reasoning	Uncertainty and resource cost
Explainable AI	Decision transparency	Improved interpretability	Fidelity/complexity tradeoff
Formal verification	Assurance	Mathematical guarantees	Scalability

This multidimensional taxonomy forms the basis for the architectural and methodological analysis that follows.

4. System Architecture / Algorithms

A representative next-generation CPS architecture can be organized into nine interconnected layers: physical process, sensing and actuation, connectivity, edge computing, cloud/continuum computing, intelligence, digital twin, orchestration/security, and human interaction. Unlike a strictly vertical layered model, these layers should be understood as partially interconnected because modern CPS applications frequently bypass traditional boundaries.

FIGURE 1: PROPOSED REFERENCE ARCHITECTURE FOR NEXT-GENERATION CPS



The physical layer contains the plant, machine, vehicle, infrastructure, human body, or environmental process. Its dynamics determine the fundamental timing and safety requirements. The sensing and actuation layer translates physical phenomena into digital observations and computational commands into physical actions. Sensor fusion

may combine multiple modalities to reduce uncertainty. Actuation is particularly important because it closes the feedback loop and therefore represents one of the most safety-sensitive interfaces.

The edge layer performs low-latency computation. Edge nodes may execute anomaly detection, feature extraction, machine-learning inference, local optimization, and control. Model compression techniques such as pruning, quantization, knowledge distillation, and efficient neural architectures become important because edge devices cannot always support large models.

The cloud and continuum layer provides computational resources for model training, global optimization, historical analytics, and fleet-level coordination. Rather than treating edge and cloud as competing alternatives, next-generation CPS should use adaptive workload placement. A task can be executed at the device if latency requirements are strict, at the edge if moderate computational capacity is required, or in the cloud if computational intensity dominates latency requirements.

The intelligence layer contains AI services. A practical CPS may employ multiple models rather than one universal model. A vision model may identify defects, a time-series model may predict machine failure, a GNN may model system-level relationships, an optimization algorithm may schedule production, and an anomaly detector may identify cyberattacks. The architecture should therefore support model composition and lifecycle management.

The digital-twin layer maintains synchronized representations of physical entities. A useful twin architecture should support four functions: state synchronization, simulation, prediction, and interaction. The twin may receive sensor streams, update model parameters, simulate future trajectories, and expose predictions to control and decision services.

The security and orchestration layer cuts across all components. Identity management, secure boot, encryption, access control, intrusion detection, anomaly detection, software attestation, model integrity, and incident response should be integrated into the architecture rather than added after deployment. NIST CSF 2.0 provides a broad cybersecurity-risk framework that can complement CPS-specific security engineering [11].

A representative intelligent control algorithm can be expressed through the following sequence:

1. Acquire synchronized sensor data.
2. Perform sensor validation and preprocessing.
3. Estimate physical state.
4. Update the digital twin.
5. Detect cyber-physical anomalies.
6. Predict future system states.
7. Generate candidate control actions.
8. Evaluate actions against physical, safety, and security constraints.
9. Execute the selected action.
10. Record outcomes for model adaptation and lifecycle monitoring.

5. Comparative Analysis

The architecture selected for a CPS directly influences latency, scalability, resilience, privacy, energy consumption, and assurance. Centralized cloud architectures are relatively simple to manage because computational resources and models can be maintained at a central location. They are particularly effective for historical analytics, large-scale training, and cross-site optimization. However, they may introduce communication delays and dependence on network availability.

Edge-centric architectures move intelligence toward the physical process. Their principal advantage is reduced response time. This is particularly important for robotic control, industrial safety, autonomous driving, and medical monitoring. Edge architectures also reduce the need to transmit raw high-volume sensor streams. Their limitations include constrained computational capacity, heterogeneous hardware, difficult model updates, and increased security exposure at distributed endpoints.

Fog and continuum architectures provide an intermediate solution by distributing computational tasks among devices, gateways, edge servers, regional infrastructure, and cloud platforms. Their principal strength is flexible resource placement. Their weakness is orchestration complexity. The system must continuously determine where a task should execute while considering latency, bandwidth, energy, cost, privacy, and reliability.

