

A Multi-Domain AI Analytics Framework for Smart Manufacturing Across the Product Lifecycle

Aisha Johansson¹; Daniel Carlsson²; Maja Lindqvist^{3, *}

¹ Department of Industrial Engineering and Management, KTH Royal Institute of Technology, Stockholm, Sweden

² School of Business, Economics and Law, University of Gothenburg, Gothenburg, Sweden

³ Department of Production Engineering, Chalmers University of Technology, Gothenburg, Sweden

* Corresponding author: maja.lindqvist@chalmers.se

ARTICLE INFO Received July 12, 2023 Revised September 28, 2023 Accepted November 18, 2023 Available Online December 30, 2023 DOI 10.63646/jaiaa.2023.010401 License Creative Commons Attribution 4.0 International Licence (CC BY 4.0) Publisher INATGI, United States of America Journal JAIAA - ISSN 3067-7386	Abstract Smart manufacturing (SM) has emerged as the operational expression of Industry 4.0 and of the broader convergence of artificial intelligence (AI), the industrial internet of things (IIoT), cyber-physical systems (CPS) and digital twins (DTs). While the cumulative number of SM publications has grown roughly eighty-fold over the last decade, most contributions remain locally scoped: a single algorithm in a single workshop, a single sensor family on a single asset class, or a single management process in a single enterprise. The field lacks a view that places those contributions inside a common product-lifecycle (PLM) architecture and that identifies, for each lifecycle stage, which analytics tools are mature, which are experimental, and which are missing. This article develops such a view. We conduct a systematic review of 214 peer-reviewed works published between 2012 and 2025, organise their AI contributions into four interlocking domains—strategy and organisation, value-chain intelligence, management support, and infrastructure and capabilities—and map the resulting topics onto a five-stage PLM backbone. Using the coded corpus we report a maturity radar, a barrier frequency analysis and a topic-share breakdown. We then translate the diagnostic into prescriptive guidance in the form of a five-layer AI-enabled transformation roadmap that covers foundations of AI capability, enterprise-wide integration, process optimisation, system reconfiguration and enterprise innovation. Two cross-cutting findings emerge. First, the dominant weakness of current SM systems is not algorithmic but architectural: data fragmentation, missing standards and poor horizontal-vertical-end-to-end integration explain most of the gap between pilot-scale success and enterprise-scale deployment. Second, the sustainability of SM will be decided less by the raw capability of frontier AI models than by how disciplined enterprises are in building knowledge-management (KM) capability, governance of data and workforce AI literacy. The framework is offered as a diagnostic grid for practitioners and as a structured research agenda for scholars. Keywords: smart manufacturing, Industry 4.0, artificial intelligence, product lifecycle management, industrial internet of things, digital twin, systematic literature review, enterprise transformation.
--	--

I. INTRODUCTION

Manufacturing has entered a phase in which information technology is no longer an overlay on top of physical production but a constitutive part of how production is designed, scheduled and controlled. Industry 4.0 and the parallel national programmes in Germany, Japan, the United States and China framed this shift as a generational reconfiguration of the factory [1], [2], [3]. Smart manufacturing (SM) is the operational label for the result. In SM, physical assets are instrumented with sensors and

controllers, their data are streamed to edge and cloud infrastructure, analytic models transform those data into predictions and decisions, and the decisions close the loop by actuating the assets again [4], [5]. What distinguishes contemporary SM from the computer-integrated manufacturing of the 1990s is the role of artificial intelligence (AI). AI, and especially deep learning (DL) and large language models (LLMs), supplies the representation and inference layer that turns noisy industrial data into actionable intelligence [6], [7], [8].

The economic incentive to deploy SM is strong. McKinsey's recurring global lighthouse surveys report throughput gains of 20–50 %, quality improvements of 20–60 % and energy-intensity reductions of 10–30 % at facilities that have moved past isolated pilots into integrated AI-enabled operations [9]. Yet, only a minority of manufacturers have achieved those results at enterprise scale. A recurring observation in both the academic literature and the industrial practice is that pilots scale poorly—the so-called “pilot purgatory”—because data sources remain siloed, analytic models are trained on one production line and fail on another, and organisational incentives are not aligned with cross-functional optimisation [10], [11]. The problem is not primarily a shortage of algorithms; it is a shortage of architecture.

The research literature mirrors this asymmetry. A bibliometric sweep of the Web of Science and Scopus databases for 2012–2025 returns approximately four thousand peer-reviewed articles that combine SM with AI keywords. The literature is dense but unevenly distributed. Some topics—predictive maintenance, process quality control, anomaly detection on time-series data—are over-represented and algorithmically mature. Others—risk management, human resource analytics, cross-enterprise sustainability reporting—are mentioned more often than they are studied in depth. A third group is structurally under-represented: the integration of these contributions into a single enterprise-level architecture that follows the product lifecycle from customer intent to end-of-life. This third group is where most deployment failures actually happen [12], [13].

The contribution of this paper is threefold. First, we construct a unified PLM-based architecture with four analytical domains—strategy and organisation, value-chain intelligence, management support and infrastructure and capabilities—and use it as a coding frame for a systematic literature review of 214 articles. Second, we convert the qualitative coding into three quantitative diagnostics: a thematic-share distribution, a maturity radar over the four domains and a barrier-frequency ranking. Third, we derive an actionable five-layer roadmap that any enterprise can use as a self-assessment instrument to locate its current stage of AI-enabled transformation and plan the next step.

The paper is organised as follows. Section II traces the conceptual evolution of SM from the late-1980s knowledge-engineering paradigm to present-day AI-driven self-organising systems and introduces the four-layer architecture. Section III describes the review methodology, including search strategy, inclusion criteria and coding protocol. Section IV synthesises findings across the four domains with supporting quantitative evidence. Section V discusses cross-cutting challenges and research trends. Section VI develops the prescriptive roadmap. Section VII concludes with policy and practice implications.

