INTERNATIONAL JOURNAL OF SCIENTIFIC INFORMATION

Review | ISSN: 2583-8512

Artificial Intelligence-Enabled Digital Twins for Smart Computing and Real-Time Decision-Making

Sakshi Sharma

Department of Computer Science, Jayoti Vidyapeeth Women's University, Jaipur, Rajasthan, India

Email: sakkshi.sharma1991@gmail.com

Abstract

Artificial intelligence (AI) and digital twin (DT) technologies are increasingly converging to create computational replicas that do more than visualize physical systems. Modern AI-enabled DTs integrate sensor streams, edge-cloud computing, simulation, machine learning, optimization and human feedback to support monitoring, prediction and action in near real time. This review examines how smart computing architectures transform DTs from passive digital shadows into adaptive decision-support systems for computer science and cyber-physical applications. A structured literature synthesis was conducted around peer-reviewed studies on AI in digital twins, real-time analytics, edge intelligence, cybersecurity, smart manufacturing, healthcare, smart cities and energy systems. The paper proposes a layered conceptual architecture, compares AI techniques used across DT functions, and analyses the decision cycle from sensing to actuation. The findings show that supervised learning and deep learning are dominant for predictive maintenance and anomaly detection, while reinforcement learning, optimization, federated learning, knowledge graphs and explainable AI are becoming essential for autonomous, privacy-aware and trustworthy decision-making. However, major barriers remain: data quality, model drift, interoperability, validation, cybersecurity exposure, explainability, latency, regulatory compliance and human accountability. The review argues that future research should move from isolated demonstrators toward standardized, secure, explainable and self-adaptive digital twins that combine physics-based models, data-driven intelligence and governance-by-design. The study contributes a computer-science-oriented synthesis and a future research agenda for AI-DT systems intended to make reliable real-time decisions in complex environments.

Keywords: *Artificial intelligence; digital twin; smart computing; real-time decision-making; cyber-physical systems; edge computing; machine learning; explainable AI*

Address for Correspondence
Sakshi Sharma
Department of Computer Science,
Jayoti Vidyapeeth Women's University (JVWU),
Jaipur, India
sakkshi.sharma1991@gmail.com

Article Received on: 12.03.26

Revised on: 27.03.26

Approved for publication: 15.04.26

How to Cite this Article: Sharma S. Artificial Intelligence-Enabled Digital Twins for Smart Computing and Real-Time Decision-Making. Int., J. Sci. Info. 2026; 4 (1): 80-95

1. Introduction

Digital twins have become one of the most important computational ideas in contemporary computer science because they connect software representations with the behaviour of physical, cyber and socio-technical systems. A digital twin is generally understood as a dynamic virtual representation of a real entity that is updated through data exchange and used for monitoring, simulation, prediction or control. Earlier digital models were often static, offline and descriptive. By contrast, a mature digital twin is synchronized with a physical twin through live data, uses computational models to estimate current and future states, and returns insights or actions that influence operational decisions. This bidirectional relationship gives digital twins a central role in cyber-physical systems, Industry 4.0, smart cities, healthcare, energy, robotics and next-generation networks [1-5].

Artificial intelligence increases the value of digital twins by allowing them to learn from data, recognize patterns, predict system degradation, detect abnormal behaviour, optimize resource allocation and recommend interventions under uncertainty. In computer science terms, the DT becomes a distributed intelligent system in which sensors, stream processing, databases, machine-learning pipelines, simulation engines, semantic models and user interfaces work as an integrated decision loop. This evolution is important because many modern systems are too complex, fast and data-intensive for manual monitoring alone. Manufacturing plants produce high-frequency machine signals, hospitals generate heterogeneous patient data, energy grids balance volatile renewable sources, and networked vehicles exchange time-sensitive information. In such settings, real-time decisions require low latency, contextual intelligence and trustworthy reasoning [6-9].

The selected topic, AI in digital twins for smart computing and real-time decision-making, is involves software architecture, data engineering, artificial intelligence, distributed computing, human-computer interaction, cybersecurity, cloud and edge systems, knowledge representation and algorithmic governance. The scientific problem is not simply whether AI can improve a digital twin. The deeper issue is how AI should be embedded inside the twin so that the resulting system is accurate, secure, explainable, scalable and useful for decision-makers. Poorly designed AI-DT systems may create false confidence, propagate biased data, expose sensitive operational information or issue unsafe recommendations. Consequently, research must balance computational performance with validation, accountability and human oversight [10-13].

2. Background

The digital twin concept has been shaped by several related terms: digital model, digital shadow, cyber-physical system and virtual commissioning. A digital model is usually an offline representation with no automatic data flow from the real system. A digital shadow receives data from the physical system but may not return decisions or actions. A digital twin, in the stronger sense used in this paper, has automatic or semi-automatic synchronization and supports reasoning about the physical entity across its lifecycle. This distinction is important because many systems called digital twins are actually dashboards, simulations or monitoring platforms. The real research challenge begins when the twin must remain synchronized, learn from evolving data and influence operational choices under time constraints [1,3].

In computer science, a digital twin may be described as a multi-layer information system. At the lowest level, sensors and controllers measure state variables such as temperature, vibration, pressure, location, energy consumption, network traffic or patient indicators. Communication technologies such as MQTT, OPC UA, REST APIs, 5G, Wi-Fi, industrial Ethernet and message brokers move these data to edge or cloud platforms. Data engineering components clean, timestamp, align and transform the streams. The twin then combines data-driven and physics-based models to estimate the current state and simulate future behaviour. Above this modelling layer, AI modules perform classification, forecasting, anomaly detection, optimization, reinforcement learning, natural-language interaction or causal reasoning. Finally, decision services translate model outputs into alerts, recommended actions, control policies or visual explanations for humans [5,14].

