International Journal of Scientific Research in Engineering and Management (IJSREM)
Volume: 10 Issue: 10 | Oct - 2026 SJIF Rating: 8.659 ISSN: 2582-3930

Human-Centric Smart Manufacturing in Industry 5.0

Chethan R[1]

[1]Automotive Europe, L&T Technology Services, India

Abstract - Industry 5.0 shifts the aim of smart manufacturing toward human-centricity, sustainability, and resilience, building on the digital base of Industry 4.0. This paper is a structured narrative review of how three enabling technologies, artificial intelligence (AI), digital twins (DTs), and collaborative robots (cobots), are being developed for those aims. Sources come from peer-reviewed journals, standards, and official institutional documents, with recent work favored and foundational papers retained. We compare findings across studies, separate established practice from emerging research, and describe a layered architecture in which physical assets, DTs, AI services, and human operators exchange data and decisions. The reviewed work is mature at the level of individual components, but evidence for integrated deployments with humans in the loop is limited. Cobot safety standards cover physical contact much better than cognitive and organizational factors. Explainability, interoperability, cybersecurity, and workforce adaptation come up repeatedly as barriers. We discuss applications in automotive, aerospace, electronics, machine-tool, and general discrete manufacturing in qualitative terms. The gaps we identify concern human-state modeling, management of DT fidelity, trust calibration, adaptive task allocation, and measurement of sustainability and wellbeing outcomes.

Key Words: Industry 5.0, human-centric manufacturing, digital twin, collaborative robots, human–robot collaboration.

1. INTRODUCTION

Histories of industrial development usually list the same stages: mechanized production, electric mass production, and programmable automation, followed by the networked cyber-physical systems now called Industry 4.0 [1], [2]. The Industry 4.0 label came out of a German national initiative that proposed linking machines, products, and information systems so that factories could make individualized goods flexibly [2]. Later surveys describe its technical base as sensing, connectivity, cloud and edge computing, and analytics, used mostly to raise productivity, flexibility, and quality [1].

Much Industry 4.0 work has been criticized for putting technology first. The factory is treated as something to instrument and optimize, and the worker shows up either as a source of variability or as a role to shrink. Human-centered alternatives were proposed early on. The Operator 4.0 idea described workers equipped with wearables, assistive devices, and better interfaces [4]. Human–cyber–physical systems (HCPS) went further and argued that human and machine intelligence should be designed as parts of a single system [3]. Both proposals rest on the same observation: automating a task does not guarantee a better system when the system still depends on human judgment, adaptability, and willingness to use it.

Industry 5.0 grew out of that criticism. Nahavandi presented it as cooperation between workers and intelligent machines, with customization as a main goal [6]. The European Commission later turned it into a policy vision that sits alongside Industry 4.0, puts worker wellbeing at the center, and treats sustainability and resilience as goals equal to productivity [5]. Scholarly discussion has since covered its foundations and enabling technologies [7]–[9]. The idea is still disputed. Some authors ask whether it is a separate stage at all or simply a change of priorities within Industry 4.0 [10], and others say its three themes need a unifying framework, for example one based on viability [11].

Existing reviews of Industry 5.0 mostly survey enabling technologies in broad terms [8], [9]. Few of them ask how AI, digital twins (DTs), and collaborative robots (cobots) actually depend on each other, or which links between them have little supporting evidence. That question matters because human-centered goals are met or missed at those links. An AI recommendation reaches an operator through a DT interface, and a cobot's behavior is limited by safety logic and shaped by predictions.

This review synthesizes and critically compares the literature on AI, DTs, and cobots in human-centric manufacturing; proposes a layered reference architecture that makes their technical relationships explicit; compares Industry 4.0 and Industry 5.0 in a structured way; and identifies research gaps and realistic future directions. The architecture in Section VI is a synthesis of the reviewed literature, not a validated implementation, and no experimental results are reported. The work is a structured narrative review, not a PRISMA systematic review, and no counts of screened papers are claimed. Following general guidance for narrative reviews [12], sources were found through topic-based searching and citation chaining from influential papers and from policy and standards documents. Inclusion required peer-reviewed or established-publisher status (or official institutional or standards status), direct relevance to manufacturing, and technical substance. Recent work was prioritized, and foundational publications were kept where they define concepts still in use. Evidence is classed as established or emerging. The selection is non-exhaustive, so conclusions reflect the reviewed set.

2. INDUSTRY 5.0

There is no agreed definition. The European Commission report, the most cited policy source, describes Industry 5.0 as a complement to Industry 4.0 that points technological progress at societal goals, organized around human-centricity, sustainability, and resilience [5]. Xu et al. place the idea in a longer line of development and observe that Industry 4.0 stresses technology and value chains, whereas Industry 5.0 stresses societal values [7]. Leng et al. gather existing definitions, organize the field around the same three values, and list approaches such as human-centered AI, bio-inspired technology, and DTs as enablers [8]. Maddikunta et al. cover a wider range of enabling technologies and applications [9]. All of these sources agree that Industry 5.0 keeps the Industry 4.0 technology base. They disagree mainly about how new the idea is and how it should be put into practice, a point we return to in Section II-E.

The literature uses human-centricity in two senses. In the first, technology augments the worker: cobots take over strenuous tasks, interfaces show information in context, and assistive systems reduce mental load. This reading dominates the technical papers, including the Operator 4.0 work [4] and HCPS frameworks [3], [16]. In the second sense, the system is designed so that workers keep meaningful roles, build skills, and share in the benefits, which is how the policy documents frame it [5]. Many technical papers borrow the vocabulary of the second sense but measure only the first, through productivity, task time, or error rate. Wang et al. point out that the human's role in HCPS changes as the system develops and that human factors have to be built in from the start [16]. This mismatch between stated aims and measured outcomes comes up again in Section IX.

