Article

Decision-Making in Industry 4.0: A Layered Framework for Converting Cyber-Physical Information into Governed Action

Suzhen Huang [1,*]

[1] Independent Researcher, China
* Correspondence: Suzhen Huang, Independent Researcher, China

Abstract: Industry 4.0 has significantly expanded the amount, speed, and granularity of information available to modern industrial organizations, but richer information does not automatically produce better or more effective decisions. The fundamental practical problem is to systematically convert cyber-physical observations into accountable actions under conditions of uncertainty, strict operational constraints, and necessary human governance. To address this gap, this paper develops a comprehensive conceptual framework for Industry 4.0 decision-making that seamlessly connects sensing, state representation, prediction, optimization, authorization, execution, and continuous learning. The proposed framework is organized around three core claims. First, industrial decision-making should be treated as a dynamic, closed-loop process rather than a static, one-way analytics pipeline. Second, decision models must be carefully aligned with the specific time horizon, observability, uncertainty, and reversibility of the action being selected. Third, reliability-oriented applications—such as inspection, maintenance, inventory control, mission abort, and system reconfiguration—serve as highly useful integration cases because they clearly expose the critical connection between physical degradation, cyber representation, prescriptive optimization, and governed execution. Ultimately, the paper contributes a novel layered architecture, a horizon-based model alignment table, and a robust closed-loop decision cycle that can be utilized to organize future academic research and practical implementation. The resulting perspective shifts critical attention away from isolated technological solutions and toward the holistic quality of the complete decision loop, ensuring the system observes relevant states, represents uncertainty properly, recommends feasible actions, assigns authority clearly, and learns effectively from realized outcomes.

Keywords: industry 4.0; decision-making; cyber-physical systems; prescriptive analytics; digital twin

Received: 05 April 2026
Revised: 15 May 2026
Accepted: 25 May 2026
Published: 01 June 2026

1. Introduction

Industry 4.0 is commonly associated with cyber-physical systems, connected machines, industrial data platforms, cloud-edge infrastructure, and autonomous resources. These technologies have changed the information base of industrial management by making production, maintenance, logistics, and service operations more observable and more responsive. The core transformation, however, is not simply technological [1]. It is also a change in how industrial organizations formulate and execute decisions. Early descriptions of Industry 4.0 emphasized vertical and horizontal integration, embedded intelligence, and the coupling of physical and digital processes. This technological foundation implies that decision-making can no longer be treated as a delayed managerial activity that occurs after production data have been collected. Instead, decision-making becomes part of the operating system itself.

The distinction matters because many implementations still stop at visibility. Sensors, dashboards, and data lakes can reveal asset condition, process variation, or demand shifts, but they do not determine which action should be taken, when it should be taken, who should authorize it, or how the action should be evaluated afterward [2]. The design principles of interconnection, information transparency, technical assistance, and decentralized decisions make this point clear: interconnection and transparency are necessary, but not sufficient, for accountable decisions. A connected factory may therefore remain decision-poor if it lacks state models, optimization logic, execution protocols, and governance rules.

This paper addresses that gap by treating Industry 4.0 decision-making as a structured conversion process. The process begins with physical events and ends with governed actions, but it requires intermediate layers that translate raw observations into decision-relevant state variables, risk estimates, feasible policies, and audit-ready outcomes. Cyber-physical manufacturing architectures provide the technical basis for this conversion because they link physical assets with computation and communication resources [3]. At the same time, the broader state of the art shows that Industry 4.0 involves not only automation but also systems integration, lifecycle data, supply-chain coordination, and organizational change.

The paper develops a conceptual framework rather than a new empirical model. Its purpose is to organize existing decision-making problems and to clarify where different analytical tools fit [4]. Three questions guide the argument. What is the industrial decision object in Industry 4.0? Which layers are required to transform cyber-physical data into action? How should model families be matched to decision horizons such as strategic investment, tactical maintenance planning, operational scheduling, and real-time control? A data-driven manufacturing perspective provides the background for these questions because industrial value is created when data become actionable intelligence rather than when data are merely stored or visualized.