Smart computing refers to the ability of a computing environment to adapt its processing, storage, communication and decision logic to the needs of the task. In AI-enabled digital twins, smart computing appears in several forms. Edge AI supports low-latency inference near the physical asset. Cloud computing provides scalable storage, simulation and model training. Federated learning allows distributed models to learn without centralizing sensitive data. Knowledge graphs connect heterogeneous entities and make relationships machine-readable. Explainable AI helps users understand why a twin recommends a particular decision. Reinforcement learning and optimization transform the twin from a predictive engine into a prescriptive system that can search for better actions [6,8,15].

Real-time decision-making should not be interpreted as instantaneous automation in every case. In safety-critical environments, a digital twin may provide near-real-time situational awareness while a human remains responsible for the final decision. In other settings, the system may autonomously adjust machine speed, allocate network resources or schedule maintenance tasks. The correct level of autonomy depends on domain risk, regulatory context, model confidence, the cost of error and the availability of human expertise. Therefore, the most mature AI-DT systems are not necessarily the most autonomous ones. They are the systems that match decision speed, transparency and control to the operational context.

The literature shows that AI-DT integration is moving through four broad stages. The first stage uses AI for modelling the physical system. The second stage uses AI for mirroring and synchronization through state estimation and data fusion. The third stage uses AI for intervention through prediction, diagnosis and optimization. The fourth stage explores autonomous management through intelligent agents, foundation models and self-learning policies. These stages are not strictly sequential; many practical systems combine a mature monitoring twin with an emerging optimization module. Nevertheless, the staged view helps clarify how AI changes the purpose of a digital twin from representation to reasoning, and from reasoning to action [7,16].

3. Review Methodology

The search logic focused on combinations of terms such as artificial intelligence, digital twin, machine learning, edge computing, real-time analytics, smart computing, cyber-physical systems, explainable AI, federated learning, smart manufacturing, healthcare digital twin, smart city digital twin, energy digital twin and cybersecurity digital twin. Preference was given to journal articles, systematic reviews, standards documents and highly relevant conference papers published from 2018 onward, with special emphasis on 2020-2026 sources.

The inclusion criteria were: (i) the study discussed digital twins with a live-data, simulation, AI, decision-support or cyber-physical component; (ii) the work was related to computer science, software architecture, AI algorithms, data engineering, cybersecurity, smart systems or real-time analytics; and (iii) the source provided conceptual, methodological or application-specific value. Studies were excluded when they used the term digital twin only as a marketing label, lacked a connection to computational methods, or focused only on static 3D visualization. Standards and institutional reports were included when they clarified interoperability, architecture or implementation challenges.

The selected studies were coded into six themes: enabling technologies, AI techniques, architecture, application domain, real-time decision function and trustworthiness challenge. The review also extracted the decision type supported by each system: monitoring, diagnosis, prediction, optimization, actuation or autonomous learning. The coding was qualitative because the reviewed literature uses different experimental settings, datasets, performance metrics and domain assumptions. As a result, the graphs in this paper represent synthesis scores and conceptual trends rather than direct statistical meta-analysis. This approach is appropriate for a rapidly developing research field where terminology, maturity and implementation detail remain inconsistent.

4. AI-Enabled Digital Twin Architecture for Smart Computing

A practical AI-enabled digital twin can be understood as a layered architecture. The physical layer contains the asset, process or environment being twinned. It may be a production line, a patient pathway, a building, a robot, a smart intersection, a power grid component or a network service. The sensing layer captures raw data through IoT sensors, controllers, logs, cameras, enterprise databases or human input. This layer must solve problems of sampling frequency, timestamp accuracy, missing values, sensor drift, device identity and secure communication. Without reliable sensing, even the most sophisticated AI model will produce weak decisions.

The connectivity and edge layer moves data from the physical system to computational services. Real-time twins cannot depend only on distant cloud infrastructure because network delay, bandwidth cost and privacy constraints may prevent rapid feedback. Edge computing places inference, filtering and short-term storage closer to the source. For example, a vibration anomaly can be detected on an industrial gateway before full data are archived in the cloud. In healthcare, patient data may be processed locally or through privacy-preserving federated learning before aggregated insights are shared. In smart traffic, edge servers near intersections can update signal strategies faster than a centralized platform [8,15].

The data-fabric layer performs ingestion, cleaning, integration, feature generation and governance. This layer is often underestimated, but it determines whether the twin can maintain a truthful state. Digital twins require data from multiple sources, including time-series sensors, enterprise systems, engineering models, maintenance logs, geographic information systems, electronic health records, network telemetry and simulation outputs. Semantic models, ontologies and knowledge graphs help integrate such heterogeneous data because they represent entities and relationships explicitly. They also support reasoning, query answering and interoperability across tools [17,18].

The modelling layer is the core of the twin. It may include physics-based equations, agent-based simulation, discrete-event simulation, finite-element models, statistical models, machine-learning surrogates or hybrid models. Hybridization is especially important because physical principles provide interpretability and stability, while data-driven learning captures patterns that are difficult to model analytically. A smart building twin, for instance, may combine heat-transfer models with neural-network forecasts of occupancy and energy use. A manufacturing twin may combine process simulation with anomaly-detection models trained on machine sensor data. The modelling layer must also include uncertainty representation, because decision-makers need confidence intervals or risk scores rather than single deterministic outputs.