Industry 4.0 technologies have been tied to economic, environmental, and social sustainability benefits, but those benefits depend on how the technologies are used [13]. For Industry 5.0, Ghobakhloo et al. built a roadmap from a content-centric literature review, identifying functions through which Industry 5.0 can support inclusive sustainable manufacturing and how those functions relate to each other [14]. Beyond these studies, two weaknesses stand out in what we read. Efficiency gains are usually argued from what a system could achieve and are rarely measured on a real shop floor. And the footprint of the digital infrastructure itself, meaning sensors, computation, and data storage, is almost never counted. We therefore treat sustainability as a design requirement with measurable indicators, not as something digitalization delivers by default.

In the policy framing, resilience means staying robust when disruptions occur, such as supply shocks or sudden changes in conditions [5]. Ivanov links resilience, sustainability, and human-centricity through viability, which asks whether a system can survive and adapt as conditions change [11]. Ivanov and Dolgui describe a digital supply chain twin for handling disruption risk [15], which shows how DTs connect to resilience at the supply-chain level. On the shop floor, resilience is closely tied to reconfigurability, meaning how easily cobots, flexible cells, and software-defined control can be redeployed. There is also a human side: experienced operators often absorb disturbances that no model anticipated. These levels are studied separately, and work that connects them is rare.

Most of the sources we reviewed treat the two as complementary. Industry 5.0 keeps the cyber-physical infrastructure, connectivity, and analytics of Industry 4.0 and changes the goal and the role given to people [5], [7], [8]. Critics note that human-centered and sustainability concerns already appeared in parts of the Industry 4.0 literature, which blurs any sharp break [10]. We take a practical position here. Industry 5.0 is a reordering of design priorities on top of Industry 4.0 technology, and our comparison in Section VII describes differences of emphasis.

3. ARTIFICIAL INTELLIGENCE IN INDUSTRY 5.0

Wuest et al. reviewed machine learning (ML) in manufacturing and listed difficulties that later work still reports: high-dimensional and mixed data, few labeled examples, poor data quality, and the need to include process knowledge [17]. Supervised classification and regression cover most applications, such as quality prediction and fault classification. Across these studies, the choice of algorithm seems to matter less than having suitable data and a well-posed problem, which helps explain why results from laboratory benchmarks often fail to carry over to production lines.

Deep learning builds layered representations straight from raw data and has driven progress in vision and signal processing [18]. Wang et al. reviewed its use in smart manufacturing, including convolutional networks, autoencoders, and recurrent models for defect detection, fault diagnosis, and prognostics [19]. The reviews agree that deep models cut down on hand-crafted features, but they need large datasets and are hard to interpret. The second drawback becomes serious once model outputs are shown to people who must act on them (Section III-F).

Vision is the most mature AI technique in manufacturing, because inspection tasks are well defined and images are relatively easy to collect. Defects, however, are rare and varied, which pushes researchers toward unsupervised anomaly detection. The MVTec AD dataset was released to test such methods on industrial-style images [20]. Benchmarks like this allow comparison, yet they use carefully prepared images, and few studies report how methods cope with lighting changes, new product variants, or drift over time. In Industry 5.0 settings, vision also serves people-related functions, for instance recognizing assembly progress or watching a shared workspace, which ties it to cobot safety and task allocation (Section V). Evidence for these uses under production conditions is still emerging.

Predictive maintenance is the most studied analytics application. Carvalho et al. reviewed ML for predictive maintenance and found that a handful of method families recur, with limited failure data and little real-world validation as common weaknesses [21]. Lei et al. lay out the prognostics pipeline, which runs from data acquisition through health-indicator construction to remaining-useful-life prediction [22]. DT-based maintenance (Section IV-D) now commonly follows this structure. Overall, diagnosis works well on controlled data, but moving a model across machines, operating conditions, and sites remains unsolved.

AI supports decisions at several levels, including scheduling, quality control, maintenance planning, and process-parameter selection. Deep reinforcement learning has been applied to production scheduling [23], usually in simulation. Two techniques matter especially for Industry 5.0. Federated learning trains a shared model across sites without pooling raw data [24], which eases data-sharing worries and, to some extent, protects worker privacy. Physics-informed neural networks add governing equations as constraints, which can lower the amount of data needed and keep predictions consistent with physical laws [25]. Both look promising for manufacturing, though in the sources we reviewed the case for them is stronger on paper than in industrial trials.

Parasuraman et al. split automation into four stages: acquiring information, analyzing it, selecting a decision, and carrying out an action, and they argue that each stage can be automated to a different degree [26]. The idea is useful for Industry 5.0 because it allows designs where AI does the analysis and a person keeps the decision. How much people rely on automation depends on calibrated trust, which in turn depends on how reliable the system appears, how transparent it is, and how well the user understands it [27]. Too much trust and too little trust both hurt performance. Explainable AI targets the transparency part. A

large survey separates models that are interpretable by design from post-hoc explanations and stresses that an explanation should be written for a particular audience [28]. In manufacturing, explanation methods are seldom tested with real operators, so claims that explanations improve decisions or trust on the factory floor are still assumptions.