Digital-twin architectures provide capabilities that conventional edge/cloud architectures do not inherently provide. The twin can simulate system evolution and support predictive maintenance and testing. However, its effectiveness depends on model fidelity and synchronization quality. A digital twin that becomes stale may produce systematically misleading recommendations.

AI-native architectures provide adaptive capabilities but create assurance challenges. A deterministic controller can often be verified mathematically, whereas a neural network's behavior may be difficult to characterize under unseen inputs. The appropriate strategy for safety-sensitive CPS is therefore increasingly hybrid: AI components should be embedded within a constrained control architecture rather than given unrestricted authority.

Table 3. Comparison of CPS Architectural Paradigms

Criterion	Centralized cloud	Edge-centric	Fog/continuum	Digital-twin-enabled	AI-native hybrid
Latency	Low-moderate	Very low	Low	Low-moderate	Low
Scalability	High	Moderate	High	High	High
Local autonomy	Low	High	High	High	Very high
Bandwidth efficiency	Low	High	High	High	High
Privacy	Moderate	High	High	Moderate-high	Variable
Computational capacity	Very high	Limited	High	High	High
Physical interpretability	Moderate	High	High	Very high	Variable
Predictive simulation	Moderate	Limited	Moderate	Very high	High
Adaptability	High	High	High	Very high	Very high
Verification difficulty	Moderate	Moderate	High	High	Very high
Deployment complexity	Moderate	High	Very high	Very high	Very high

A second comparison concerns intelligent methodologies.

Table 4. Comparison of Intelligent Methods for CPS

Method	Strength	Weakness	Suitable CPS task
CNN	Spatial feature extraction	Data intensive	Machine vision
LSTM/GRU	Temporal dependencies	Sequential computation	Fault prediction
Transformer	Long-range dependencies	High resource demand	Multimodal/time-series analytics
GNN	Relational modeling	Graph construction complexity	Grid/factory/network CPS
Reinforcement learning	Sequential optimization	Unsafe exploration	Adaptive control
Federated learning	Data locality	Non-IID and communication cost	Distributed monitoring
Physics-informed ML	Physical consistency	Requires domain knowledge	Dynamic modeling
Digital-twin simulation	Predictive evaluation	Synchronization cost	Testing/maintenance
Foundation models	Generalization/multimodality	Reliability and compute	Engineering assistance
Hybrid AI-control	Safety + adaptability	Architecture complexity	Safety-sensitive CPS

6. Applications

6.1 Smart Manufacturing

Smart manufacturing is one of the most mature application areas for CPS. Modern production systems integrate programmable logic controllers, robots, machine vision, autonomous mobile robots, industrial networks, edge servers, enterprise systems, and digital twins. CPS allows production processes to become observable and dynamically optimized. Machine condition can be continuously monitored through vibration, temperature, acoustic, electrical, and vision sensors. Machine-learning models can estimate remaining useful life and support predictive maintenance.

Digital twins are particularly relevant because manufacturing equipment can be represented virtually and evaluated under alternative operating conditions. Production scheduling can therefore incorporate machine health, energy prices, material availability, and expected demand. Recent work also distinguishes digital twins from related concepts such as digital threads and asset administration shells, emphasizing the importance of clear semantic and architectural boundaries [12].

The major limitation is interoperability. Manufacturing plants frequently contain equipment from different vendors and generations. Legacy industrial controllers may not expose data through modern interfaces. Standards such as OPC UA can provide semantic communication, but integration remains difficult at the organizational and lifecycle levels. Recent CPS research has demonstrated the potential of portable OPC-UA-based digital-twin connectors for distributed robotic systems [13].

6.2 Autonomous Transportation

Autonomous vehicles represent highly complex CPS because perception, localization, planning, communication, control, and physical motion must operate under strict safety constraints. Cameras, radar, lidar, GPS, inertial sensors, and vehicle telemetry generate heterogeneous data. AI models interpret these observations and generate decisions, while vehicle controllers execute steering, acceleration, and braking commands.