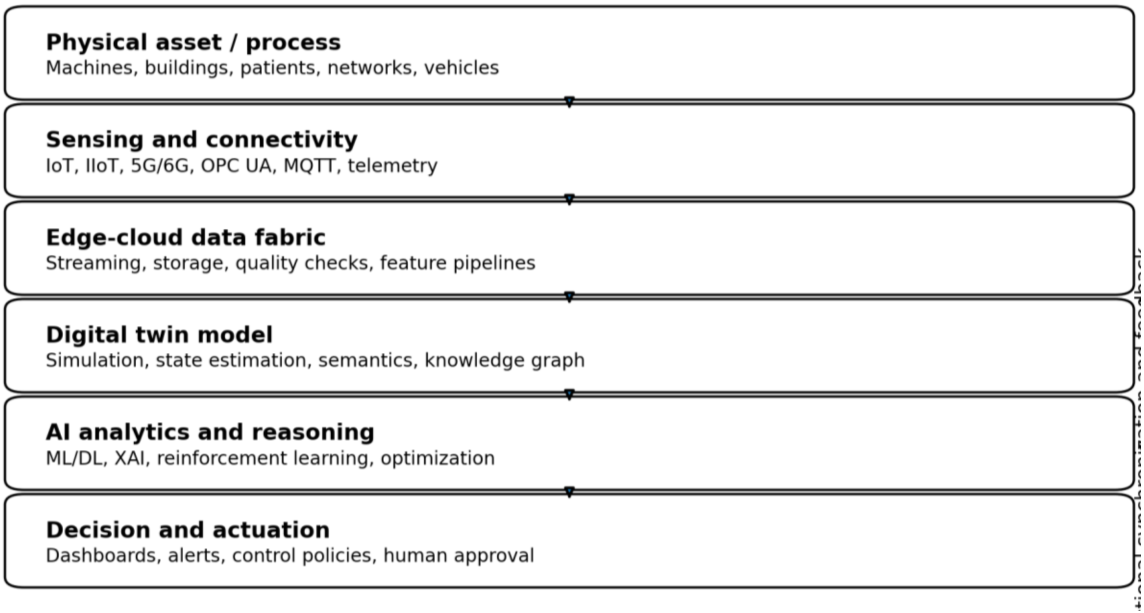

Figure 1: The proposed architecture. It emphasizes that real-time decision-making is not a single algorithm but an ecosystem of data, models, computation, user interaction and feedback. The quality of the final decision depends on all layers working together.

The AI analytics layer transforms the twin into a decision-support system. Supervised learning can classify faults and forecast quality outcomes. Unsupervised learning can discover abnormal operating regimes. Deep learning can process images, vibration sequences, spatial-temporal signals and multimodal data. Reinforcement learning can explore control strategies in simulation before deployment. Optimization algorithms can find feasible schedules, resource allocations or maintenance plans. Explainable AI can reveal which variables influenced a prediction. Large language models and agentic AI are beginning to support documentation, scenario generation, conversational querying and operator assistance, but they require strong grounding and safety controls before they can be trusted in operational settings [7,16].

The decision layer presents insights to humans or sends control actions to the physical system. It includes dashboards, alerts, APIs, automated controllers, workflow systems, digital work instructions and audit logs. Human-computer interaction matters here because a technically accurate recommendation may still fail if it is hard to interpret, poorly timed or not aligned with user responsibility. Decision services should therefore provide context, confidence, explanation, possible alternatives and expected consequences. In high-risk settings, human-in-the-loop or human-on-the-loop control should be preserved.

Finally, the feedback and learning layer closes the loop. After an action is taken, the physical response is observed and used to update models, rules or parameters. This continuous learning capability distinguishes AI-enabled DTs from conventional analytics systems. However, it also creates governance risks: model drift must be detected, updates must be validated, and unsafe policies must be blocked. A reliable architecture therefore combines automation with version control, monitoring, model cards, access control, cybersecurity testing and periodic human review.

5. AI Techniques for Real-Time Digital Twin Decision-Making

AI techniques contribute differently across the digital twin lifecycle. For monitoring, statistical process control, clustering and anomaly detection identify deviations from normal behaviour. These methods are useful when labelled failure data are rare. Autoencoders, isolation forests, one-class support vector machines and density-based clustering have been widely used for anomaly detection in industrial and cyber-physical contexts. For diagnosis, supervised classifiers such as random forests, support vector machines, convolutional neural networks and graph neural networks can map sensor patterns to fault categories, security attacks or clinical states. The effectiveness of these models depends on labelled data, balanced classes and domain-specific feature engineering.

For prediction, time-series forecasting models estimate future states such as remaining useful life, traffic density, energy demand, patient deterioration, temperature deviation or network load. Classical models remain useful where data are limited, but recurrent neural networks, temporal convolutional networks, transformers and hybrid physics-informed models are increasingly used for nonlinear and multimodal systems. Prediction is the bridge between monitoring and decision-making because it gives operators time to act before a failure, congestion event or security incident occurs.

For optimization, AI-enabled twins can evaluate alternative actions through simulation and search. Genetic algorithms, particle swarm optimization, Bayesian optimization, mixed-integer programming and model predictive control are used to identify action plans under constraints. Reinforcement learning is particularly attractive because an agent can learn policies through repeated interaction with a simulated environment. A digital twin provides a safer training environment than a real plant, hospital or road network. However, policies learned in simulation may fail under real-world uncertainty, which means sim-to-real validation is essential before deployment.