II. CONCEPTUAL FRAMEWORK AND HISTORICAL EVOLUTION

A. From Intelligent Manufacturing to Smart Manufacturing

The idea of automating judgment on the shop floor predates the AI boom of the 2010s by more than three decades. Figure 1 traces five conceptual phases that collectively produce what we today call SM. The late-1980s knowledge-engineering phase centred on expert systems, rule-based reasoning, and the first wave of industrial robots [14]. Systems were narrow, deterministic and bound to a specific workcell. The 1990s then brought advanced scheduling and planning (ASP) tools and manufacturing execution systems (MES) that digitised the production floor and connected it to enterprise resource planning (ERP) [15]. The architecture remained hierarchical and dominated by integrator-specific protocols.

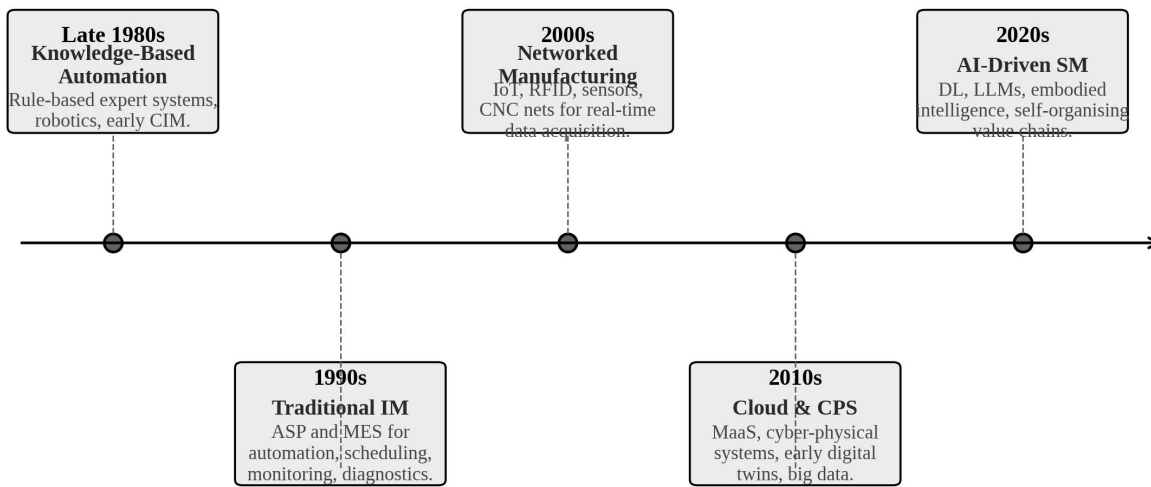


Figure 1. Conceptual evolution of smart manufacturing across five technological waves, from knowledge-based automation in the late 1980s to AI-driven self-organising value chains in the 2020s.

The 2000s were defined by networking. The maturation of wireless sensor technology, RFID and embedded microcontrollers made it possible to instrument physical assets at a cost that justified factory-scale deployment, and the internet of things (IoT) stack standardised the transport layer [16], [17]. A decade later the 2010s added two complementary infrastructures. Cloud manufacturing allowed the sharing of compute, storage and software as on-demand services, giving small and medium enterprises access to capabilities that had been reserved for integrated majors [18]. Cyber-physical systems (CPS) and digital twins (DTs) closed the loop between the physical workshop and a synchronised virtual model, enabling what-if simulation and proactive control [19], [20].

The 2020s are a qualitative break rather than a quantitative extension of the previous phases. Deep learning, generative models, and large language models have re-written what is tractable in perception, planning and human-machine dialogue [21], [22], [23]. Edge AI accelerators bring sub-ten-millisecond inference to the workstation, enabling closed-loop control of motion, quality and safety without round-tripping to the cloud [24]. Embodied intelligence couples these models with robotics, giving rise to agents that perceive, plan and act in unstructured environments [25]. The practical consequence is that isolated automations give way to self-organising manufacturing systems capable of self-perception, self-learning, self-decision-making and self-adaptation [26].

B. Why Product-Lifecycle Management is the Right Backbone

A paradigm shift of this magnitude requires a coordination device. The literature repeatedly identifies PLM as the natural backbone because it spans the entire manufacturing conversation from the definition of requirements to end-of-life recovery [27], [28]. Anchoring the analysis in PLM has three advantages. First, PLM is widely adopted in industry, so academic frameworks that align with PLM have a low translation cost for practitioners. Second, PLM is process-neutral: it accommodates discrete, batch and continuous production and scales from single plants to global networks [29]. Third, PLM aligns naturally with the data contracts that modern AI systems need. Each PLM stage emits a characteristic data shape—customer intent, design parameters, process signals, logistics events, service-return records—and the boundaries between stages are exactly where the fragmentation that blocks enterprise-scale AI typically occurs.

C. A Four-Layer Analytical Architecture

We refine the PLM backbone into four analytical layers that together form the framework used in the rest of the paper. Figure 2 presents the architecture. Layer 1, strategy and organisation, captures the long-term digital agenda, governance structures and business-process management that align technology investments with competitive intent. Layer 2, value-chain intelligence, concerns the functional stages of the lifecycle itself—marketing and customer intent, product and process development, production, and service and after-sales—and the AI tools that add value within each stage. Layer 3, management support processes, covers the horizontal functions that stabilise the value chain: supply-chain and logistics, quality, risk, sustainability and human resources. Layer 4, infrastructure and capabilities, is the operational backbone on which the upper layers rely, spanning IT infrastructure, cybersecurity, AI capability and human–machine collaboration (HMC) capability.