Figure 1 summarizes the logic of the manuscript. The argument begins with the industrial problem, moves to a layered decision architecture, aligns models with decision horizons, applies the framework to reliability-oriented cases, and closes with governance requirements. This sequence is deliberate [4]. A paper that begins with algorithms may overlook whether the decision problem is properly framed. A paper that begins with governance may lack a clear account of the technical loop that governance must control. The framework therefore proceeds from the industrial decision problem to the full loop that converts predictions into actions.

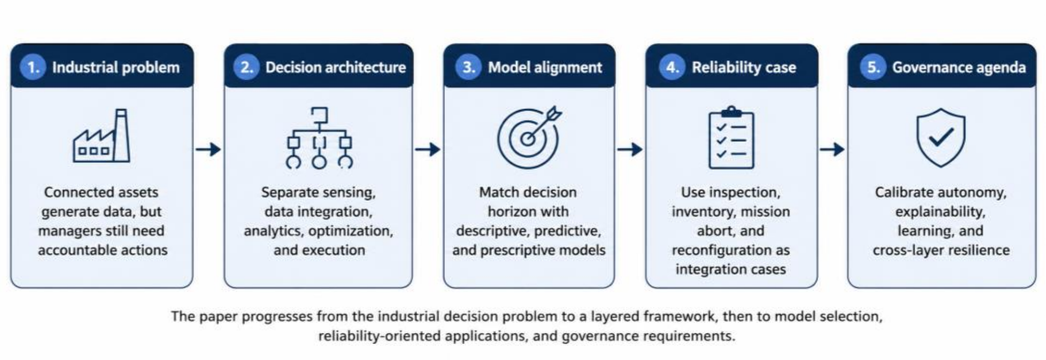

The paper progresses from the industrial decision problem to a layered framework, then to model selection, reliability-oriented applications, and governance requirements.

Figure 1. Sequential Logic of the Manuscript and the Industry 4.0 Decision-Making Problem.

2. From Industrial Visibility to Decision Requirements

The first step is to distinguish visibility from decision quality. Visibility refers to the ability to observe or infer what is happening in an industrial system. Decision quality refers to whether the organization selects an action that is feasible, timely, defensible, and aligned with objectives. Industry 4.0 improves visibility by expanding the sensor network and by connecting asset-level data with enterprise information. Decision quality, however, requires a model of action. The organization must know what can be changed, which constraints matter, how uncertainty affects consequences, and how responsibility is assigned [5].

A decision can be represented as the selection of an action from a feasible action set under a state description and an evaluation criterion. In industrial settings, the state may include equipment degradation, workload, spare-part inventory, production priorities, operator availability, cyber risk, energy price, or customer demand. The feasible actions may include inspection, maintenance, repair, replacement, rescheduling, load redistribution, mission continuation, abort, or reconfiguration. The evaluation criterion may be cost, availability, throughput, safety, carbon performance, service reliability, or a weighted combination. This formulation explains why a dashboard is not a decision system. A dashboard may display an alarm, but it rarely defines the state transition model, action space, reward or loss function, and governance boundary needed for decision-making [6].

A second requirement is state representation [7]. Industrial systems are rarely fully observable. A machine can be degraded even when no alarm is active; a protection component can be unavailable until a demand event occurs; a demand stream can change before planners recognize the shift. The representation layer therefore converts signals into beliefs, risk scores, digital-twin states, or sufficient statistics. In performance-based reliability contexts, representation may also include the configuration of a multi-component system and the possibility that the system can dynamically reconfigure itself to preserve performance. Such cases show that the state is not only a sensor reading. It is an operational interpretation of what the system can still do and what action remains possible.

A third requirement is the separation between prediction and prescription. Predictive analytics estimates what is likely to happen. Examples include remaining useful life, demand, defect probability, energy use, production bottlenecks, and service loss [8]. Prescriptive decision-making evaluates what should be done under objectives and constraints. In smart manufacturing, the difference is central because the value of prediction is realized only when it changes an action. A highly accurate failure prediction is useful only if it improves inspection timing, maintenance scheduling, spare provisioning, mission planning, or risk acceptance.