For distributed privacy-aware learning, federated learning allows multiple organizations, devices or edge nodes to train shared models without centralizing raw data. This is useful for healthcare, smart grids and industrial IoT, where privacy, ownership and bandwidth restrictions limit data sharing. Digital twins can assist federated learning by representing device states, communication conditions and resource constraints. Conversely, federated learning can improve digital twins by enabling collaborative model training across similar assets or sites [19,20].

For knowledge representation, ontologies and knowledge graphs provide a semantic layer that connects assets, sensors, events, models, decisions and constraints. They are important because real-time decisions are rarely based on raw sensor values alone. A machine alarm must be connected to maintenance history, operating mode, product batch, spare-parts availability and safety rules. Knowledge graphs support such contextual reasoning and make AI outputs easier to interpret. They also address interoperability, which is one of the most persistent obstacles in DT deployment [17,18].

For trustworthiness, explainable AI, uncertainty quantification and causal reasoning are increasingly important. Decision-makers need to know not only what the system predicts but why it predicts it, how confident it is, what evidence supports the output and what could change the recommendation. Explainability is especially necessary in healthcare, cybersecurity and safety-critical industrial control. The challenge is to provide explanations that are technically correct and understandable to users. A feature-importance plot may be useful for a data scientist but insufficient for a maintenance engineer or clinician who needs actionable context.

The maturity of AI techniques also differs by decision function. Monitoring and prediction are relatively mature because they can operate with historical data and measurable performance metrics. Optimization and autonomy remain harder because they require causal understanding, safe exploration, constraint handling and responsibility allocation. Figure 2 presents a qualitative synthesis of AI contribution across five real-time decision functions. The scores indicate that prediction and optimization currently receive the strongest research attention, while full autonomy remains less mature because of validation and governance concerns.

Table 1: AI techniques and their main roles in digital twins.

AI technique	Typical DT role	Decision contribution
Supervised learning	Fault classification, quality prediction, risk scoring	Supports diagnosis and targeted intervention
Unsupervised learning	Anomaly detection, clustering, novelty discovery	Identifies abnormal states without many labels
Deep learning	Image, video, vibration, multimodal and time-series modelling	Improves prediction in nonlinear environments
Reinforcement learning	Policy learning through simulated interaction	Supports adaptive control and optimization
Federated learning	Distributed model training across sites or devices	Improves privacy-aware collaboration
Knowledge graphs	Semantic integration of assets, events and constraints	Improves context, reasoning and interoperability
Explainable AI	Feature attribution, rule extraction, confidence communication	Improves trust and accountability

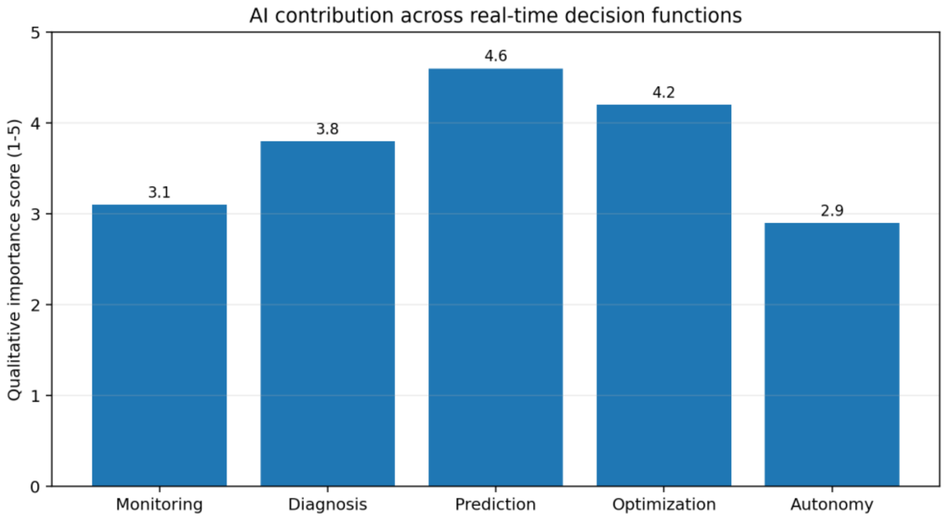

Figure 2. Qualitative synthesis of AI contribution across core real-time decision functions.

6. Application Domains

Smart manufacturing is the most established application area for AI-enabled digital twins. Production systems generate rich sensor data, have measurable performance indicators and benefit directly from predictive maintenance, quality control, scheduling and process optimization. AI-DT systems can forecast tool wear, detect machine anomalies, simulate process changes and support adaptive production. Standards such as ISO 23247 and NIST work on manufacturing DTs are important because they reduce fragmentation and provide reference architectures for implementation [21-23]. The main challenge is that many manufacturing plants contain legacy machines, heterogeneous protocols and proprietary systems that complicate interoperability.

Healthcare digital twins aim to represent patients, organs, treatment pathways, hospital resources or clinical workflows. AI-enabled patient twins may combine physiological models, medical images, electronic health records, wearable data and predictive algorithms to support diagnosis or treatment planning. Healthcare twins are promising because they could personalize care and anticipate deterioration. However, they also face the highest barriers: privacy, bias, clinical validation, regulatory approval, model explainability and ethical responsibility. A wrong recommendation in healthcare can directly harm a patient, so AI-DT systems must support clinicians rather than replace professional judgement [24].

Smart cities use digital twins to integrate urban data from transport, buildings, energy, pollution, weather, utilities and public services. AI can forecast traffic congestion, simulate policy scenarios, optimize energy use, evaluate emergency response and support public participation. Urban DTs are difficult because cities are open systems with social, political and environmental complexity. Data ownership is distributed across agencies and private companies. Citizens may also be affected by algorithmic decisions without knowing how those decisions were made. Therefore, urban digital twins require transparent governance, cybersecurity and participatory design [25].