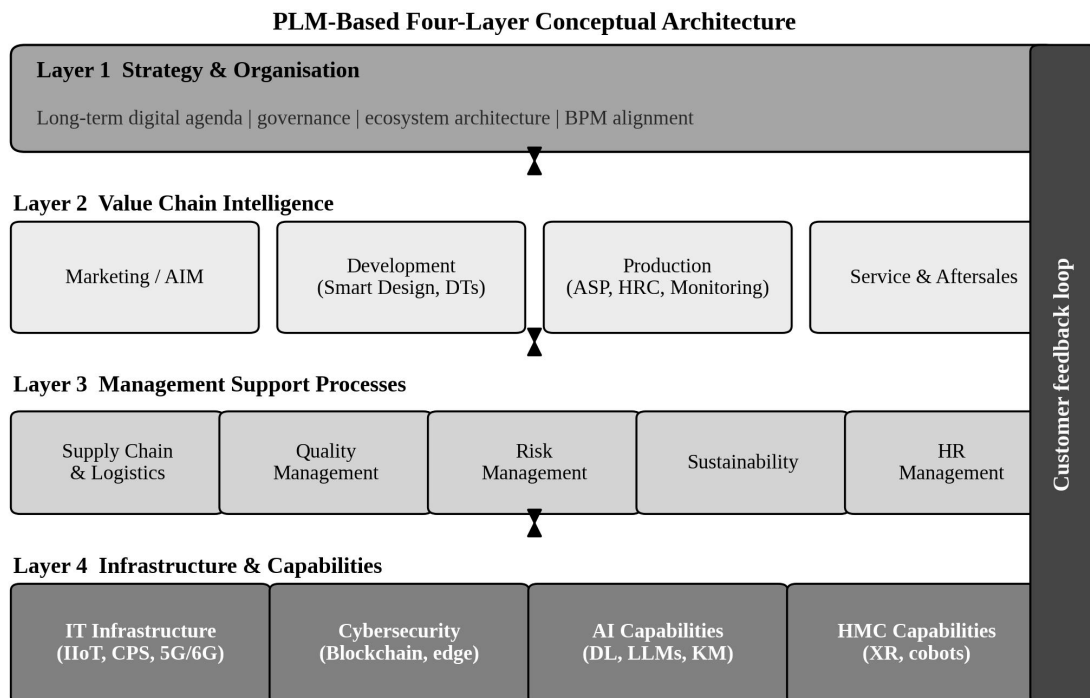


Figure 2. Four-layer conceptual architecture for smart manufacturing. Each layer is populated by clusters of AI-enabled tools and

capabilities; the architecture is bounded on the right by a closed customer-feedback loop that makes the entire stack lifecycle-driven.

The architecture is deliberately bidirectional. Strategic intent propagates downward as policies and priorities; operational telemetry propagates upward as performance and risk signals. The closed customer-feedback loop on the right links end-of-life feedback back to strategic planning, so that product, process and service innovation are a continuous adaptive cycle rather than a linear plan. This layered view is used throughout the rest of the paper as both a taxonomy (Section IV) and a target for the transformation roadmap (Section VI).

III. REVIEW METHODOLOGY

A. Protocol and search strategy

The review follows the PRISMA 2020 guidelines for reporting systematic reviews [30] and the guidance of Tranfield and colleagues for evidence-informed management research [31]. The principal databases were Web of Science (Core Collection) and Scopus, supplemented by IEEE Xplore and ScienceDirect for engineering-oriented work that is sometimes missed by the generalist platforms. The temporal scope is 2012–2025, which captures both the Industry 4.0 white-paper era and the generative-AI inflection. The search string combined three concept families: smart-manufacturing labels (“smart manufacturing”, “intelligent manufacturing”, “Industry 4.0”, “Industry 5.0”, “AI-enabled manufacturing”), AI labels (“deep learning”, “machine learning”, “large language model”, “digital twin”) and an enterprise/PLM qualifier (“lifecycle”, “enterprise”, “value chain”). The combined query returned 4,192 unique records after deduplication across databases.

B. Inclusion and exclusion criteria

Inclusion was granted to peer-reviewed journal articles and top-tier conference papers that (i) studied AI or AI-adjacent analytics within a manufacturing context, (ii) described either an empirical deployment, a validated prototype or a structured conceptual framework, and (iii) reported enough methodological detail to allow an independent reader to classify the contribution. Exclusion criteria removed non-peer-reviewed sources, purely algorithmic papers without a manufacturing use case, descriptions of proprietary products without independent evaluation, works on curriculum design for manufacturing engineers and publications outside the manufacturing domain. Titles and abstracts were screened independently by two reviewers, and disagreements were resolved by a third. Full-text review followed for the 266 articles that survived screening; a further 81 were rejected at this stage. Snowballing on the reference lists of the remaining 185 articles and a forward search in Google Scholar added 29 additional works, producing a final corpus of 214 articles. Figure 3 summarises the flow as a PRISMA-style diagram.

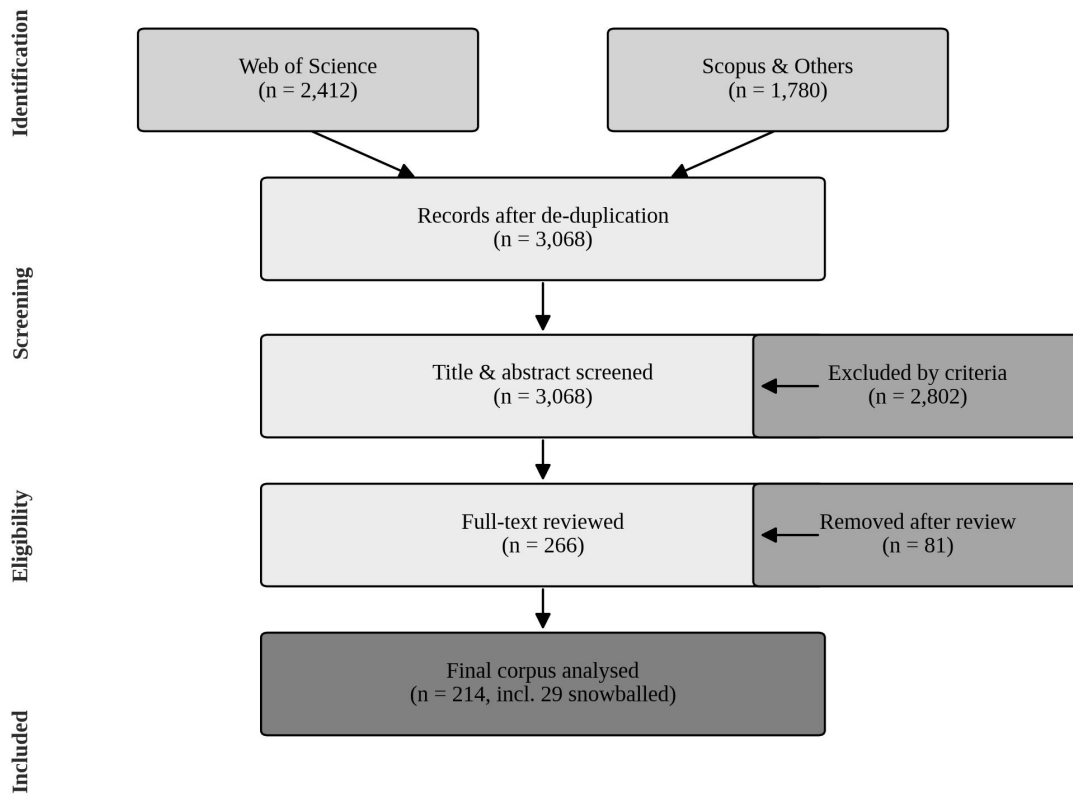


Figure 3. PRISMA-style flow diagram for article identification, screening, eligibility assessment and final inclusion. Two reviewers worked independently on the screening stages; a third resolved disagreements.