A fourth requirement is execution governance [9]. The output of a model may be a recommendation, but the industrial decision is not complete until the action is authorized, executed, monitored, and recorded. This is particularly important when decisions are decentralized or semi-autonomous. Industry 4.0 technologies can enhance the decision process, yet the enhancement depends on whether organizations redesign the flow from information to action rather than simply adding more analytics. Governance is therefore not a separate administrative layer after the decision. It is part of the decision design because it determines which recommendations can become actions and under what conditions.

These requirements create a practical definition: Industry 4.0 decision-making is the governed conversion of cyber-physical information into feasible, timed, and accountable action [2]. The word governed is essential. It captures the need for human authority, rule-based constraints, explainability, auditability, and post-action learning. Without governance, an optimizer may produce an efficient action that is unacceptable in practice. Without optimization, governance may produce cautious but inefficient rules. The architecture must therefore connect both sides.

3. A Layered Architecture for Industry 4.0 Decision-Making

The second logical step is to specify the architecture that supports the conversion from observation to action [10]. Figure 2 presents the proposed architecture. It is intentionally layered because industrial decisions require information to move across different levels of abstraction. A vibration signal, for example, is not a maintenance policy. It must be cleaned, contextualized, mapped into a health state or belief state, combined with operational constraints, processed through a decision model, authorized, and then converted into a work order or control instruction.

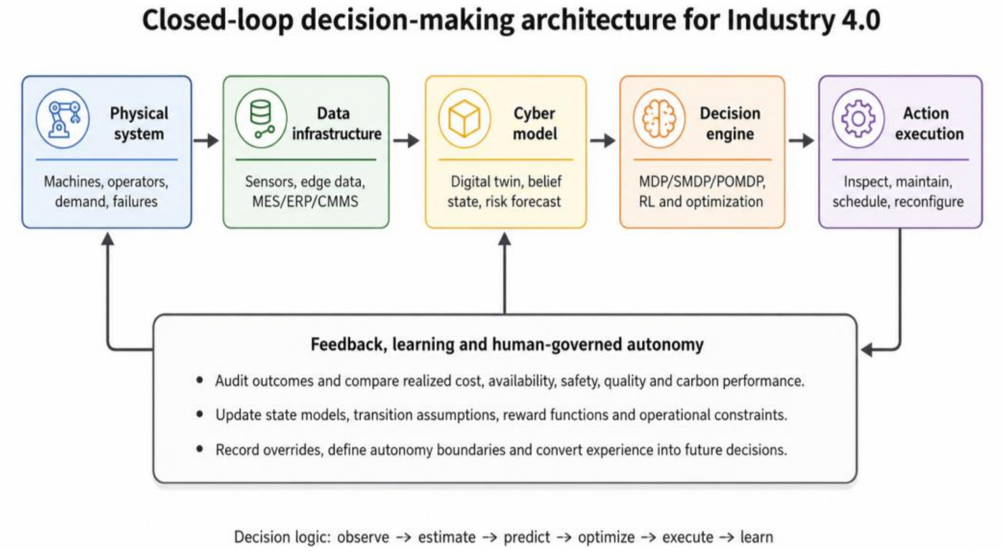

Figure 2. Layered Architecture for Transforming Cyber-Physical Observations into Governed Industrial Actions.

The physical and human layer contains machines, products, operators, service users, inventories, work orders, energy inputs, failures, and environmental conditions. This layer is the source of operational value and operational risk. It also places limits on digital decision-making. A recommendation that ignores downtime windows, operator workload, tool availability, or safety restrictions may be analytically optimal but infeasible in practice. The architecture therefore begins with the physical system rather than with the data platform [11].

The data infrastructure layer converts physical events into usable digital information. It includes sensors, controllers, edge devices, industrial networks, manufacturing execution systems, enterprise systems, maintenance management systems, and cloud services. Its decision role is to provide timely, reliable, and contextual data. The issue is not only whether data are collected, but whether the data are synchronized, traceable, sufficiently granular, and linked to the decision context. In reliability applications, for example, asset condition data must be combined with usage intensity, operating environment, inventory status, and planned production to support maintenance decisions [12].

The cyber model layer translates data into states [13]. This layer may contain a digital twin, a degradation model, a belief-state estimator, a simulation model, or a risk representation. Digital twins and big data approaches are useful because they can connect lifecycle information with current operating states. They are most valuable when the representation is decision-relevant: the model should describe not merely what the asset looks like, but what actions remain feasible and how future outcomes may change under different actions. Recent digital-twin applications show the breadth of possible use cases, but the decision architecture requires that the twin be coupled to policy evaluation rather than used only for visualization.