Energy and smart grid twins support forecasting, fault detection, distributed energy resource management and resilience planning. AI can predict demand, optimize storage, detect equipment degradation and schedule grid tasks. Real-time operation is especially important because renewable energy creates variability and because grid failures can cascade. Edge-cloud task scheduling and reinforcement learning are promising for reducing computational delay and cost in power-grid twins [26]. Yet energy DTs must satisfy strict reliability and security requirements because they are linked to critical infrastructure.

Cybersecurity is both an application and a requirement of digital twins. On the application side, DTs can model network behaviour, simulate attack scenarios, detect anomalies and test defensive strategies. AI can improve threat detection by learning normal and malicious patterns. On the risk side, the digital twin itself expands the attack surface. If adversaries manipulate sensor data, steal the twin model or compromise control actions, the twin may become a pathway to physical harm. Research on AI-based cybersecurity for DTs therefore emphasizes secure architecture, authentication, intrusion detection, explainable threat analysis and privacy-preserving learning [11,12].

Transportation and autonomous systems use digital twins for vehicles, intersections, logistics networks and wireless environments. Real-time DTs can support route planning, vehicle-to-infrastructure communication, autonomous driving validation and traffic signal control. In 6G research, digital twins of wireless environments are proposed to guide communication and sensing decisions through continuously updated 3D maps and machine learning [27]. These systems require fast synchronization because outdated state information can create unsafe or inefficient decisions.

Table 2: Application domains of AI-enabled digital twins.

Domain	Real-time decision	Typical data	Main challenge
Smart manufacturing	Predictive maintenance, scheduling, quality control	Machine sensors, PLC logs, process data	Legacy integration and model validation
Healthcare	Risk prediction, treatment support, resource planning	EHRs, imaging, wearables, labs	Privacy, bias, clinical approval
Smart cities	Traffic, energy, emergency and planning decisions	GIS, IoT, mobility, weather, utilities	Governance and data ownership
Energy systems	Demand response, fault detection, grid balancing	Smart meters, SCADA, DER data	Reliability and critical infrastructure security
Cybersecurity	Threat detection and defence simulation	Network telemetry, logs, attack traces	Adversarial manipulation and privacy
Transport	Routing, signal control, fleet coordination	Vehicle, road, camera and communication data	Latency and safety assurance

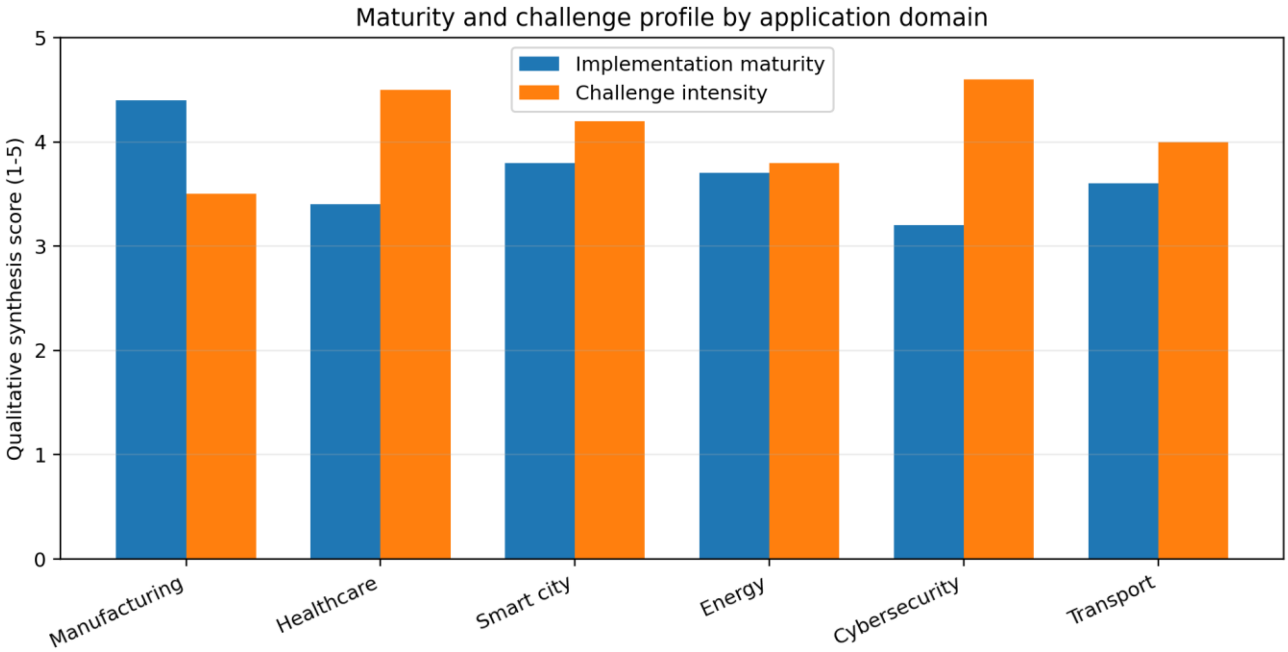

Figure 3. Qualitative comparison of implementation maturity and challenge intensity across application domains.