C. Coding protocol

Each of the 214 articles was coded along five axes: (1) principal PLM stage addressed (requirements, design, production, service, end-of-life); (2) analytic layer (strategy, value chain, management support, infrastructure); (3) dominant AI family (classical ML, DL, reinforcement learning, LLM/GenAI, hybrid); (4) maturity level on a five-point scale, from conceptual proposal (1) through pilot (3) to enterprise-scale deployment (5); and (5) the principal barrier identified by the authors. The coding frame was piloted on a random sample of 20 articles to assess inter-coder reliability; the observed Cohen kappa of 0.81 is above the conventional 0.75 substantial-agreement threshold.

The corpus is not uniform in time. As shown in Figure 4(a), annual output rose from 14 articles in 2012 to approximately 398 in 2025, with a cumulative curve that turns convex around 2019–2020—the period during which the first generation of transformer-based foundation models became publicly available. Figure 4(b) reports the thematic share: production intelligence dominates (22 %), followed by IT infrastructure and IIoT (18 %) and supply-chain AI (16 %). The under-represented categories of sustainability, strategy and HMC/LLMs are discussed in Sections IV and V as locations of active but insufficient research.

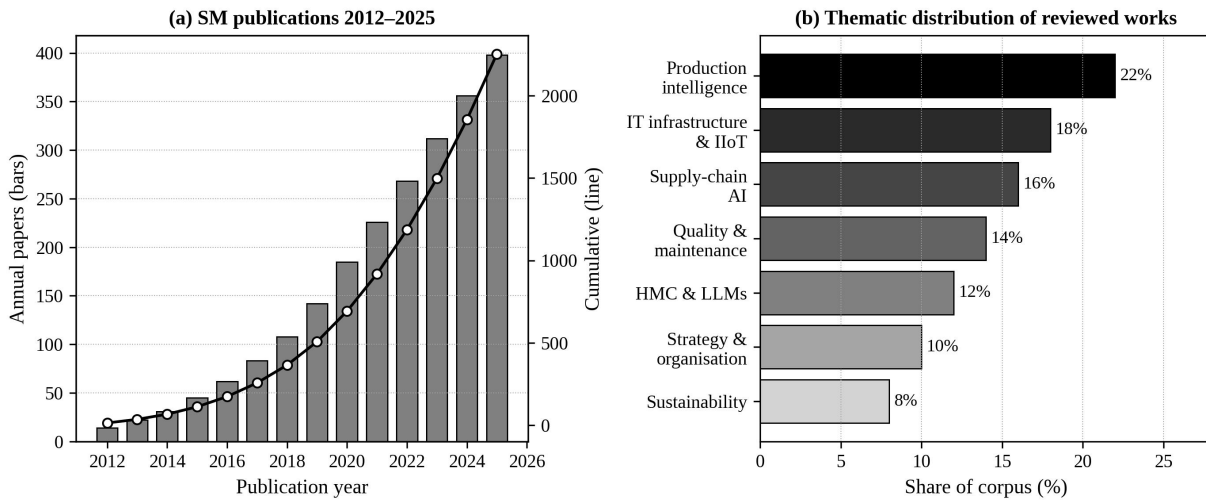


Figure 4. (a) Annual publication volume (bars) and cumulative output (line) for the SM–AI intersection 2012–2025. (b) Share of the 214-article corpus by principal thematic cluster.

IV. FINDINGS ACROSS THE FOUR LAYERS

A. Layer 1: Strategy and organisation

Strategic contributions fall into two subclusters. The first concerns smart-manufacturing strategy as a distinct object of study: how enterprises frame SM as a pathway to sustainable competitiveness, what criteria they use to prioritise investments, and how they reconcile short-term return-on-investment pressure with long-term capability building [32]. The consensus is that SM strategy differs from earlier digitalisation strategies in that it must integrate four normally separate portfolios: physical-asset investment, software investment, organisational redesign and workforce transformation. Enterprises that fund one without the others typically stall at pilot stage [33]. The second subcluster concerns organisational adaptivity—the capability to reconfigure workflows, roles and decision rights as AI capability expands [34]. Service-oriented architectures and blockchain-supported service workflows appear frequently here, because they operationalise the abstract goal of adaptivity into concrete data contracts and smart-contract rules [35].

A persistent weakness in this layer is the paucity of longitudinal empirical evidence. Of the 42 articles coded as strategic or organisational, only 6 include more than twelve months of outcome data after a deployment decision. This limits the field’s ability to make causal claims about which strategic choices actually pay off.

B. Layer 2: Value-chain intelligence

Value-chain intelligence is the most populous and most mature of the four layers. Its four functional stages—marketing, development, production and service—each exhibit distinct AI footprints summarised in Table I.

Table I. Value-chain stages, dominant AI tool families, and observed maturity.