Table II links AI techniques to Industry 5.0 manufacturing tasks. Governance has become part of the engineering picture as well. The European guidelines on trustworthy AI call for human agency and oversight, technical robustness and safety, transparency, and accountability [29]. The EU AI Act sets up a risk-based legal framework, with duties for high-risk systems that include human oversight [30]. A manufacturer who puts AI into a safety-relevant function will face requirements beyond accuracy, and that affects architecture decisions (Section VI).

Table-1: AI technologies and their applications in industry 5.0 manufacturing

AI technology	Typical application	Human-centric relevance	Principal limitation (reviewed)
Supervised ML	Quality prediction, fault classification [17]	Decision support to engineers	Labeled-data scarcity; generalization
Deep learning	Defect detection, diagnosis, prognostics [19]	Reduces manual inspection burden	Data demand; limited transparency
Computer vision / anomaly detection	Surface inspection, workspace monitoring [20]	Operator-facing quality assistance; safety perception	Robustness to variants and conditions
Predictive analytics	Remaining useful life, maintenance planning [21], [22]	Maintenance workload planning	Few failure samples; limited field validation
Reinforcement learning	Scheduling and control optimization [23]	Adaptive allocation	Mostly simulation-based evidence
Federated learning	Cross-site model training [24]	Data governance and privacy support	Limited industrial validation
Physics-informed models	Hybrid modeling in DTs [25]	Consistent, interpretable behavior	Modeling effort; scaling
Explainable AI	Operator-facing explanations [28]	Trust calibration [27]	Rarely evaluated with operators

4. DIGITAL TWINS FOR INDUSTRY 5.0

The DT idea is usually traced to a model with three parts: a physical product, its virtual counterpart, and the link between them [31]. A common taxonomy separates a digital model, with no automatic data exchange, a digital shadow, with automatic flow from the physical to the digital side, and a digital twin, with automatic flow in both directions [32]. By this definition, much of the work called “digital twin” is really a model or a shadow. Tao et al. reviewed industrial DTs, covering enabling technologies and application areas [33], and Lu et al. proposed a reference model for DT-driven smart manufacturing along with open research issues [35]. The ISO 23247 series defines a framework for manufacturing DTs, including a reference architecture [34]. The sources broadly agree on a layered structure: data acquisition, communication, a core modeling layer, and application services. They differ on how faithful the model must be, how often it must update, and how much autonomy justifies the word “twin” [36].

Monitoring is the simplest DT function. Sensor and controller data are mapped onto a model of the state of an asset or process, which depends on good data acquisition and interoperable communication (Section IX). The central technical question is synchronization: how closely and how quickly the virtual state follows the physical one, since this decides whether later predictions can be trusted. Reviews report that latency and accuracy requirements are described inconsistently [36], so results from different studies are hard to compare.

Simulation is what separates a DT from a dashboard. A twin can test what-if scenarios, such as a new schedule, a layout change, or a parameter adjustment, before anything is applied to the real system [33], [35]. For machine tools, the Twin-Control project documents a DT approach that combines machine, process, and monitoring models across the machine’s life [38]. In Industry 5.0, simulation also helps check human–robot layouts and ergonomics in advance (Section V-E). What limits all of this is how far the model can be believed. Rasheed et al. stress that the value of a DT depends on modeling choices, handling of uncertainty, and validation, all of which remain difficult [37].

DT-based maintenance combines state estimation from monitoring with the prognostics methods of Section III-D. A physical model inside the twin can provide features or constraints that purely data-driven prognostics lack, which makes hybrid physics-and-data approaches attractive [25], [37]. Reviews warn that many studies of DT-based maintenance are conceptual or demonstrated only on laboratory rigs [36]. Condition monitoring and data-driven prognostics are the established part [21], [22]. The emerging part is closed-loop use, where a twin’s prediction automatically triggers maintenance scheduling that accounts for production and staff availability.

For people, the value of a DT is that it shows the consequences of an action before it is taken. We see two uses. One is as a

decision-support environment where operators and engineers try out options; this calls for interfaces that expose assumptions and uncertainty, and such interfaces are rarely tested with users. The other is the *human digital twin* (HDT), which models a worker's physiological, behavioural, or ergonomic state. Wang et al. surveyed HDTs for Industry 5.0, proposing a conceptual framework and architecture and analysing enabling technologies and applications [39]. HDTs speak directly to the personalization and wellbeing goals of Industry 5.0, but they are the least mature kind of DT and they raise the ethical issues covered in further Section.

AI and DTs support each other. AI gives the twin perception, surrogate models, and optimization, while the twin gives AI synthetic data, a safe test bed, and current context [35], [37]. When models are trained on simulated data and then deployed on real equipment, the gap between simulation and reality appears, so the fidelity of the twin sets an upper limit on how reliable the AI can be. Connecting twins to execution and control systems needs shared information models and communication standards [34]; without them every twin becomes a one-off integration. Table III summarizes DT applications and how mature they are.