The central challenge is the mismatch between AI benchmark performance and safety assurance. A vision model can achieve high recognition accuracy but still fail under unusual weather, sensor degradation, adversarial manipulation, or rare events. Foundation models are increasingly being investigated for autonomous driving because of their multimodal reasoning capabilities, but integrating them into safety-critical control requires new validation mechanisms [14].

Vehicle-to-everything communication further extends the CPS boundary. Vehicles may exchange information with infrastructure, pedestrians, other vehicles, and cloud services. This creates opportunities for cooperative perception and traffic optimization but simultaneously increases cybersecurity and privacy risks.

6.3 Smart Grids

Smart grids represent macro-scale CPS in which generation, transmission, distribution, storage, and consumption interact dynamically. Phasor measurement units, smart meters, distributed energy resources, batteries, and automated substations generate large quantities of time-series data. CPS intelligence can support load forecasting, fault detection, demand response, voltage control, and renewable-energy integration.

The distributed nature of smart grids makes graph-based learning particularly attractive. The grid can be represented as a graph whose nodes correspond to buses and whose edges represent transmission lines. GNNs can capture spatial dependencies that conventional independent time-series models cannot represent efficiently.

Cybersecurity is critical because false measurements or malicious control commands can cause physical instability. Secure state estimation, anomaly detection, resilient control, and coordinated incident response therefore require joint cyber-physical modeling.

6.4 Healthcare

Healthcare CPS combines physiological sensing, wearable devices, medical equipment, clinical software, communication infrastructure, and human decision-making. Remote patient monitoring is a prominent example. Sensors may collect heart rate, oxygen saturation, glucose, blood pressure, motion, and other signals. Edge intelligence can identify abnormal conditions locally, while cloud platforms provide longitudinal analysis.

Healthcare CPS differs from industrial CPS because human safety and privacy dominate performance considerations. False positives can produce unnecessary interventions, while false negatives can delay treatment. Models must therefore be calibrated and clinically validated. Privacy-preserving approaches such as federated learning are attractive because medical data may remain within institutional boundaries.

6.5 Smart Agriculture

Agricultural CPS integrates soil sensors, weather stations, irrigation controllers, drones, satellite imagery, farm machinery, and decision-support systems. Smart irrigation can combine soil moisture observations, weather predictions, crop models, and learned evapotranspiration estimates to determine irrigation schedules.

The major challenge is environmental heterogeneity. A model trained on one soil type, climate, or crop may perform poorly elsewhere. This makes transfer learning, federated learning, domain adaptation, and physics-informed modeling especially relevant. Agricultural CPS must also operate under intermittent connectivity and constrained energy availability, increasing the importance of edge computing.

6.6 Smart Buildings and Cities

Buildings can be treated as CPS when heating, ventilation, air conditioning, lighting, occupancy, security, and energy systems are dynamically coordinated. AI-based building control can optimize comfort and energy consumption simultaneously. At city scale, CPS integrates transportation, energy, water, waste, emergency response, and public infrastructure.

The city-scale problem is fundamentally a system-of-systems problem. Different agencies and infrastructures may have separate ownership, standards, and data policies. Interoperability therefore becomes as important as algorithmic intelligence.

7. Case Study: Intelligent Digital-Twin-Enabled Industrial CPS

Consider a manufacturing cell containing a CNC machine, robotic manipulator, conveyor, machine-vision camera, vibration sensors, temperature sensors, programmable logic controller, edge gateway, and cloud analytics platform. The objective is to develop a CPS capable of real-time monitoring, predictive maintenance, anomaly detection, adaptive production scheduling, and secure operation.

The physical layer generates vibration, current, temperature, position, and vision data. The edge gateway synchronizes these streams and performs preprocessing. A lightweight anomaly-detection model continuously estimates the probability of abnormal operation. A digital twin receives the same data and estimates machine health and production state.

Cybersecurity is integrated into the same architecture. Suppose an attacker manipulates a temperature sensor. A purely data-driven anomaly detector may detect an unusual value, but a digital twin can additionally determine whether the observed temperature is physically consistent with machine load and power consumption. The combination of cyber and physical consistency therefore provides a stronger detection mechanism.