The decision engine layer converts states and predictions into actions. Depending on the problem, it may use mathematical programming, simulation-optimization, stochastic dynamic programming, Markov decision processes, semi-Markov decision processes, partially observable models, reinforcement learning, multi-criteria decision analysis, or rule-based logic [10]. In multi-component systems, this layer must often combine component-level information with system-level consequences. Inspection and maintenance policies with protection components illustrate this requirement because the decision must account for hidden degradation, protection failure, interaction among components, and the economic consequence of inspecting or maintaining at different levels.

The action execution layer converts selected actions into operational changes. Execution may involve issuing a maintenance work order, changing a production schedule, allocating spare parts, adjusting load, authorizing a mission abort, or reconfiguring the system. The execution layer must include confirmation that the action was completed and that its outcome was observed [6]. This feedback is required for model updating and for accountability. A decision architecture without feedback can recommend actions, but it cannot learn whether the recommendations worked.

The architecture also shows why a purely technical view of Industry 4.0 is incomplete. Data infrastructure, cyber models, decision engines, and action execution are interdependent. Better prediction may not improve performance if execution is slow [6]. More automation may increase risk if authorization rules are unclear. A digital twin may be impressive but weak if it is not connected to optimization. The next step is therefore to align models with decision horizons and decision properties.

4. Decision Horizons and Model Alignment

Industry 4.0 decisions differ in time scale, reversibility, observability, and risk. Strategic decisions concern technology adoption, platform architecture, supplier relationships, sustainability priorities, and organizational capability. Tactical decisions translate strategy into maintenance plans, inventory policies, workforce allocation, and capacity choices. Operational decisions assign work, schedule inspections, allocate resources, and coordinate production. Real-time decisions control machines, reroute tasks, adjust loads, or trigger emergency actions. Treating all of these decisions as one generic analytics problem leads to weak model selection.

Table 1 clarifies that model alignment should begin with the decision horizon. Strategic decisions tolerate slower computation and may use qualitative information because their purpose is to compare long-term alternatives. Tactical decisions require stronger stochastic modeling because they convert uncertainty into policy. Operational decisions require timeliness and integration with execution systems [6]. Real-time decisions require strict safety and latency constraints. Learning-loop decisions require evidence about whether the model or policy should be changed.

Table 1. Decision Horizons, Suitable Model Families, and Governance Emphasis in Industry 4.0

Decision horizon	**Typical question**	**Suitable model family**	**Governance emphasis**
Strategic	Which Industry 4.0 capabilities should be built or purchased?	Scenario analysis, portfolio models, multi-criteria decision-making	Investment accountability, sustainability, interoperability
Tactical	How should maintenance, inventory, capacity, and service policies be configured?	Stochastic programming, simulation-optimization, MDP/SMDP, reliability models	Cost-risk tradeoff, resource feasibility, policy robustness
Operational	Which job, inspection, repair, or spare allocation should be executed now?	Scheduling, dispatching, rolling-horizon optimization, predictive analytics	Human approval, explainability, traceability

Real-time	Should the system continue, stop, reroute, or reconfigure?	Control, reinforcement learning, threshold policies, anomaly response	Safety envelope, override rights, cyber integrity
Learning loop	How should models and policies be updated after outcomes?	Bayesian updating, drift detection, causal evaluation, policy learning	Audit trail, model lifecycle, change management

The same predictive model can therefore be useful or irrelevant depending on the decision horizon. A remaining-useful-life estimate may inform long-term spare-part planning, weekly maintenance scheduling, same-day dispatching, or real-time shutdown. Each use requires a different action space and a different tolerance for error [3]. Predictive-maintenance research has identified a wide range of models and challenges, but the decision value of such models depends on whether predictions are linked to decisions with explicit constraints and consequences.

Observability is another model-selection criterion. If the true system state is directly observed, deterministic optimization or standard stochastic models may be sufficient. If the system is partially observed, the model must represent uncertainty through beliefs, hidden states, or information structures. Belief-based inspection and maintenance models show how partial observability changes decision timing and action selection because the policy is based on what is known and what remains uncertain, not only on the physical state itself [3].