Table-2: Digital twin applications in industry 5.0

Application	DT function	Link to Industry 5.0 pillar	Maturity (per reviewed sources)
Real-time monitoring	State mapping of assets and processes [33], [35]	Resilience (visibility)	Established at digital-shadow level [32]
Process simulation and optimization	What-if evaluation [35], [38]	Sustainability; resilience	Established in selected domains; fidelity issues [37]
Predictive maintenance	State estimation plus prognostics [21], [22]	Sustainability; resilience	Mixed; closed-loop use emerging [36]
Virtual commissioning of cobot cells	Interaction and layout testing [47]	Human-centricity (safety, ergonomics)	Case-study level
Human digital twin	Worker state and ergonomics modeling [39]	Human-centricity	Emerging
Supply-chain twin	Disruption analysis [15]	Resilience	Emerging

5. COLLABORATIVE ROBOTICS

Traditional industrial robots work behind guarding, away from people, on repetitive high-speed tasks. Krüger et al. examined how humans and machines can cooperate on assembly lines, compared their complementary strengths, and pointed to the potential of shared workspaces [40]. Cobots are built for, or allowed in, shared space, and standards set the conditions. ISO 10218 gives safety requirements for industrial robots and robot systems [43], and ISO/TS 15066 adds requirements for collaborative operation [44]. Both describe four collaborative operation types: safety-rated monitored stop, hand guiding, speed and separation monitoring, and power and force limiting. Villani et al. surveyed the field and organized it around safety, intuitive interfaces, and applications [41].

Human–robot collaboration (HRC) covers a range, from people and robots working in separate zones to joint work on one task with physical contact. Wang et al. reviewed symbiotic HRC in assembly, covering perception, planning, and control, and argued that closer collaboration requires the robot to understand human intent and the state of the task [42]. The surveys [40]–[42] suggest that industrial deployments sit mostly at the lower interaction levels, while simultaneous work on the same part is still largely a research topic. Keeping this difference between established and emerging practice in mind is useful when reading Industry 5.0 claims.

Safety is the main constraint on cobot deployment. The standards apply to the whole application and not just the robot: a robot labeled collaborative can still be unsafe with a sharp tool or heavy payload, so each installation needs its own risk assessment [43], [44]. ISO/TS 15066 gives guidance on permissible contact in power and force limiting, including limits for different body regions [44]. Speed and separation monitoring slows the robot as a person gets closer, and Marvel and Norcross discuss how stopping distance and sensor latency determine the separation needed [45]. Two limits follow. The standards deal mainly with physical hazards, while psychological effects such as stress or a sense of lost control get little formal treatment, even though they are central to human-centricity [5]. And ML-based perception or decision logic inside a safety function is hard to certify. The EU Machinery Regulation addresses safety components with self-evolving behavior based on ML [48], which supports our position in Section VI that safety-critical functions should stay verifiable.

Task allocation decides which agent, human or robot, does which operation and when. Fixed allocation based on assumed abilities is the simplest approach; adaptive allocation responds to a worker's state, skill, and preferences, and to disruptions. Gombolay et al. formulated mixed-initiative scheduling for human–robot teams that takes into account factors such as workload and workflow preferences [46]. In the allocation work we reviewed, the optimization side is mature and the weak spot is the human model, meaning validated descriptions of fatigue, learning, and shifting preferences. An allocation scheme is only as human-centered as the human model behind it, and most of those models are simple.

Cobots add flexibility through fast redeployment, lead-through programming, and tolerance for shared layouts. DTs help by enabling virtual commissioning. Malik and Brem present a case study in which a DT of a cobot is used to design and test human–robot interaction before physical deployment [47]. The

expected benefits are less commissioning effort and safer design iterations. One caution is that virtual checks of safety behavior are only as good as the human and contact models in the simulation.

Interaction can happen through hand guiding, gestures, speech, graphical or augmented-reality interfaces and, more experimentally, physiological signals. The surveys identify easy programming and clear communication of the robot's intent as keys to acceptance [41], [42]. For work to flow, the robot has to be predictable to the person, which limits how far learning-based adaptation should change its behavior. This tension between adaptivity and predictability is unresolved in the sources we reviewed and remains a design trade-off.

6. INTEGRATION OF AI, DIGITAL TWINS, AND COBOTS

Drawing on the reviewed sources, we propose a layered reference architecture. It is our own synthesis and has not been validated. The layers are as follows. The physical layer holds machines, cobots, sensors, wearables, and human operators. The connectivity and edge layer moves data between them; platform-independent information modeling such as OPC UA [50] supports interoperable exchange, and edge computing handles latency-sensitive processing. The digital twin layer provides state estimation, process and robot models, and simulation, in line with the reference structures in [34], [35]. The AI services layer provides perception, prognostics, surrogate modeling, and optimization (Section III). The decision and interaction layer is where people see predictions, explanations, and options, and where they approve, change, or override. A governance layer cuts across all of these and covers safety, cybersecurity, data protection, and ethical oversight [29], [30], [48], [49].

Data move up from the physical layer as sensor streams, controller states, and operator-related signals, and down as setpoints, schedules, recommendations, and robot task parameters. In the shadow-and-twin taxonomy [32], the automatic downward path is what makes a system a twin and not a shadow. It is also the riskiest path, because an error in the twin can turn into a physical action. For this reason, downward commands that affect safety should pass through validated, deterministic logic and should never come straight from an AI or simulation output.

One way to understand how the technologies fit together is to look at three nested loops running at different speeds. We use this as an analytical device and not as a standard model.

Safety loop (milliseconds). Protective stops, speed and separation monitoring, and force limiting run on certified controllers and safety sensors [43]–[45]. This loop should not depend on AI inference or on the twin.

Tactical loop (seconds to minutes). AI services use the twin's state to adjust task allocation, robot parameters, or operator guidance within the limits set by the safety loop. Adaptive allocation [46] and human-state estimation [39] work at this level, and twin predictions of operator and robot behavior can be used to anticipate conflicts, as in DT-supported collaborative execution [47].