A representative experimental architecture would use an industrial PC or GPU-enabled edge server, OPC UA for semantic communication, MQTT for selected telemetry streams, TSN-enabled Ethernet for deterministic control traffic, a digital-twin environment for simulation, and containerized deployment for model lifecycle management. The cloud environment would be responsible for historical analytics and model retraining, whereas the edge node would execute safety-critical inference.

The case demonstrates the central principle of next-generation CPS: no individual technology is sufficient. Predictive maintenance without secure sensing may be manipulated; digital twins without reliable synchronization may become inaccurate; AI without physical constraints may generate unsafe decisions; and cloud analytics without edge control may violate latency requirements. The value emerges from the integration of these technologies.

8. Tools, Technologies and Implementation

Implementation of next-generation CPS requires a heterogeneous software and hardware ecosystem. At the physical layer, microcontrollers, PLCs, industrial PCs, robotics platforms, embedded GPUs, smart sensors, and actuators provide the computational and control foundation. Platforms such as Raspberry Pi, NVIDIA Jetson, industrial edge computers, ARM-based systems, and FPGA-based controllers are commonly appropriate depending on computational and timing requirements.

Communication technologies include Ethernet, industrial Ethernet, Wi-Fi, Bluetooth Low Energy, LoRaWAN, 5G, private 5G, and TSN. TSN is particularly important for applications requiring deterministic communication because it supports time synchronization, traffic scheduling, and bounded latency [7]. In wireless industrial environments, private 5G can complement TSN by providing flexible connectivity for mobile devices and robots. Middleware is responsible for data exchange and service integration. MQTT is useful for lightweight publish/subscribe telemetry, while OPC UA provides richer industrial interoperability and semantic modeling. DDS is suitable for distributed real-time systems requiring data-centric communication. ROS 2 is widely used for robotics and supports distributed robotic applications.

Digital-twin implementation can use simulation and modeling environments such as MATLAB/Simulink, Modelica-based tools, Unity, Unreal Engine, Gazebo, and specialized industrial platforms. The choice depends on whether the objective is physical simulation, visualization, control design, robotics testing, or industrial deployment.

Machine-learning implementation can use PyTorch, TensorFlow, scikit-learn, ONNX, TensorRT, and related inference frameworks. Edge deployment frequently requires converting large models into optimized representations. Quantization can reduce precision from floating-point representations to lower-bit formats, while pruning can remove redundant parameters. Knowledge distillation can produce a smaller student model that approximates a larger teacher model.

Cybersecurity implementation should include secure boot, hardware roots of trust, certificate-based authentication, network segmentation, encrypted communication, access control, software signing, vulnerability management, intrusion detection, and incident response. NIST CSF 2.0 provides a broad organizational framework for cybersecurity risk management [11], while CPS-specific protection research emphasizes the need to consider both cyber and physical effects [8], [15].

Table 5. Representative Implementation Stack

Layer	Technologies/tools	Typical function
Sensors	MEMS, cameras, lidar, industrial sensors	Physical observation
Controllers	PLC, MCU, FPGA, RTOS	Real-time control
Edge	Industrial PC, Jetson, ARM	Local AI/control
Messaging	MQTT, DDS	Telemetry/data exchange
Industrial interoperability	OPC UA	Semantic interoperability
Deterministic networking	TSN	Low-latency control
Wireless	5G/private 5G, Wi-Fi	Mobile connectivity
AI	PyTorch, TensorFlow	Model development
Edge inference	ONNX, TensorRT	Optimized deployment
Digital twin	MATLAB, Simulink, Modelica, simulation platforms	Modeling/simulation
Containers	Docker/Kubernetes	Deployment/orchestration
Monitoring	Prometheus/Grafana	Operational monitoring
Security	PKI, TLS, IDS, secure boot	Protection
Cloud	AWS/Azure/GCP/private cloud	Training/analytics

The implementation challenge is not simply selecting tools. It is ensuring that interfaces between tools preserve timing, semantics, security, and lifecycle information. A model developed in a laboratory environment may not

execute with the same timing or numerical characteristics on an embedded platform. Consequently, deployment validation must include hardware-in-the-loop, software-in-the-loop, and digital-twin-based testing.