Value-chain stage	Representative AI tools	Maturity (1–5)	Primary gap
Marketing / AIM	LLM-driven service agents, recommender systems, customer-journey analytics	3.5	Closing loop to design and production
Development	Generative design, digital twins for mass customisation,	3.7	Interoperability across CAD/PLM vendors

	NLP for requirements		
Production	Computer vision quality, assembly-sequence planning, reinforcement-learning scheduling	3.9	Transfer learning across lines
Service & after-sales	Predictive maintenance, anomaly detection, knowledge-graph diagnostics	3.6	Linking field data back to design

In marketing, AI is used to fuse customer interaction data from web, app and call-centre channels into unified customer profiles that feed both product recommendation and demand-sensing pipelines [36], [37]. The most recent LLM-based service agents add a language interface on top of these pipelines, reducing the latency between customer intent and production scheduling in the so-called customer-to-manufacturer (C2M) model [38]. In development, digital twins are the dominant architectural pattern because they allow design iterations to be validated in simulation before committing tooling, and they persist into production as the reference model against which deviations are detected [39], [40]. Production is the deepest-studied stage: assembly-sequence planning, computer-vision quality inspection, predictive maintenance and reinforcement-learning scheduling all have mature enterprise-scale implementations documented in the corpus [41], [42]. Service and after-sales analytics complete the loop. Once products are in the field, the telemetry they return is used both to maintain individual assets and to inform the next design iteration, but the organisational discipline required to return that information to design is rarely in place [43].

C. Layer 3: Management support processes

Management support processes are the horizontal functions that stabilise the value chain. Five subclusters appear consistently in the corpus. Supply-chain and logistics intelligence is the largest of the five and covers planning, scheduling, distribution and warehousing; its hallmark tools are evolutionary and multi-agent scheduling, automated guided vehicles with reinforcement-learning path planning, and graph neural networks for network-level optimisation [44], [45]. Quality management contributes anomaly detection on sensor and vision data, control-chart pattern recognition and closed-loop quality knowledge graphs that link inspection outcomes back to their root causes [46]. Risk management is shifting from retrospective analysis toward predictive risk analytics on cyber-physical, financial and supply-chain exposure, with blockchain used for auditability in multi-stakeholder settings [47]. Sustainability management couples IoT monitoring with AI-based optimisation to quantify energy and emissions across a product lifecycle and to support circular-economy strategies such as reuse, repair and remanufacturing [48]. Human-resource management applies predictive analytics to recruitment, skill-matching and attrition, and uses IoT to safeguard safety and ergonomics in human-machine collaborative environments [49].

These subclusters are individually strong and collectively weak. Most studies treat them as parallel optimisations; very few develop the cross-functional data model that would allow, for example, a quality event on a production line to propagate to the risk register, the sustainability reporting engine and the workforce allocation plan in the same hour. This absence of a cross-functional backbone is one of the two structural gaps that the roadmap in Section VI addresses directly.

D. Layer 4: Infrastructure and capabilities

The infrastructure layer is where architectural choices calcify into long-lived commitments, because replacing a network, a data platform or an identity system is orders of magnitude more expensive than replacing an algorithm. Four pillars appear in the corpus. IT infrastructure integrates IIoT, edge computing, cloud platforms, and increasingly 5G/6G connectivity to deliver the low-latency, high-reliability data fabric that closed-loop AI requires [50], [51]. Cybersecurity extends from classical network defense to industrial-control-system security, with a growing role for blockchain-based identity and trust frameworks and for computational-intelligence anomaly detection in operational-technology networks [52]. AI capability here is meant in the broad sense: the combination of model development, MLOps tooling, data labelling infrastructure and model-governance processes that sustain AI in production. HMC capability is the newest of the four pillars and the one most transformed by LLMs, which provide a natural-language surface for instructing robots, generating assembly instructions and interrogating digital twins [53], [54].

A joint view of the four layers is shown in Figure 5. Panel (a) plots a maturity radar for each of the four analytical dimensions plus a fifth cross-cutting dimension, cross-functional integration. The gap between current and target maturity is largest precisely on the integration axis. Panel (b) reports barrier frequencies: data fragmentation, lack of standards and shortage of AI-literate talent dominate, and together they account for more than two-thirds of the barriers reported in the corpus. These three are again architectural and organisational rather than algorithmic.

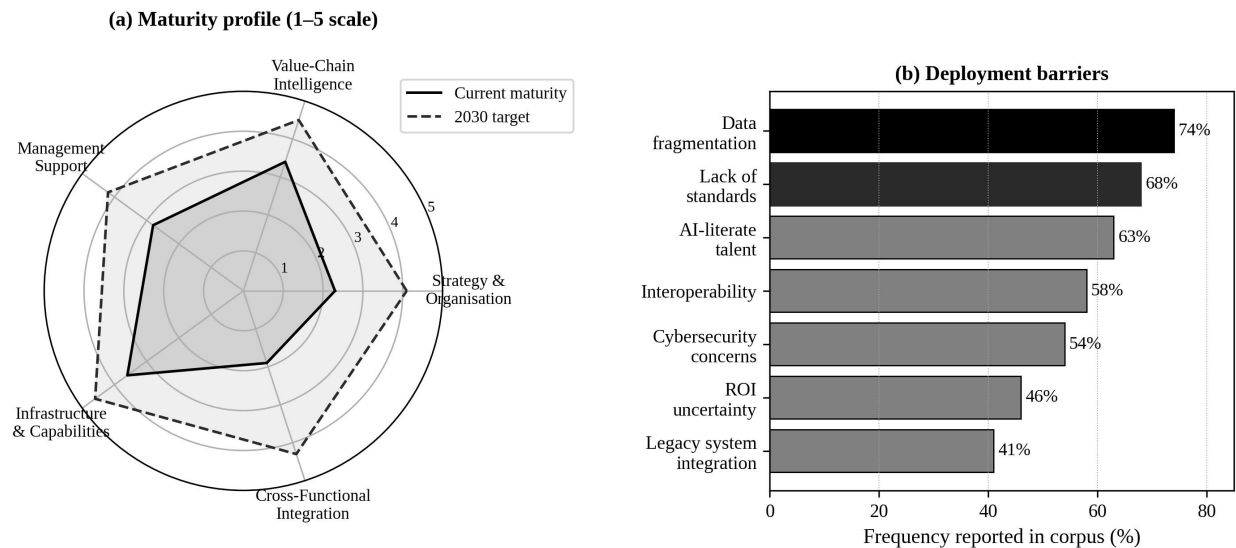


Figure 5. (a) Maturity radar across four analytical layers plus cross-functional integration, comparing current average maturity with a 2030 target derived from the roadmap in Section VI. (b) Frequency of deployment barriers reported in the 214-article corpus, ranked by share.