Strategic loop (hours to months). The twin evaluates schedules, layouts, maintenance plans, and reconfiguration options. AI proposes candidates and people choose among them. Resilience and sustainability enter here as explicit objectives [11], [14], [15].

The dependencies then run as follows. The twin gives AI context and a place to test; AI gives the twin predictive and surrogate models [37]; cobots both carry out decisions and produce data; and people act as supervisors, collaborators, and sources of data.

Because a twin can simulate, plans generated by AI can be screened virtually, and only candidates that meet the constraints reach the physical system. This makes it possible to use reinforcement learning and optimization methods that would be unsafe to explore directly on production equipment [23], [35]. The approach is only as valid as the twin is accurate, so the architecture should monitor the gap between twin and plant. If that gap exceeds a set tolerance, the autonomy given to AI recommendations should shrink, which is consistent with graded levels of automation [26].

Human oversight has to be designed in. It needs a suitable level of automation at each stage [26], explanations that match the decision [28], interfaces that support calibrated trust [27], and a workable way to override the system. The oversight provisions of the AI Act make this a compliance matter for high-risk uses [30]. We found few integrated studies that measure a supervisor's workload, situational awareness, or reaction to automation errors in combined AI–DT–cobot settings.

7. COMPARATIVE ANALYSIS

Table 3 compares Industry 4.0 and Industry 5.0 on ten dimensions. It follows the complementary view of Section II-E and describes differences of emphasis drawn from [5], [7], [8]; it does not claim that Industry 4.0 research ignored people, sustainability, or resilience. One pattern is worth stating plainly. The dimensions where Industry 5.0 differs most, namely the role of humans, sustainability, resilience, and human–machine collaboration, are also the ones with the least developed indicators and validated methods in the technical literature. The table therefore describes an agenda more than an achieved state.

Table-3: Comparison of industry 4.0 and industry 5.0

Dimension	Industry 4.0	Industry 5.0
Primary objective	Productivity, flexibility, efficiency through digitalization [1], [2]	Human wellbeing, sustainability, and resilience alongside productivity [5]
Role of humans	Operator or supervisor; augmentation in some work [4]	Central collaborator; agency, skills, and wellbeing as design goals [5], [8]
Automation	Emphasis on increasing automation and connectivity	Automation chosen by task and human role; graded levels [26]

Dimension	Industry 4.0	Industry 5.0
AI	Analytics, prediction, optimization for performance	Human-centric AI: explainable, trustworthy, overseen [28]–[30]
Robotics	Industrial robots, increasingly flexible	Cobots in shared workspaces with adaptive behavior [41], [42]
Sustainability	Opportunity via efficiency; conditional [13]	Explicit objective with societal framing [5], [14]
Resilience	Mostly supply-chain and system robustness via visibility	Explicit pillar; viability perspective [5], [11], [15]
Decision-making	Data-driven, often centralized or automated	Shared human–AI; oversight and override [26], [30]
Customization	Mass customization via flexible production	Personalization extended to workers and tasks [6], [39]
Human–machine collaboration	Interfaces and assistance [4]	Collaboration as a design principle across physical and cognitive levels [3], [42]

8. INDUSTRIAL APPLICATIONS

We discuss application areas in qualitative terms. We cite no company case studies or performance numbers, because the reviewed literature supports mostly statements about capability or pilot results. In every area the evidence is stronger for single technologies than for integrated human-centered systems.

Automotive manufacturing. Many product variants and long-established assembly lines make this the standard reference area for HRC; early analyses of human–machine cooperation on assembly lines were motivated by it [40]. Typical candidates are ergonomically demanding or repetitive operations, where robots assist and people keep the dexterous or judgment-heavy steps [41], [42]. DT-supported design and testing of cobot cells has been shown at case-study level [47]. Open questions include whether safe collaborative speeds fit within takt time and how to allocate tasks adaptively when the variant mix changes [46].

Aerospace manufacturing. Low volumes, high mix, large structures, and strict quality documentation make this a natural fit for flexible, human-in-the-loop systems in principle. DTs are relevant for process validation and for tracing product state [33], [35], and cobots are of interest where fixed automation does not pay off. We found little evidence on integrated deployment here, so our statements are forward-looking and rest on general capability.

Electronics manufacturing. Vision-based inspection and anomaly detection fit most directly [19], [20], since components are small, defect types are many, and inspection produces a lot of data. Frequent changeovers match the flexibility argument for cobots and the use of federated learning to reuse models across lines [24]. The main practical worry is how well vision models survive design changes.

Machine-tool manufacturing. DTs of machine tools for monitoring, process supervision, and lifecycle support are documented, with Twin-Control as a published example [38]. Predictive maintenance and prognostics are mainstream topics in this area [21], [22]. The human-centered side lies in operator-facing decision support, with the twin helping during setup and diagnosis. Evidence on what this does for operators is sparse.

General discrete manufacturing. In small-batch and job-shop settings, scheduling and flexible cells matter most. Learning-based scheduling [23] and mixed-initiative human–robot scheduling [46] address this, mostly in simulation or controlled experiments. Cobot cells suit settings that are reconfigured often, and DT-based commissioning aims to reduce the engineering effort involved [47].

9. CHALLENGES AND RESEARCH GAPS

Cybersecurity. Linking operational technology to analytics and DTs enlarges the attack surface, and because of the downward control path (Section VI-B), a compromised AI or twin component can have physical effects. The IEC 62443 series defines system security requirements and security levels for industrial automation [49]. The DT and AI papers we reviewed rarely treat security as a design requirement, and secure synchronization of twins and the integrity of models are under-studied.