9. Performance Evaluation

Performance evaluation of CPS requires metrics that span computation, communication, physical control, intelligence, reliability, and security. Traditional machine-learning metrics such as accuracy, precision, recall, and F1-score are insufficient because a model can achieve high classification accuracy while violating real-time constraints.

Recent benchmark studies illustrate both the usefulness and limitations of public datasets. SWaT and WADI are widely used for industrial anomaly detection, and recent deep-learning systems report strong benchmark performance [16]. However, benchmark performance can be misleading if models are evaluated under identical data distributions. Real deployments exhibit concept drift, sensor replacement, maintenance events, process changes, and previously unseen attack patterns. Synthetic dataset generation has consequently become an active research area, although synthetic data must be validated against physical realism [17].

Table 6. Recommended CPS Evaluation Framework

Dimension	Metrics
Control	IAE, ISE, settling time, overshoot
Communication	Mean latency, worst-case latency, jitter, packet loss
AI	Accuracy, precision, recall, F1, AUROC
Anomaly detection	Detection delay, false-positive rate
Predictive maintenance	RUL error, warning horizon
Energy	Joules/task, decisions/Joule
Reliability	MTBF, MTTR, availability
Security	Detection rate, attack impact, recovery time
Digital twin	Synchronization error, model fidelity
Scalability	Throughput, resource utilization
Safety	Constraint violations, unsafe-state probability
Human interaction	Workload, response time, trust calibration

A rigorous evaluation should therefore combine laboratory benchmarks, simulation, digital-twin testing, hardware-in-the-loop experiments, and real-world pilot deployments. This multi-level evaluation is necessary because performance measured at one abstraction level may not translate to another.

10. Challenges and Limitations

The first major challenge is real-time intelligence. Deep neural networks and foundation models can be computationally expensive, while CPS often requires deterministic response times. Model compression can reduce computation but may degrade accuracy. The resulting problem is a three-way trade-off among model complexity, prediction quality, and latency.

The second challenge is data quality. CPS data are noisy, asynchronous, incomplete, and often non-stationary. Sensors may drift, devices may be replaced, and physical operating conditions may change. Machine-learning systems trained under historical distributions can therefore experience concept drift. Continual learning provides one potential solution, but uncontrolled adaptation may introduce catastrophic forgetting or unexpected behavior.

Third, CPS requires strong guarantees. Conventional machine learning primarily optimizes statistical performance, whereas CPS engineering requires safety constraints. A classifier with 99% accuracy can still be unacceptable if the remaining 1% contains dangerous failure cases. Research must therefore move toward risk-sensitive learning, formal verification, runtime assurance, and uncertainty-aware decision-making.

Fourth, interoperability remains unresolved. CPS deployments combine legacy equipment, proprietary protocols, modern cloud platforms, and heterogeneous AI frameworks. Digital twins add additional semantic requirements. Recent research identifies interoperability challenges at technical, syntactic, semantic, pragmatic, dynamic, and organizational levels [10]. Solving interoperability therefore requires more than standardizing communication protocols.

Fifth, cybersecurity is becoming increasingly difficult because the attack surface is expanding. AI models themselves may become attack targets through data poisoning, adversarial examples, model extraction, backdoors, and prompt injection in foundation-model-based systems. Digital twins can also become attack surfaces if their models or synchronization channels are compromised.

Sixth, explainability remains difficult for complex models. A human operator may need to know why an AI system recommended shutting down a machine. Post-hoc explanation methods can provide useful insights but do not necessarily guarantee that the explanation accurately represents the model's internal decision process.

Seventh, CPS systems often contain resource-constrained devices. Energy, memory, storage, and processing limitations restrict the complexity of algorithms that can be deployed at the edge. This limitation becomes particularly significant for distributed AI and foundation models.

Eighth, data sharing remains problematic. Industrial organizations may consider sensor data commercially sensitive, healthcare organizations must protect patient data, and critical infrastructure operators may restrict data access for security reasons. Federated learning can reduce raw-data sharing, but it introduces communication overhead and vulnerability to malicious participants.