Figure 3 compares application maturity and challenge intensity across domains. Manufacturing appears most mature because of clearer operational boundaries and measurable outputs. Healthcare and cybersecurity show high challenge intensity because of privacy, safety and adversarial risk. Smart cities and transport require large-scale interoperability and governance. Energy systems require reliability and critical-infrastructure protection. The comparison shows that AI-DT adoption is not limited by algorithms alone; it is limited by socio-technical integration.

7. Benefits, Evaluation and Design Principles

The primary benefit of AI-enabled digital twins is improved situational awareness. A conventional dashboard may display the current state of a system, but an AI-DT can interpret the state, compare it with historical patterns, simulate future trajectories and recommend responses. This capability supports faster and more informed decisions. In manufacturing, early fault detection can reduce downtime. In healthcare, patient-risk prediction can support timely intervention. In smart cities, scenario simulation can help planners evaluate policy options. In cybersecurity, a twin can provide a safe environment for testing attack and defence strategies before they affect production systems.

The second benefit is predictive and prescriptive capability. Prediction estimates what is likely to happen; prescription recommends what should be done. AI-DT systems move from descriptive analytics to decision intelligence when they combine forecasting with optimization. For example, predicting that a machine will fail is useful, but recommending the best maintenance time based on production schedule, spare-parts availability and failure risk is more valuable. This transformation requires the twin to integrate AI models with constraints, objectives and domain rules.

The third benefit is experimentation without physical disruption. Digital twins allow what-if analysis, synthetic data generation and virtual testing. This is important for AI because real-world experimentation may be expensive, dangerous or unethical. A twin can simulate rare faults, emergency events, cyberattacks or alternative operating conditions. However, the value of simulation depends on fidelity. A low-fidelity model may produce misleading results. Therefore, validation against real data is a core requirement for responsible DT-based decision-making.

Evaluation of AI-enabled digital twins should include both AI metrics and system-level metrics. AI metrics include accuracy, precision, recall, F1-score, area under the curve, forecast error, calibration, uncertainty quality and explanation quality. System-level metrics include latency, availability, throughput, interoperability, decision usefulness, energy consumption, cost reduction, downtime reduction, safety improvement and user trust. In real-time settings, a slightly less accurate model may be preferable if it is faster, more robust and easier to explain. Evaluation should therefore be multi-objective rather than based only on model accuracy.

Design principles can guide implementation. First, build the twin around decisions, not around technology. Developers should identify what decision the system supports, who makes it, what time window is available and what evidence is needed. Second, maintain a clear digital thread from data source to decision output so that recommendations can be audited. Third, combine physics-based and data-driven models where possible because hybrid modelling improves generalization and trust. Fourth, design for interoperability using standards, APIs and semantic models. Fifth, implement security and privacy controls from the beginning rather than as later additions. Sixth, keep human responsibility explicit, especially in high-risk decisions.

Table 3: Risks and mitigation strategies.

Risk	Impact	Mitigation
Poor data quality	Incorrect state estimation and weak predictions	Data validation, redundancy, drift detection and uncertainty propagation
Model drift	Declining accuracy after system changes	Continuous monitoring, retraining triggers and version control
Cyberattack	Manipulated data, stolen models or unsafe actuation	Zero-trust design, encryption, intrusion detection and access control
Lack of explainability	Low user trust and weak accountability	Role-specific XAI, audit logs and confidence communication
Interoperability failure	High cost and limited scaling	Standards, APIs, ontologies and modular architecture
Unsafe autonomy	Physical, clinical or social harm	Human-in-the-loop control, safety constraints and validation gates

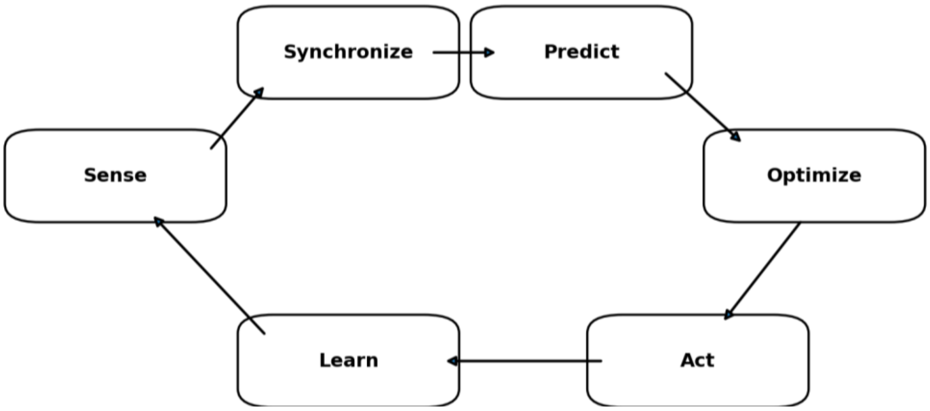

Figure 4: The closed-loop decision cycle. The twin senses the physical state, synchronizes data, predicts future behaviour, optimizes possible actions, acts through human or automated channels and learns from the outcome. The cycle is repeated continuously, but every loop should be governed by validation checks and risk controls.

8. Challenges and Future Research Agenda

Data quality is the first challenge. Digital twins depend on continuous data streams, yet real systems produce missing, noisy, delayed, biased and inconsistent data. Sensor drift may cause the twin to represent a false state. Data from different systems may use incompatible formats or semantics. Future work should develop robust data-quality monitoring, automated data repair, uncertainty propagation and semantic validation methods for DT pipelines.

Interoperability is the second challenge. Digital twins often connect engineering tools, IoT platforms, enterprise software, AI models and visualization systems. Without common standards, implementation becomes expensive and difficult to scale. ISO 23247, NIST initiatives and industrial reference architectures are important steps, but more work is needed for cross-domain interoperability, model exchange, semantic mapping and lifecycle governance [21-23].