Data interoperability and standardization. OPC UA provides a platform-independent standard for information modeling and communication [50], and ISO 23247 offers a DT framework [34]. In practice, vendors use different data models and twins are defined inconsistently [36], which hinders reuse. The standards for information models also say little about the semantics needed for human-related data and for AI model metadata.

Legacy-system integration. Existing equipment often has no sensors or open interfaces. Retrofitting with gateways and edge devices is practical, but we found few systematic methods for deciding which legacy assets are worth instrumenting or for estimating the data quality that retrofits can reach.

Workforce skills. Operator-focused frameworks assume that workers can use advanced interfaces and interpret AI output [4], [8]. Training, role redesign, and long-term skill development are discussed in concept, but empirical studies of upskilling in AI–DT–cobot environments are scarce.

Implementation cost. Building, validating, and maintaining twins and AI models carries ongoing costs. Rigorous lifecycle cost–benefit methods are rare in the reviewed literature, and

many claimed benefits come without cost accounting. For that reason we make no quantitative claim about cost.

Human–machine safety. Beyond the physical safety covered by the standards [43], [44], open problems include certifying adaptive or ML-enabled behavior [48] and psychological safety. Virtual safety assessments also need more comparison against physical measurements.

AI explainability and trust. Explanation techniques exist [28] and trust theory offers design principles [27], but explanations have seldom been evaluated with manufacturing operators working under time pressure in noisy environments.

Data requirements. Rare failures, rare defects, and scarce labels limit ML [17], [21]. Federated and physics-informed methods are promising [24], [25], but there is little industrial validation.

Ethical considerations. Worker monitoring, human digital twins, and algorithmic task assignment raise questions of privacy, autonomy, and fairness [29], [30], [39]. Technical papers often describe the benefits of sensing human state without discussing consent, data governance, or the possibility that the data will be used for surveillance.

Research gaps. Putting these points together, the reviewed literature supports the following gaps.

- Few integrated evaluations of AI–DT–cobot systems with human participants under realistic conditions [42], [46].
- Human models of fatigue, learning, preference, and intent are simple and rarely validated in allocation and safety settings [39], [46].
- There is no shared method for specifying and verifying the fidelity, latency, and uncertainty a DT needs for a given decision [36], [37].
- Safety assurance for adaptive, learning-enabled collaborative behavior is unresolved [44], [48].
- Frameworks for measuring how technical systems affect wellbeing, sustainability, and resilience are underdeveloped [5], [14].
- Explanation methods have not been evaluated for their effect on operator decision quality and trust calibration [27], [28].
- Most DT and human-DT proposals do not treat security and privacy as first-order architectural requirements [39], [49].

Table-4: Research challenges and future directions

Challenge / gap	Technically realistic direction
Integrated human-participant evaluation	Controlled human-in-the-loop studies of AI–DT–cobot cells
Simplified human models	Validated fatigue, intent, and skill models for allocation
DT fidelity specification	Fidelity, latency, and uncertainty reporting; divergence monitoring
Assurance of adaptive behavior	Bounded, verifiable adaptation within safety envelopes
Outcome measurement	Indicators for wellbeing, energy, and resilience linked to system design
Explanation evaluation	Role-specific explanations tested with operators
Security and privacy	Security-by-design for DTs; privacy-preserving human DTs
Interoperability and standards	Mapping DT, OPC UA, and AI metadata models

10. FUTURE RESEARCH DIRECTIONS

Human–AI collaboration. Evaluation should shift from algorithm-centered tests to controlled studies with operators, measuring decision quality, workload, and trust calibration at different levels of automation [26], [27]. Explanations should be designed for specific roles, such as operator, maintenance engineer, and planner, and tested for usefulness.

AI-enabled digital twins. Realistic directions include hybrid physics-and-data models with quantified uncertainty [25], [37], surrogate models for faster what-if analysis, and automatic detection of divergence between twin and plant that triggers recalibration or reduced autonomy. Federated approaches [24] could allow learning across sites with limited data sharing.

Adaptive cobots. Useful work includes intent recognition, adaptive allocation based on validated human models, and adaptation kept within predictability limits [42], [46]. Adaptation must stay inside envelopes that can be verified, in line with the safety architecture. Standards for adaptive behavior are still needed [48].

Sustainable manufacturing. DT-supported modeling of energy and resource use, including the computational footprint of the DT itself, is a realistic direction [14]. Multi-objective optimization that weighs environmental and ergonomic indicators alongside throughput would put the Industry 5.0 goals into practice.

Resilient manufacturing. Resilience models that span scales, linking shop-floor reconfigurability to supply-chain twins [15], and viability-based evaluation [11], are plausible lines of work. Libraries of disruption scenarios for testing inside twins could support both.

Human-centric decision-support systems. Such systems would show uncertainty, assumptions, and alternatives and let users override and learn from them, which draws the earlier directions together. HDTs could personalize the support if privacy-preserving design and governance are built in from the start [29], [30], [39]. Table V summarizes the challenges and directions.

11. CONCLUSION

We reviewed AI, digital twins, and collaborative robots as enabling technologies for human-centered, sustainable, and resilient manufacturing. Each is well developed at the component level, with ML-based inspection and prognostics, standards-based DT frameworks, and defined collaborative operation types. The main finding is that human-centered outcomes depend on the links between the technologies, and those links have the thinnest evidence. The literature has more algorithms and architectures than human-subject validation, lifecycle cost analysis, or measurement of effects on wellbeing and sustainability.