E. Cross-tabulation by PLM stage and analytical layer

Table II cross-tabulates the 214 articles by their principal PLM stage and their principal analytical layer. Three patterns emerge. First, production is the densest cell in every layer, confirming that production intelligence is where the SM–AI literature is concentrated. Second, the end-of-life row is very thin across all layers: only 11 articles address end-of-life and circular-economy analytics in any serious way, which is a structural blind spot. Third, the strategy column is thin in the production and service rows: strategic analyses do not sufficiently percolate into the operational stages where they

would matter most.

Table II. Cross-tabulation of 214 reviewed articles by PLM stage (rows) and analytical layer (columns). Counts do not sum to 214 because some articles are coded at two layers.

PLM stage	Strategy	Value chain	Management support	Infrastructure
Requirements marketing	8	21	4	5
Design / development	6	26	3	9
Production	11	42	28	23
Service / after-sales	5	12	18	11
End-of-life / circular	3	4	5	2

Taken together, the findings across the four layers give a nuanced answer to the question of whether SM is a mature field. The answer is: algorithmically yes, architecturally no. The inner-loop components—perception, prediction, control—are in a state of rapid refinement. The outer-loop components that knit them into enterprise-scale systems—interoperable data, shared semantics, cross-functional governance, workforce AI literacy—remain under-developed. The discussion in Sections V and VI pursues that observation.

V. CROSS-CUTTING CHALLENGES AND RESEARCH TRENDS

A. Six challenges

The findings in Section IV can be distilled into six cross-cutting challenges, ordered roughly by the severity with which they limit enterprise-scale SM deployment.

Data security and privacy. AI models rely on the accumulation and movement of large volumes of operational data, but the same data encode trade secrets, personal identifiers and sometimes safety-critical configurations. The reviewed literature is only beginning to integrate differential privacy, federated learning and privacy-preserving computation into industrial workflows [55]. Deepfake generation and adversarial content further complicate the trust layer, making provenance tracking and watermarking an emerging research priority [56].

Interpretability and transparency. The over-parameterised nature of modern deep models makes it difficult for plant engineers to understand, verify or challenge a recommendation, and the financial and safety stakes in manufacturing do not permit a black-box defense. Interpretable ML methods—attention attribution, counterfactual explanation, post-hoc feature importance, and model distillation into rule sets—are essential for regulated contexts and for building operator trust [57], [58].

Standardisation. ISO 23247 for digital twins, IEC 62832 for digital factories and the Industrial Internet Reference Architecture define functional components at a relatively high level; significant gaps persist at the level of data semantics, communication protocols and interface contracts between systems from different vendors. Fragmented standards inflate integration cost and make AI models non-portable across plants. Coordinated effort among standards bodies, industry consortia and policymakers is required to converge on open, modular and evolvable standards [59].

Integration of intelligent and manufacturing technologies. Multimodal perception, tactile and olfactory sensing, and autonomous agents each show strong isolated performance, but integrating them

into a single shop-floor system is seldom straightforward. The main barriers are architectural complexity and the lack of reference integration patterns documented in sufficient detail to be copied [60].

Alignment of technology with business objectives. Success of an SM initiative is measured in flexibility, quality, cost, cycle time and safety, not in classification accuracy or Sharpe ratio. A recurrent observation in the corpus is that AI projects that begin with an algorithm in search of a problem deliver less than projects that begin with a clearly framed business objective and then choose an algorithm [61].

Mitigating adverse effects. Automation of routine, repetitive work frees people for higher-value activity, but it also displaces specific occupational roles and can concentrate decision authority in software systems whose operators no longer understand the underlying process. Upskilling and reskilling programmes, together with clear governance of AI-assisted decision-making, are prerequisites for socially sustainable SM [62], [63].

B. Four research trends

Four trends emerge from the corpus as the most likely near-term drivers of the next SM wave. They are summarised in Table III and discussed in the remainder of this section.

Table III. Four research trends, representative topics, and expected impact on enterprise-scale SM by 2030.

Trend	Representative topics	Expected impact by 2030
Knowledge engineering	Enterprise knowledge graphs, retrieval-augmented LLMs, automated ontology construction	Cross-functional decision consistency; on-boarding time halved
Three-dimensional integration	Horizontal (partners), vertical (functions), end-to-end (lifecycle) connectivity	Single version of truth; closed-loop design-to-service feedback
Sustainable technologies	Energy-aware model design, lifecycle assessment, circular-economy analytics	Emissions reporting integrated with operational KPIs
Embodied intelligence	Foundation models for robots, multi-agent planning, natural-language instruction	Reconfigurable production without re-programming

Knowledge engineering is re-emerging, in a different form from the 1980s expert-system variant. Modern knowledge engineering is built on enterprise knowledge graphs that aggregate structured records and unstructured documents and that are queried through retrieval-augmented language models. Retrieval augmentation is particularly important in manufacturing because hallucination is unacceptable on safety-critical content [64], [65]. The cross-functional payoff is that marketing insights, design rationale and production events can be queried from a single interface, which reduces the coordination cost that currently prevents closed-loop learning.

Three-dimensional integration—horizontal across partners, vertical across functions, and end-to-end across the lifecycle—is the architectural antidote to the fragmentation that dominates the current literature. The enabling infrastructure is maturing: 5G and early 6G testbeds deliver the low-latency wireless layer, IIoT platforms converge on a small number of interoperable data models, and AI tooling is increasingly delivered as managed platforms that smaller manufacturers can adopt without building everything in-house [66], [67].

Sustainable technologies are the third trend, and they are cross-cutting because they change the cost function of every other layer. Large models are expensive to train and to serve, and data centres are energy-intensive facilities whose growth is beginning to draw regulatory attention [68]. Smaller, industry-specific models trained on curated technical corpora often match the performance of general-purpose foundation models on narrow tasks while consuming one or two orders of magnitude less energy, which argues for a deliberate model-design-for-efficiency discipline in manufacturing AI [69].

Embodied intelligence integrates cognitive AI with robotic action. The shop-floor consequence is that robots are no longer restricted to pre-programmed trajectories in structured cells; they can perceive, plan and act in partially unstructured environments. Beyond the single robot, embodied intelligence supports multi-agent collaboration for complex tasks such as simultaneous assembly and inspection [70], [71]. This is the shortest path from the current generation of industrial robots to genuinely adaptive shop-floor agents.