Trustworthiness is the third challenge. AI-DT recommendations may affect safety, privacy, finance and public services. Users must understand how decisions are generated and when the system should not be trusted. Future research should integrate explainable AI, uncertainty quantification, model validation and audit trails into DT architectures. Explanations should be designed for specific user roles rather than presented as generic technical outputs.

Cybersecurity is the fourth challenge. Digital twins collect sensitive data, expose APIs and may connect to control systems. Attackers can target sensors, communication channels, models, dashboards or feedback commands. The twin can also be used by defenders to simulate threats and improve detection. Future work should treat security as a co-equal design layer, including zero-trust access, secure edge deployment, adversarial robustness, encrypted data exchange and continuous threat monitoring [11,12].

Latency and scalability are the fifth challenge. Real-time decision-making requires computing resources to be placed and orchestrated intelligently. Edge computing reduces delay but has limited resources. Cloud computing provides scale but increases latency and dependency on connectivity. Future DT architectures should support adaptive partitioning, where inference, simulation and storage move between edge, fog and cloud based on decision urgency, privacy and computational cost.

Human accountability is the final challenge. AI-enabled digital twins may create an illusion of objectivity because they combine advanced models, live data and visual interfaces. However, all models are abstractions and all decisions involve values. Future research should study human-DT collaboration, interface design, operator trust, cognitive load, responsibility allocation and ethical governance. The goal is not to remove humans from decisions but to create digital twins that make human decisions faster, safer and better informed.

Table 4: Future research agenda

Research direction	Expected contribution
Hybrid physics-AI modelling	Improves prediction, generalization and interpretability
Semantic digital twin platforms	Enables interoperable and machine-readable system knowledge
Privacy-preserving DT learning	Supports multi-site intelligence without raw data centralization
Real-time edge-cloud orchestration	Balances latency, cost, privacy and computational load
Trustworthy and explainable AI	Improves adoption in safety-critical and regulated settings
DT cybersecurity testbeds	Supports realistic evaluation of cyber-physical threats
Human-centered interfaces	Improves operator understanding, trust and decision quality

8.1 Integrative Research Note

A careful interpretation of these findings is necessary because digital twins are not a single product category but a system-of-systems pattern. Their performance depends on data governance, synchronization, domain modelling, model lifecycle management, interface quality and organizational readiness. For this reason, future evaluations should report not only algorithmic results but also deployment assumptions, latency budgets, update frequency, human roles, cybersecurity controls and evidence of real operational benefit. A careful interpretation of these findings is necessary because digital twins are not a single product category but a system-of-systems pattern. Their performance depends on data governance, synchronization, domain modelling, model lifecycle management, interface quality and organizational readiness. For this reason, future evaluations should report not only algorithmic results but also deployment assumptions, latency budgets, update frequency, human roles, cybersecurity controls and evidence of real operational benefit. A careful interpretation of these findings is necessary because digital twins are not a single product category but a system-of-systems pattern. Their performance depends on data governance, synchronization, domain modelling, model lifecycle management, interface quality and organizational readiness. For this reason, future evaluations should report not only algorithmic results but also deployment.

9. Conclusion

AI-enabled digital twins represent a major direction in computer science because they connect live data, modelling, machine learning, simulation, optimization and human decision-making. They are increasingly used to monitor complex systems, forecast future states, test scenarios and recommend actions. This review has shown that the real value of AI-DT systems lies not in visualization alone but in smart computing architectures that support trustworthy real-time decisions. The proposed layered architecture highlights the importance of sensing, connectivity, data fabric, modelling, AI analytics, decision services and feedback learning. Across application domains, manufacturing is currently the most mature, while healthcare, smart cities, energy, cybersecurity and transport reveal important opportunities and challenges.

The field must move beyond isolated prototypes. Future AI-DT systems need standards, interoperability, secure architectures, explainable models, uncertainty-aware predictions, privacy-preserving learning and human-centered governance. Digital twins should be designed around decisions, validated continuously and evaluated through both AI metrics and operational outcomes. When implemented responsibly, AI-enabled digital twins can become powerful platforms for real-time decision-making in smart computing environments. When implemented carelessly, they can amplify errors, bias and cyber risk. The next stage of research should therefore focus on reliable, secure and explainable digital twins that combine the strengths of physical modelling, artificial intelligence and responsible system design.

References

[1] Jones, D., Snider, C., Nassehi, A., Yon, J. and Hicks, B. (2020). Characterising the digital twin: A systematic literature review. CIRP Journal of Manufacturing Science and Technology, 29, 36-52. https://doi.org/10.1016/j.cirpj.2020.02.002

[2] Tao, F., Zhang, H., Liu, A. and Nee, A.Y.C. (2019). Digital twin in industry: State-of-the-art. IEEE Transactions on Industrial Informatics, 15(4), 2405-2415. https://doi.org/10.1109/TII.2018.2873186

[3] Kritzinger, W., Karner, M., Traar, G., Henjes, J. and Sihn, W. (2018). Digital twin in manufacturing: A categorical literature review and classification. IFAC-PapersOnLine, 51(11), 1016-1022. https://doi.org/10.1016/j.ifacol.2018.08.474

[4] Fuller, A., Fan, Z., Day, C. and Barlow, C. (2020). Digital twin: Enabling technologies, challenges and open research. IEEE Access, 8, 108952-108971. https://doi.org/10.1109/ACCESS.2020.2998358

[5] VanDerHorn, E. and Mahadevan, S. (2021). Digital twin: Generalization, characterization and implementation. Decision Support Systems, 145, 113524. https://doi.org/10.1016/j.dss.2021.113524

[6] Kreuzer, T., Papapetrou, P. and Zdravkovic, J. (2024). Artificial intelligence in digital twins: A systematic literature review. Data & Knowledge Engineering, 151, 102304. https://doi.org/10.1016/j.datak.2024.102304

[7] Zhou, R. et al. (2026). Digital twin AI: Opportunities and challenges from large language models to world models. arXiv preprint arXiv:2601.01321.