Coupling is a third criterion [11]. Many industrial decisions are coupled across assets, resources, and time. A maintenance decision changes asset availability; an inventory decision changes the cost of future downtime; a mission abort decision changes both safety risk and revenue; a reconfiguration decision changes performance and degradation exposure. Joint sizing, maintenance, and inventory policies illustrate why a single-layer decision is often inadequate when load-sharing, self-announcing failures, and spare availability interact.

Digital twins can support model alignment if they are treated as decision interfaces. A twin that only reproduces the physical system is descriptive. A twin that estimates future trajectories is predictive. A twin that supports policy comparison is prescriptive [9]. The distinction between cyber-physical systems and digital twins is therefore not semantic; it affects whether the model is used for monitoring, simulation, or decision optimization. Model alignment should make that role explicit before implementation begins.

5. Reliability-Oriented Decision-Making as an Integrative Application

Reliability-oriented decision-making is a useful application for the framework because it forces the full loop to be specified. Reliability decisions are not only about whether an asset has failed. They concern when to inspect, whether to maintain, how many spares to hold, whether to continue a mission, whether to reconfigure the system, and how much autonomy should be allowed. These decisions involve hidden states, uncertain degradation, time-dependent costs, operational constraints, and human accountability. They therefore expose the weakness of treating Industry 4.0 as a collection of sensors or algorithms.

Reliability-oriented decisions also connect technical availability with productivity, safety, service reliability, and sustainability performance, which is why they fit the broader Industry 4.0 value proposition.

Figure 3 represents the closed-loop decision cycle. The cycle begins with sensing and ends with learning, but the central point is that prediction becomes valuable only when it is converted into governed action. This is why the loop includes authorization between optimization and execution. Without authorization, the system may act faster than the organization can justify. Without learning, the same errors can be repeated under the appearance of automation [4] (As shown in Table 2).

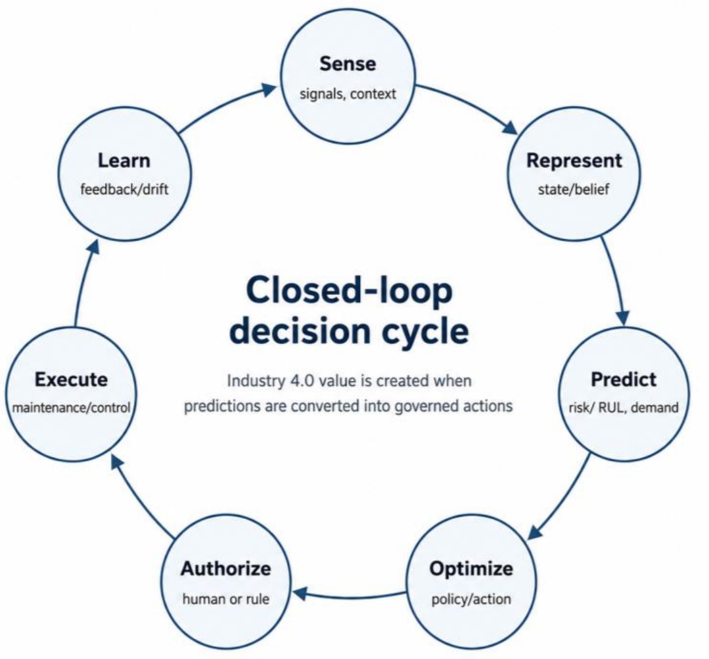

Figure 3. Closed-loop Decision Cycle for Industry 4.0 Reliability-Oriented Applications.