Ninth, digital twins themselves may suffer from model mismatch. A twin is only useful if its virtual state accurately corresponds to the physical system. Aging equipment, undocumented modifications, environmental changes, and sensor calibration errors can cause divergence.

Finally, there is an assurance gap for autonomous CPS. As systems become increasingly autonomous, responsibility becomes distributed across software developers, hardware manufacturers, operators, AI model providers, infrastructure providers, and users. Technical assurance must therefore be complemented by organizational and governance mechanisms.

11. Conclusion & Future Research

Cyber-Physical Systems are evolving into distributed intelligent ecosystems through the integration of AI, edge computing, digital twins, advanced networks, autonomous systems, and distributed learning. The key challenge is no longer technology availability but trustworthy integration. Next-generation CPS must combine real-time control, synchronized digital representations, low-latency computing, cybersecurity, and physical-system awareness. Hybrid architectures integrating physics-based models, machine learning, digital twins, edge intelligence, and formal assurance are particularly promising. Future systems should emphasize trustworthy autonomy, resilience, interoperability, explainability, predictable timing, and safe adaptation. Ultimately, next-generation CPS should deliver not merely greater intelligence and connectivity, but **trustworthy, adaptive, resilient, secure, and sustainable integration between cyber intelligence and physical reality**.

Table 7. Priority Research Directions

Research direction	Current limitation	Desired capability
Physics-grounded foundation models	Hallucination and weak physical grounding	Physically consistent reasoning
Autonomous digital twins	Static/partially synchronized models	Self-updating predictive twins
Self-healing CPS	Manual recovery	Safe autonomous reconfiguration
AI verification	Weak formal guarantees	Provably bounded behavior
Semantic interoperability	Heterogeneous meanings	Machine-understandable semantics

6G CPS	Emerging infrastructure	Integrated sensing/computing/communication
Federated CPS	Non-IID and poisoning	Robust decentralized learning
Secure AI	New attack surfaces	AI-native cyber resilience
Uncertainty-aware control	Point predictions	Risk-aware decisions
Sustainable CPS	High computational energy	Energy-aware intelligence
Human–AI collaboration	Automation bias	Calibrated human supervision
Benchmarking	Fragmented datasets	Standardized cyber-physical benchmarks

The overarching research direction is therefore the development of CPS that can learn and adapt without compromising safety, security, interpretability, and human agency.