[8] Alkhateeb, A., Jiang, S. and Charan, G. (2023). Real-time digital twins: Vision and research directions for 6G and beyond. arXiv preprint arXiv:2301.11283.

[9] Thelen, A., Zhang, X., Fink, O., Lu, Y., Ghosh, S., Youn, B.D., Todd, M.D., Mahadevan, S., Hu, C. and Hu, Z. (2022). A comprehensive review of digital twin: Part 1: Modeling and twinning enabling technologies. arXiv preprint arXiv:2208.14197.

[10] Sharma, A., Kosasih, E., Zhang, J., Brintrup, A. and Calinescu, A. (2020). Digital twins: State of the art theory and practice, challenges, and open research questions. Journal of Industrial Information Integration, 22, 100189.

[11] Homaei, M.H., Mogollon-Gutierrez, O., Sancho-Nunez, J.C., Vegas, M.A. and Lindo, A.C. (2024). A review of digital twins and their application in cybersecurity based on artificial intelligence. Artificial Intelligence Review, 57, 201.

[12] Khan, M.A., Salah, K., Jayaraman, R., Arshad, J., Omar, M. and Ellahham, S. (2024). Explainable AI for cybersecurity automation, intelligence and trustworthiness in digital twin: Methods, taxonomy, challenges and prospects. ICT Express, 10(5), 1010-1024.

[13] Minerva, R., Lee, G.M. and Crespi, N. (2020). Digital twin in the IoT context: A survey on technical features, scenarios, and architectural models. Proceedings of the IEEE, 108(10), 1785-1824.

[14] Qi, Q. and Tao, F. (2018). Digital twin and big data towards smart manufacturing and Industry 4.0: 360 degree comparison. IEEE Access, 6, 3585-3593.

[15] Chen, Y., Zhang, Y., Yao, X. and Hu, J. (2024). Towards the application of machine learning in digital twin technology: A multi-scale review. Discover Applied Sciences, 6, 589.

[16] Liu, X. and David, I. (2025). AI simulation by digital twins: Systematic survey, reference framework, and mapping to a standardized architecture. arXiv preprint arXiv:2506.06580.

[17] Listl, F.G., Dittler, D., Hildebrandt, G., Stegmaier, V., Jazdi, N. and Weyrich, M. (2024). Knowledge graphs in the digital twin: A systematic literature review about the combination of semantic technologies and simulation in industrial automation. arXiv preprint arXiv:2406.09042.

[18] Karabulut, E., Pileggi, S.F., Groth, P. and Degeler, V. (2023). Ontologies in digital twins: A systematic literature review. arXiv preprint arXiv:2308.15168.

[19] Yang, W., Zhang, Y., Wang, H. and Zhang, J. (2024). Adaptive optimization federated learning enabled digital twins in industrial IoT. Journal of Industrial Information Integration, 41, 100645. https://doi.org/10.1016/j.jii.2024.100645

[20] Sun, W., Lei, S., Wang, L., Liu, Z. and Zhang, Y. (2020). Adaptive federated learning and digital twin for Industrial Internet of Things. arXiv preprint arXiv:2010.13058.

[21] ISO. (2021). ISO 23247-1:2021: Automation systems and integration - Digital twin framework for manufacturing - Part 1: Overview and general principles. International Organization for Standardization.

[22] ISO. (2021). ISO 23247-2:2021: Automation systems and integration - Digital twin framework for manufacturing - Part 2: Reference architecture. International Organization for Standardization.

[23] Shao, G. and Helu, M. (2024). Manufacturing digital twin standards. National Institute of Standards and Technology, EDTconf 2024.

[24] Corral-Acero, J. et al. (2020). The digital twin to enable the vision of precision cardiology. European Heart Journal, 41(48), 4556-4564.

[25] Deng, T., Zhang, K. and Shen, Z.J.M. (2023). A systematic review of a digital twin city: A new pattern of urban governance toward smart cities. Expert Systems with Applications, 214, 119531.

[26] Wang, Y., Chen, Y., Yang, H., Wang, L. and Zhang, D. (2024). Real-time scheduling of power grid digital twin tasks in cloud via deep reinforcement learning. Journal of Cloud Computing, 13, 106.

[27] Yigit, Y., Panitsas, I., Maglaras, L., Tassiulas, L. and Canberk, B. (2024). Cyber-Twin: Digital twin-boosted autonomous attack detection for vehicular ad-hoc networks. arXiv preprint arXiv:2401.14005.

[28] Adreani, L., Bellini, P., Fanfani, M., Nesi, P. and Pantaleo, G. (2023). Smart city digital twin framework for real-time multi-data integration and wide public distribution. arXiv preprint arXiv:2309.13394.

[29] Zhang, Z., Zhang, X., Zhu, G., Wang, Y. and Hui, P. (2023). Efficient task offloading algorithm for digital twin in edge/cloud computing environment. arXiv preprint arXiv:2307.05888.

[30] NIST. (2025). Digital twin standardization. National Institute of Standards and Technology.