Table 2. Elements of the Closed-Loop Decision Cycle and Common Failure Modes

Loop element	Decision role	Common failure mode	Practical control
Sense	Collect signals and context	Incomplete or biased data	Calibration, redundancy, metadata
Represent	Convert observations into states or beliefs	State mismatch or hidden degradation	Belief updating, digital-twin validation
Predict	Estimate risk, demand, RUL, or performance	Point forecasts used as certainty	Uncertainty quantification
Optimize	Select feasible action under objectives	Objective ignores constraints or side effects	Constraint checks, sensitivity analysis
Authorize	Assign human or rule-based approval	Ambiguous responsibility	Approval thresholds, audit trail
Execute	Implement action and record outcome	Recommendation not operationally feasible	Workflow integration, completion confirmation
Learn	Update models and policies	Model drift or unobserved consequences	Lifecycle monitoring, post-action review

Inspection and maintenance are natural examples. Sensing provides condition indicators, usage histories, environmental exposure, and failure reports. Representation converts these observations into health states, beliefs, or degradation parameters. Prediction estimates failure risk, remaining useful life, and downtime consequences.

Optimization selects inspection and maintenance actions. Authorization determines whether the recommendation is approved automatically, by a rule, or by a human manager. Execution creates the work order and updates the operational record. Learning compares the realized outcome with the predicted outcome.

The same loop extends to production and service systems with multiple demands. When an asset serves several task types, the maintenance decision must consider not only asset degradation but also demand composition, service delay, and the possibility of competing sudden and deterioration-induced failures. Such settings support the broader point that Industry 4.0 decision models should integrate operating context rather than optimize a technical indicator in isolation.

Multi-component systems make the loop more demanding. A component-level failure may affect system performance only under certain configurations; a protection component may fail silently; an inspection can reveal some states but not all states; a maintenance action may restore one component while delaying production. Multi-level inspection and maintenance policies illustrate how system-level and component-level decisions must be coordinated rather than handled as independent tasks [3]. This is where Industry 4.0 architecture is valuable: connected data can support richer state representation, but the decision engine must still determine which information is worth acquiring and which action is economically justified.

Partial observability is especially important in reliability decisions. If decision makers act as if the state is fully known, they may maintain too late, inspect too often, or accept unacceptable risk. Belief-based models address this issue by allowing the decision to depend on both observed evidence and uncertainty about the hidden state [1]. In Industry 4.0 environments, this reasoning becomes more relevant rather than less relevant. More sensors do not eliminate uncertainty; they often reveal that uncertainty is multi-source, time-varying, and context-dependent.

Inventory and maintenance decisions also show why prescriptions must be coupled across resources. A spare-part decision changes the value of preventive maintenance; a maintenance action changes future spare demand; a load-sharing configuration changes component deterioration; self-announcing failures change the timing of decision epochs. Joint policies are therefore more consistent with Industry 4.0 than separated models that optimize inventory first and maintenance afterward.

Reinforcement learning and multi-agent methods add another layer of complexity. They can be useful when systems are too large or dynamic for classical optimization, but they also require clear reward definitions, safety constraints, and policy validation. Maintenance policies designed through multi-agent reinforcement learning show the potential of learning-based prescriptive models in load-sharing systems. The practical question is not whether learning is powerful, but whether its learned policy can be explained, bounded, audited, and improved after deployment.

Reliability decisions can also be mission-oriented. In mission-critical systems, the relevant action may not be maintenance alone but whether to continue, abort, restart, or shift to a safer operating mode. These decisions are policy decisions under risk, not merely predictions of failure. The inclusion of mission abort in the same framework reinforces the importance of authorization and governance because the action may protect the system while sacrificing immediate revenue or task completion.

6. Governance, Autonomy, and Human-in-the-Loop Control

The framework treats governance as part of the decision loop rather than as an external compliance activity. This is necessary because Industry 4.0 systems operate across organizational boundaries, technical layers, and time horizons. A maintenance recommendation may affect production output, spare consumption, service commitments, safety exposure, and sustainability metrics. Industry 4.0 is therefore connected to sustainability and organizational performance, but these outcomes depend on how decisions are designed and governed.

The first governance issue is autonomy calibration [11]. Not every decision should be automated to the same degree. Low-risk, frequent, reversible decisions may be suitable for automatic execution. High-risk, rare, irreversible, or legally sensitive decisions require stronger human oversight. The autonomy boundary should be defined by risk, confidence, time pressure, reversibility, and evidence quality. In practical terms, the system may automatically schedule a routine inspection, recommend a maintenance window for human approval, or require manual authorization before aborting a mission. The governance design should specify the boundary before the model is deployed.