References

- [1] E. R. Griffor, C. Greer, D. A. Wollman, and M. J. Burns, Framework for Cyber-Physical Systems: Volume 1, Overview, NIST Special Publication 1500-201, National Institute of Standards and Technology, 2017.
- [2] C. Greer, M. Burns, D. Wollman, and E. Griffor, Cyber-Physical Systems and Internet of Things, NIST Special Publication 1900-202, National Institute of Standards and Technology, 2019.
- [3] M. Satyanarayanan, "The emergence of edge computing," Computer, vol. 50, no. 1, pp. 30–39, 2017.
- [4] S. M. H. S. Al-Seghayer et al., "Digital-twin-based testing for cyber-physical systems: A systematic literature review," Information and Software Technology, vol. 156, 107145, 2023.
- [5] O. Baris, Y. Chen, G. Dong, L. Han, T. Kimura, P. Quan, R. Wang, T. Wang, T. Abdelzaher, M. Bergés, P. P. Liang, and M. Srivastava, "Foundation Models for CPS-IoT: Opportunities and Challenges," arXiv preprint arXiv:2501.16368, 2025.
- [6] C. Lu, P. Valle, J. Wu, E. Isaku, H. Sartaj, A. Arrieta, and S. Ali, "Foundation Models for Software Engineering of Cyber-Physical Systems: The Road Ahead," arXiv preprint arXiv:2504.04630, 2025.
- [7] P. Pop, M. L. Raagaard, M. Gutiérrez, and W. Steiner, "Timely survey of Time-Sensitive Networking: Past and future directions," IEEE Access, vol. 9, pp. 142506–142527, 2021.
- [8] M. Bashendy, A. Tantawy, and A. Erradi, "Intrusion response systems for cyber-physical systems: A comprehensive survey," Computers & Security, vol. 124, 102984, 2023.
- [9] M. Goh, S. Adepu, K. N. Junejo, and A. Mathur, "A dataset to support research in the design of secure water treatment systems," in Critical Information Infrastructures Security, Springer, 2016.
- [10] "Interoperability levels and challenges of digital twins in cyber-physical systems," Internet of Things and Cyber-Physical Systems, 2024.
- [11] National Institute of Standards and Technology, The NIST Cybersecurity Framework (CSF) 2.0, NIST Cybersecurity White Paper 29, 2024.
- [12] F. P. M. A. et al., "Clarifying concepts of Metaverse, Digital Twin, Digital Thread and AAS for CPS-based production systems," IFAC-PapersOnLine, vol. 56, no. 2, pp. 6351–6357, 2023.
- [13] "Systems of Digital Twins and Physical Systems: Interoperability, Decentralization, and Mobility in Robotic Applications," in Proc. 28th International Conference on Methods and Models in Automation and Robotics, 2024.
- [14] X. Wu, M. Cheng, Q. Hu, J. Chen, Y. Huang, M. Okada, M. Hayashi, T. Tsuchiya, X. Xie, and L. Ma, "Foundation Models for Autonomous Driving System: An Initial Roadmap," arXiv preprint arXiv:2504.00911, 2025.
- [15] "Industrial cyber-physical systems protection: A methodological review," Computers & Security, vol. 135, 103531, 2023.
- [16] "AI-driven real-time anomaly detection in industrial cyber-physical systems using hybrid deep learning with SWaT and WADI benchmarks," Internet of Things and Cyber-Physical Systems, 2026.
- [17] Y. Acquah and K. Roy, "Realistic synthetic dataset generation for cyber-physical systems: A performance evaluation," Discover Applied Sciences, vol. 7, 719, 2025.
- [18] S. Ali, P. Arcaini, and A. Arrieta, "Foundation Models for the Digital Twin Creation of Cyber-Physical Systems," arXiv preprint arXiv:2407.18779, 2024.
- [19] "Edge AI: A survey," Internet of Things and Cyber-Physical Systems, vol. 3, pp. 71–92, 2023.
- [20] "A systematic analysis of enhancing cyber security using deep learning for cyber physical systems," IEEE Access, vol. 12, pp. 6017–6035, 2024.
- [21] "Communications Security in Industry X: A Survey," IEEE Open Journal of the Communications Society, vol. 5, pp. 982–1025, 2024.
- [22] "Survey on Foundation Models for Prognostics and Health Management in Industrial Cyber-Physical Systems," IEEE Transactions on Industrial Cyber-Physical Systems, vol. 2, pp. 264–280, 2024.
- [23] "Cyber-Physical Systems and Artificial Intelligence: State of Art and Research Agenda," in Proc. IEEE International Multi-Conference on Smart Systems & Green Process, 2024.
- [24] "Time-Sensitive Networking: Standards and Application Prospects," in Proc. International Conference on Information Processing and Network Provisioning, 2023.
- [25] "Leveraging digital twins for advanced threat modeling in cyber-physical systems cybersecurity," International Journal of Information Security, vol. 24, Art. 151, 2025.
- [26] "A Literature Review of Digital Twins and Their Applications in the Internet of Things Based on Artificial Intelligence," in Proc. International Conference on Control, Automation and Diagnosis, 2024.
- [27] "Future Directions of Cybersecurity in Industrial Internet of Things Through Edge Computing," in Proc. International Conference on Computer Science and Engineering, 2024.
- [28] F. Li, X. Wu, and H. Han, Data-Driven Cyber Physical Systems, Springer, 2025.
- [29] "Integration of digital twins and physical AI in cyber-physical systems," Internet of Things and Cyber-Physical Systems, 2026.
- [30] "An intelligent intrusion detection system for cyber-physical systems using GAN-LSTM networks," Internet of Things and Cyber-Physical Systems, 2025.