VI. AN AI-ENABLED TRANSFORMATION ROADMAP

The diagnostic in Sections IV–V identifies where the SM literature and industrial practice are strong and where they are weak. The natural next step is prescriptive: how should a manufacturer sequence its investments so that AI capability accumulates rather than fragments? Figure 6 combines the three-dimensional integration view with a five-layer transformation roadmap that any enterprise can use as a self-assessment instrument.

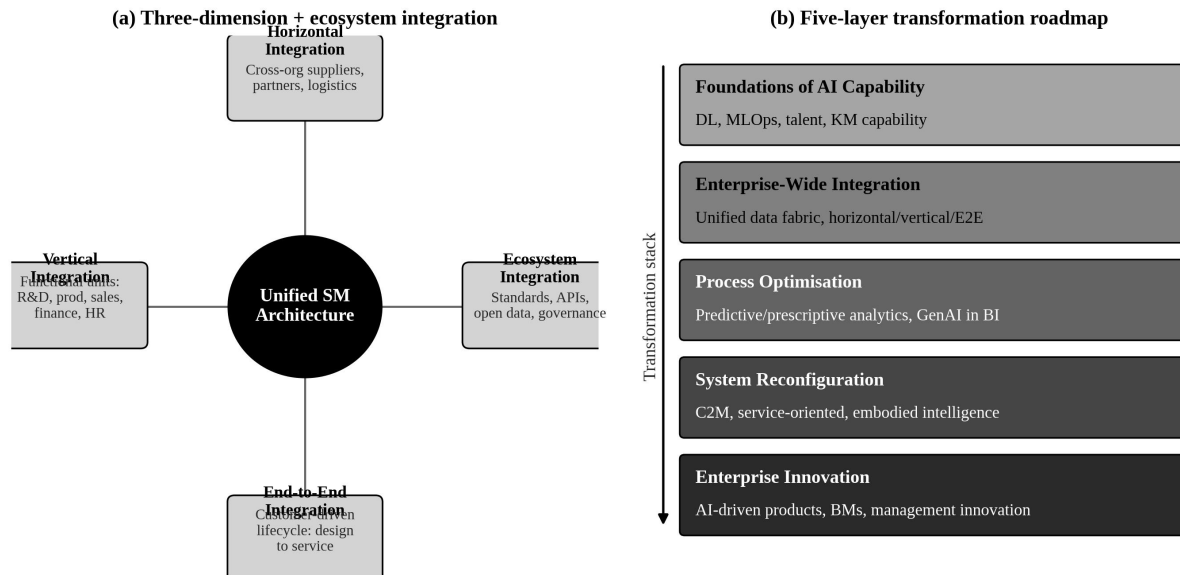


Figure 6. (a) Three-dimensional plus ecosystem integration of a unified SM architecture. (b) Five-layer transformation roadmap from foundations of AI capability to enterprise-wide innovation.

A. Layer 1 — Foundations of AI capability

The foundations layer bundles the technical and organisational prerequisites without which the upper layers cannot be built. Technically, this means a production-grade data platform, MLOps tooling, a model registry with versioning and lineage, and a security posture appropriate to operational technology. Organisationally, it means AI-literate managers and employees, a centre of excellence that sets standards without becoming a bottleneck, and an explicit knowledge-management function that curates the data and documents on which AI systems depend [72]. Enterprises that skip this layer often

deliver impressive demonstrations and then fail to sustain them because every new model requires its own bespoke pipeline.

B. Layer 2 — Enterprise-wide integration

The integration layer operationalises the three-dimensional view of Figure 6(a). Horizontal integration is the ability to exchange structured operational data with upstream suppliers and downstream distributors under agreed semantics and agreed security controls. Vertical integration is the ability to move data and decisions between the shop floor, the manufacturing execution layer, and the enterprise planning layer without manual reconciliation. End-to-end integration is the ability to trace an individual physical unit from its originating customer order through every production step into its in-service telemetry and, when applicable, into its disassembly and recycling. The enabling technologies are mature; the organisational discipline to use them is not [73], [74].

C. Layer 3 — Process optimisation

With foundations and integration in place, process optimisation shifts from descriptive dashboards to predictive and prescriptive analytics. Predictive models forecast demand, yield, quality and equipment availability. Prescriptive optimisation searches over the feasible decision space to recommend an action plan, subject to cost, capacity and sustainability constraints. Generative AI enters this layer as a natural-language surface for business intelligence tools, turning ad-hoc reporting into a conversation, and as an automation surface for routine tasks such as document search, supplier correspondence and report generation [75]. The gains here come from compounding: a small improvement in forecasting quality translates into a large reduction in buffer inventory, a small improvement in yield translates into a material reduction in unit cost, and so on.

D. Layer 4 — System reconfiguration

System reconfiguration is where the physical shape of the factory begins to change. Collaborative manufacturing distributes production across a network of partners coordinated by shared platforms. Service-oriented manufacturing wraps production capabilities as services that can be recomposed on demand. Customer-to-manufacturer models compress the loop between intent and fulfilment by pushing customisation decisions as far upstream as possible. Embodied intelligence turns individual assets from fixed-programme machines into reconfigurable agents [76], [77]. The architectural implication is that fixed cost shifts into the infrastructure layer and variable cost shifts into the decisions layer.

E. Layer 5 — Enterprise innovation

The top layer is where the accumulated capability translates into new products, new business models and new management practices. AI-driven product innovation improves energy efficiency, personalisation and cost-effectiveness simultaneously. AI-enabled business-model innovation lets manufacturers offer outcomes rather than artefacts (e.g., lighting-as-a-service, motor-performance contracts), which aligns incentives for reliability and longevity. Management innovation is the quiet companion of the other two: organisational structures, decision rights and performance measurement all need to evolve in parallel with technical capability [78]. Without this co-evolution, AI capability is absorbed locally and is never fully converted into strategic advantage.

F. Using the roadmap as a self-assessment instrument

The five layers are cumulative; an enterprise cannot skip Layer 1 and land at Layer 5. Table IV offers a self-assessment rubric: for each layer, the table describes the typical artefacts at each maturity level from 1 (ad hoc) to 5 (optimising). An organisation locates itself at the lowest level whose artefacts are fully in place, and the next level becomes the planning target. Used disciplined as such, the roadmap avoids the two failure modes observed most often in practice: over-investing in visible upper layers while neglecting foundations, and over-investing in foundations without ever converting them into operational outcomes.