The second issue is explainability. Explanations do not need to expose every mathematical detail, but they must connect the recommendation to the decision variables that matter to managers and operators. For maintenance, an explanation may identify the inferred health state, predicted risk, spare availability, downtime window, and expected cost difference between alternatives. For reinforcement learning, an explanation may describe the policy conditions under which one action is preferred over another [12]. Without such explanation, human oversight becomes symbolic rather than meaningful.

The third issue is auditability. Each industrial decision should leave a record of the state representation, prediction, action recommendation, approval path, execution result, and post-action outcome. Auditability is essential for accountability and for learning. It also helps diagnose whether poor performance came from data error, state representation, model structure, optimization objective, human override, or execution failure. This is especially important when models update over time or when multiple agents jointly influence the policy.

The fourth issue is risk-sensitive authorization. Some actions are costly but safe; others are profitable but risky; some actions protect the asset while sacrificing mission success [5, 11]. Mission-oriented maintenance and abort decisions show that the action boundary must be linked to explicit risk acceptance criteria rather than left to an opaque score. In this setting, the authorization layer is not bureaucratic. It is the mechanism that converts a probabilistic forecast into an accountable operational choice.

The fifth issue is lifecycle governance. Industrial systems drift. Demand changes, degradation patterns shift, operators adapt to recommendations, suppliers change lead times, and cyber threats evolve. A model that was valid at deployment can become unreliable after the operating environment changes. Digital twins and smart manufacturing applications can help maintain an updated representation of the system, but only if update rules and validation routines are embedded in the decision loop.

The sixth issue is data governance and cyber integrity. Because Industry 4.0 decisions depend on connected data, the decision system must protect data availability, confidentiality, and integrity [8]. A corrupted sensor stream can cause a wrong health belief; a manipulated work-order record can distort learning; a delayed communication channel can turn a good recommendation into a late action. Data analytics in Industry 4.0 therefore requires not only modeling capability but also data lineage, access control, validation, and monitoring.

7. Research Agenda

The first research direction is cross-layer optimization. Many models optimize within a single layer: sensing, prediction, scheduling, inventory, or maintenance. Real industrial value often depends on interactions across layers. For example, an inspection policy changes state information; state information changes maintenance timing; maintenance changes capacity; capacity changes service performance. Future models should make these dependencies explicit instead of assuming that a locally optimal layer produces a globally coherent decision.

The second direction is uncertainty-aware prescription. Predictive models often report point estimates, but prescriptive decisions require distributions, belief states, ambiguity sets, or robust decision rules. Uncertainty matters because a recommended action may be optimal under the mean forecast but unacceptable under tail risk. Partially observable and bi-level inspection-maintenance models point toward richer treatments of uncertainty and information value, especially when component-level observations are costly or delayed.

The third direction is digital-twin-to-policy integration. The literature on Industry 4.0 has documented many technologies and applications, but open research issues remain around integration, standards, interoperability, and decision usefulness. A digital twin should not be evaluated only by visual fidelity or data richness. It should also be evaluated by whether it improves decisions. This requires experiments and field studies that compare policies with and without twin-enabled state representation.

The fourth direction is human-centered validation. A decision policy may perform well in simulation but fail in practice if operators distrust it, if explanations are too technical, if recommendations do not fit work routines, or if responsibility is unclear. Future studies should therefore evaluate not only model accuracy and objective value but also acceptance, override behavior, workload, and organizational learning. Human-in-the-loop decision-making should be treated as a design problem rather than an afterthought.

The fifth direction is resilience-oriented decision-making. Industry 4.0 systems face cyber disruption, supply uncertainty, energy volatility, climate-related hazards, and demand shocks. Conventional efficiency objectives may not be sufficient when disruptions are recurrent and dependent. A resilience-oriented decision model should evaluate the ability to detect disruption, maintain critical functions, reconfigure resources, recover after action, and learn from the disruption. This extension is consistent with the closed-loop view because resilience depends on sensing, representation, prediction, optimization, execution, and learning together [6].

The sixth direction is publication-quality conceptualization [10]. Industry 4.0 papers can become weak when they list technologies without specifying decision variables, state representations, action spaces, objectives, constraints, and governance rules. A stronger paper should show how a technology changes a decision and how that decision changes an industrial outcome. This principle also helps prevent superficial citation practices: references should justify claims, define methods, or locate the decision problem, not replace the argument.