The layered architecture and its split into safety, tactical, and strategic loops give a way to decide where AI and DTs can act on their own and where certified deterministic logic and human authority must remain. We identified seven gaps: integrated evaluation with human participants, validated human models, specification of DT fidelity, assurance of adaptive behavior, outcome measurement, evaluation of explanations, and security and privacy. This review has limits. Its selection is narrative and incomplete, the published work it draws on leans toward laboratory and pilot evidence, and the proposed architecture is conceptual. The next step is to test it in controlled studies with people in the loop.

REFERENCES

1. H. Lasi, P. Fettke, H.-G. Kemper, T. Feld, and M. Hoffmann, “Industry 4.0,” Bus. Inf. Syst. Eng., vol. 6, no. 4, pp. 239–242, 2014.

2. H. Kagermann, W. Wahlster, and J. Helbig, Recommendations for Implementing the Strategic Initiative Industrie 4.0: Final Report of the Industrie 4.0 Working Group. Frankfurt/Main, Germany: acatech, Apr. 2013.

3. J. Zhou, Y. Zhou, B. Wang, and J. Zang, “Human–cyber–physical systems (HCPSs) in the context of new-generation intelligent manufacturing,” Engineering, vol. 5, no. 4, pp. 624–636, 2019.

4. D. Romero, P. Bernus, O. Noran, J. Stahre, and Å. Fast-Berglund, “The Operator 4.0: Human cyber-physical systems & adaptive automation towards human-automation symbiosis work systems,” in Advances in Production Management Systems, IFIP AICT, vol. 488. Cham, Switzerland: Springer, 2016, pp. 677–686.

5. M. Breque, L. De Nul, and A. Petridis, Industry 5.0: Towards a Sustainable, Human-Centric and Resilient European Industry. Luxembourg: Publications Office of the European Union, 2021.

6. S. Nahavandi, “Industry 5.0—A human-centric solution,” Sustainability, vol. 11, no. 16, Art. no. 4371, 2019.

7. X. Xu, Y. Lu, B. Vogel-Heuser, and L. Wang, “Industry 4.0 and Industry 5.0—Inception, conception and perception,” J. Manuf. Syst., vol. 61, pp. 530–535, 2021.

8. J. Leng et al., “Industry 5.0: Prospect and retrospect,” J. Manuf. Syst., vol. 65, pp. 279–295, 2022.

9. P. K. R. Maddikunta et al., “Industry 5.0: A survey on enabling technologies and potential applications,” J. Ind. Inf. Integr., vol. 26, Art. no. 100257, 2022.

10. S. Grabowska, S. Saniuk, and B. Gajdzik, “Industry 5.0: Improving humanization and sustainability of Industry 4.0,” Scientometrics, vol. 127, no. 6, pp. 3117–3144, 2022.

11. D. Ivanov, “The Industry 5.0 framework: Viability-based integration of the resilience, sustainability, and human-centricity perspectives,” Int. J. Prod. Res., vol. 61, no. 5, pp. 1683–1695, 2023.

12. H. Snyder, “Literature review as a research methodology: An overview and guidelines,” J. Bus. Res., vol. 104, pp. 333–339, 2019.

13. M. Ghobakhloo, “Industry 4.0, digitization, and opportunities for sustainability,” J. Clean. Prod., vol. 252, Art. no. 119869, 2020.

14. M. Ghobakhloo, M. Iranmanesh, B. Foroughi, E. B. Tirkolaee, S. Asadi, and A. Amran, “Industry 5.0 implications for inclusive sustainable manufacturing: An evidence-knowledge-based strategic roadmap,” J. Clean. Prod., vol. 417, Art. no. 138023, 2023, doi: 10.1016/j.jclepro.2023.138023.

15. D. Ivanov and A. Dolgui, “A digital supply chain twin for managing the disruption risks and resilience in the era of Industry 4.0,” Prod. Plan. Control, vol. 32, no. 9, pp. 775–788, 2021.

16. B. Wang, P. Zheng, Y. Yin, A. Shih, and L. Wang, “Toward human-centric smart manufacturing: A human-cyber-physical systems (HCPS) perspective,” J. Manuf. Syst., vol. 63, pp. 471–490, 2022.

17. T. Wuest, D. Weimer, C. Irgens, and K.-D. Thoben, “Machine learning in manufacturing: Advantages, challenges, and applications,” Prod. Manuf. Res., vol. 4, no. 1, pp. 23–45, 2016.

18. Y. LeCun, Y. Bengio, and G. Hinton, “Deep learning,” Nature, vol. 521, no. 7553, pp. 436–444, 2015.

19. J. Wang, Y. Ma, L. Zhang, R. X. Gao, and D. Wu, “Deep learning for smart manufacturing: Methods and applications,” J. Manuf. Syst., vol. 48, pp. 144–156, 2018.

20. P. Bergmann, M. Fauser, D. Sattlegger, and C. Steger, “MVTec AD—A comprehensive real-world dataset for unsupervised anomaly detection,” in Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR), Long Beach, CA, USA, 2019, pp. 9592–9600.

21. T. P. Carvalho, F. A. A. M. N. Soares, R. Vita, R. da P. Francisco, J. P. Basto, and S. G. S. Alcalá, “A systematic literature review of machine learning methods applied to predictive maintenance,” Comput. Ind. Eng., vol. 137, Art. no. 106024, 2019.