8. Conclusion

Industry 4.0 transforms industrial decision-making by integrating physical systems, data infrastructure, cyber models, analytics, optimization, and human governance. The primary challenge is not simply to gather more data or create more precise predictions but to design a decision loop that translates cyber-physical information into actionable, timely, and accountable outcomes. This paper has introduced a layered framework for this transformation and demonstrated how decision horizons, observability, coupling, and governance requirements influence model selection.

The framework offers three practical insights. First, visibility and action should be distinctly analyzed, as observing an industrial state differs fundamentally from selecting an action under constraints. Second, model families must align with the decision horizon and the structure of uncertainty, as strategic, tactical, operational, real-time, and learning-loop decisions demand distinct methodologies and governance controls. Third, reliability-oriented applications serve as valuable integration examples, as they encompass degradation, partial observability, inventory, mission risk, reconfiguration, and human accountability within a unified decision loop.

The broader implication is that Industry 4.0 represents a revolution in decision design rather than solely a technological advancement. Connected assets, digital twins, predictive models, and learning algorithms generate value by enhancing the quality of the entire decision loop. Future research and implementation should prioritize the disciplined conversion of observation into governed action over focusing on isolated technological concepts.

References

1. Q. Qi and F. Tao, "Digital twin and big data towards smart manufacturing and industry 4.0: 360 degree comparison," *IEEE Access*, vol. 6, pp. 3585-3593, 2018.
2. F. Rosin, P. Forget, S. Lamouri, and R. Pellerin, "Enhancing the decision-making process through industry 4.0 technologies," *Sustainability*, vol. 14, no. 1, p. 461, 2022.
3. F. Tao, Q. Qi, L. Wang, and A. Y. C. Nee, "Digital twins and cyber–physical systems toward smart manufacturing and industry 4.0: Correlation and comparison," *Engineering*, vol. 5, no. 4, pp. 653-661, 2019.
4. M. Achouch, M. Dimitrova, K. Ziane, S. Sattarpanah Karganroudi, R. Dhouib, H. Ibrahim, and M. Adda, "On predictive maintenance in industry 4.0: Overview, models, and challenges," *Applied Sciences*, vol. 12, no. 16, p. 8081, 2022.
5. M. Hermann, T. Pentek, and B. Otto, "Design principles for industrie 4.0 scenarios," in *2016 49th Hawaii International Conference on System Sciences (HICSS)*, pp. 3928-3937, 2016.
6. L. D. Xu, E. L. Xu, and L. Li, "Industry 4.0: state of the art and future trends," *International Journal of Production Research*, vol. 56, no. 8, pp. 2941-2962, 2018.
7. T. Zonta, C. A. Da Costa, R. da Rosa Righi, M. J. De Lima, E. S. Da Trindade, and G. P. Li, "Predictive maintenance in the Industry 4.0: A systematic literature review," *Computers & Industrial Engineering*, vol. 150, p. 106889, 2020.
8. J. Lee, B. Bagheri, and H. A. Kao, "A cyber-physical systems architecture for industry 4.0-based manufacturing systems," *Manufacturing Letters*, vol. 3, pp. 18-23, 2015.
9. M. Javaid, A. Haleem, and R. Suman, "Digital twin applications toward industry 4.0: A review," *Cognitive Robotics*, vol. 3, pp. 71-92, 2023.
10. H. Lasi, P. Fettke, H. G. Kemper, T. Feld, and M. Hoffmann, "Industry 4.0," *Business & Information Systems Engineering*, vol. 6, no. 4, pp. 239-242, 2014.
11. M. Ghobakhloo, "Industry 4.0, digitization, and opportunities for sustainability," *Journal of Cleaner Production*, vol. 252, p. 119869, 2020.
12. Y. Lu, "Industry 4.0: A survey on technologies, applications and open research issues," *Journal of Industrial Information Integration*, vol. 6, pp. 1-10, 2017.
13. A. Bousdekis, K. Lepenioti, D. Apostolou, and G. Mentzas, "A review of data-driven decision-making methods for industry 4.0 maintenance applications," *Electronics*, vol. 10, no. 7, p. 828, 2021.