22. Y. Lei, N. Li, L. Guo, N. Li, T. Yan, and J. Lin, “Machinery health prognostics: A systematic review from data acquisition to RUL prediction,” Mech. Syst. Signal Process., vol. 104, pp. 799–834, 2018.
[23] B. Waschneck et al., “Optimization of global production scheduling with deep reinforcement learning,” Procedia CIRP, vol. 72, pp. 1264–1269, 2018.

24. Q. Yang, Y. Liu, T. Chen, and Y. Tong, “Federated machine learning: Concept and applications,” ACM Trans. Intell. Syst. Technol., vol. 10, no. 2, Art. no. 12, 2019.

25. M. Raissi, P. Perdikaris, and G. E. Karniadakis, “Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations,” J. Comput. Phys., vol. 378, pp. 686–707, 2019.

26. R. Parasuraman, T. B. Sheridan, and C. D. Wickens, “A model for types and levels of human interaction with automation,” IEEE Trans. Syst., Man, Cybern. A, Syst., Humans, vol. 30, no. 3, pp. 286–297, May 2000.

27. J. D. Lee and K. A. See, “Trust in automation: Designing for appropriate reliance,” Hum. Factors, vol. 46, no. 1, pp. 50–80, 2004.

28. A. Barredo Arrieta et al., “Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI,” Inf. Fusion, vol. 58, pp. 82–115, 2020.

29. High-Level Expert Group on Artificial Intelligence, Ethics Guidelines for Trustworthy AI. Brussels, Belgium: European Commission, Apr. 2019.

30. Regulation (EU) 2024/1689 of the European Parliament and of the Council Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act), Off. J. Eur. Union, 2024.

31. M. Grieves and J. Vickers, "Digital twin: Mitigating unpredictable, undesirable emergent behavior in complex systems," in Transdisciplinary Perspectives on Complex Systems, F.-J. Kahlen, S. Flumerfelt, and A. Alves, Eds. Cham, Switzerland: Springer, 2017, pp. 85–113.

32. W. Kritzinger, M. Karner, G. Traar, J. Henjes, and W. Sihn, "Digital twin in manufacturing: A categorical literature review and classification," IFAC-PapersOnLine, vol. 51, no. 11, pp. 1016–1022, 2018.

33. F. Tao, H. Zhang, A. Liu, and A. Y. C. Nee, "Digital twin in industry: State-of-the-art," IEEE Trans. Ind. Informat., vol. 15, no. 4, pp. 2405–2415, Apr. 2019.

34. Automation Systems and Integration—Digital Twin Framework for Manufacturing, ISO 23247 series (Parts 1–4), 2021.

35. Y. Lu, C. Liu, K. I.-K. Wang, H. Huang, and X. Xu, "Digital twin-driven smart manufacturing: Connotation, reference model, applications and research issues," Robot. Comput.-Integr. Manuf., vol. 61, Art. no. 101837, 2020.

36. D. Jones, C. Snider, A. Nassehi, J. Yon, and B. Hicks, "Characterising the digital twin: A systematic literature review," CIRP J. Manuf. Sci. Technol., vol. 29, pp. 36–52, 2020.

37. A. Rasheed, O. San, and T. Kvamsdal, "Digital twin: Values, challenges and enablers from a modeling perspective," IEEE Access, vol. 8, pp. 21980–22012, 2020.

38. M. Armendia, M. Ghassempouri, E. Ozturk, and F. Peysson, Eds., Twin-Control: A Digital Twin Approach to Improve Machine Tools Lifecycle. Cham, Switzerland: Springer, 2019.

39. B. Wang et al., "Human digital twin in the context of Industry 5.0," Robot. Comput.-Integr. Manuf., vol. 85, Art. no. 102626, Feb. 2024.

40. J. Krüger, T. K. Lien, and A. Verl, "Cooperation of human and machines in assembly lines," CIRP Ann., vol. 58, no. 2, pp. 628–646, 2009.

41. V. Villani, F. Pini, F. Leali, and C. Secchi, "Survey on human–robot collaboration in industrial settings: Safety, intuitive interfaces and applications," Mechatronics, vol. 55, pp. 248–266, 2018.

42. L. Wang et al., "Symbiotic human-robot collaborative assembly," CIRP Ann., vol. 68, no. 2, pp. 701–726, 2019.

43. Robots and Robotic Devices—Safety Requirements for Industrial Robots—Part 1: Robots; Part 2: Robot Systems and Integration, ISO 10218-1 and ISO 10218-2, 2011.

44. Robots and Robotic Devices—Collaborative Robots, ISO/TS 15066, 2016.

45. J. A. Marvel and R. Norcross, "Implementing speed and separation monitoring in collaborative robot workcells," Robot. Comput.-Integr. Manuf., vol. 44, pp. 144–155, 2017.

46. M. Gombolay, A. Bair, C. Huang, and J. Shah, "Computational design of mixed-initiative human–robot teaming that considers human factors: Situational awareness, workload, and workflow preferences," Int. J. Robot. Res., vol. 36, no. 5–7, pp. 597–617, 2017.

47. A. A. Malik and A. Brem, "Digital twins for collaborative robots: A case study in human-robot interaction," Robot. Comput.-Integr. Manuf., vol. 68, Art. no. 102092, 2021.

48. Regulation (EU) 2023/1230 of the European Parliament and of the Council on Machinery, Off. J. Eur. Union, 2023.

49. Industrial Communication Networks—Network and System Security—Part 3-3: System Security Requirements and Security Levels, IEC 62443-3-3, 2013.

50. OPC Unified Architecture, IEC 62541 series, International Electrotechnical Commission.