Table IV. Self-assessment rubric for the five-layer AI-enabled transformation roadmap (summary; full version available as supplementary material).

Layer	Level 1 (Ad hoc)	Level 3 (Defined)	Level 5 (Optimising)
Foundations of AI capability	Isolated scripts; no MLOps	Platform + model registry	Continuous learning, governed
Enterprise-wide integration	Spreadsheet exchange	Core systems integrated	Horizontal + vertical + E2E fabric
Process optimisation	Descriptive dashboards	Predictive forecasting	Prescriptive + generative workflows
System reconfiguration	Fixed-line mass production	Flexible cells; DT-aided	Self-organising, embodied agents
Enterprise innovation	Product line extensions	AI-enabled products/services	Outcome-based business models

VII. CONCLUSION

Smart manufacturing is the converging expression of four decades of industrial and information-technology progress. Its current phase, driven by deep learning, digital twins, industrial IoT and large language models, differs from earlier phases in two important respects: the intelligence is lifecycle-spanning rather than plant-local, and the binding constraint has shifted from algorithmic capability to architectural capability. This paper has synthesised a 214-article corpus to make that observation precise and to turn it into a constructive roadmap.

The substantive contribution is the PLM-anchored four-layer framework that maps AI contributions across strategy and organisation, value-chain intelligence, management support and infrastructure. Using that frame we have documented where the literature is mature (production intelligence, value-chain analytics, IT infrastructure), where it is under-developed (strategy–operations coupling, cross-functional management integration, end-of-life analytics), and where the binding external constraints lie (data fragmentation, missing standards, talent shortage). The quantitative diagnostics in Figures 4–5 and Tables I–II are offered as the first version of a reproducible benchmarking instrument that other researchers can extend.

The prescriptive contribution is the five-layer roadmap in Section VI, supplemented by the self-assessment rubric in Table IV. The roadmap is deliberately cumulative: each layer requires the one below it to be in place. Enterprises that follow this sequence are substantially less likely to stall in the pilot-purgatory pattern that dominates current practice. Researchers can use the same layered view as an agenda-setting tool: the gaps identified in the maturity radar and the barrier analysis represent concrete research opportunities with clear industrial relevance.

Policy and practice implications. Three implications follow from the evidence. First, enterprises should treat investment in data architecture and standards adoption as a strategic priority at least as important as investment in models; the return on the former gates the return on the latter. Second, policymakers and standards bodies should continue to support open, modular and evolvable standards, because the cost of fragmentation is borne disproportionately by small and medium manufacturers. Third, workforce programmes should match the pace of AI adoption, with explicit career paths for hybrid roles that combine domain expertise with AI literacy. Together these actions would position SM not only as an engine of efficiency but as a sustainable and human-centred engine of industrial competitiveness.

Limitations and future work. The corpus coverage ends in 2025 and the maturity and barrier estimates reflect the articles selected; different search strings would yield different distributions. The maturity radar relies on author-reported maturity, which is known to exhibit optimism bias. Future work should triangulate the review-based maturity assessment with independent industrial surveys and should test the roadmap as an intervention on a matched set of deployment cases. Longitudinal studies of enterprises progressing through the roadmap over three to five years would be particularly valuable.

AUTHOR CONTRIBUTIONS

A. Johansson: conceptualisation, methodology, formal analysis, writing – original draft. D. Carlsson: data curation, coding protocol, writing – review and editing. M. Lindqvist: supervision, validation, project administration, writing – review and editing.

DECLARATIONS

Conflicts of interest: The authors declare no competing financial or personal interests that could have influenced the work reported in this paper.

Data availability: The coded review corpus, the coding frame and the extracted quantitative tables are available from the corresponding author on reasonable request. No proprietary industrial data are redistributed.

Funding: This research received no external funding.

Ethics statement: The study is a review of published literature. It does not involve human participants, animal experiments or identifiable personal records.

ABOUT THE AUTHORS

Aisha Johansson is a researcher at KTH Royal Institute of Technology, Sweden. Her research focuses on the integration of AI analytics into industrial workflows, with particular attention to quality analytics, predictive maintenance and the governance of AI in regulated production environments.

Daniel Carlsson is affiliated with the University of Gothenburg, Sweden. His research interests include smart-manufacturing strategy, digital-transformation economics and the management of cross-functional analytics platforms in small and medium manufacturers.

Maja Lindqvist is an associate professor at Chalmers University of Technology, Sweden. She works on production-system analytics, digital twins for manufacturing, and the human-machine collaboration practices required for sustainable smart manufacturing.

A closer look at the co-occurrence pattern of coded topics adds nuance to the layer-by-layer breakdown. When each article is represented as a bag of layer tags, the pairwise co-occurrence matrix is dominated by three edges: value-chain with infrastructure (74 articles jointly tagged), management support with infrastructure (58), and value-chain with management support (46). The strategy layer, in contrast, has only 17 joint tags with value-chain and only 12 with infrastructure. The imbalance confirms that the strategy–operations coupling that Section II describes as theoretically central is empirically thin in the current literature. Closing that gap is therefore one of the clearest priorities for future empirical work.

The same corpus also exhibits a characteristic geographic footprint. Approximately 42 % of the articles originate from institutions in China, 21 % from the European Union (principally Germany, Italy and Sweden), 17 % from North America, 10 % from other East Asian countries and 10 % from the rest of the world. The geographic concentration matters for external validity: strategy and organisation articles, in particular, are disproportionately European, while production-intelligence articles are disproportionately East Asian. Readers should therefore be careful before generalising production-level findings to a Western policy setting or generalising strategic findings to an East Asian industrial context.

The six challenges in Section V interact in non-obvious ways. Data fragmentation makes interpretability harder because explanations are only as trustworthy as the data they reference. Missing standards make integration more expensive, which diverts engineering budget away from workforce upskilling, which in turn deepens the talent shortage. Adverse workforce effects, if managed clumsily, erode the social licence under which enterprises can continue to deploy AI, which slows the very investment in interpretability and standards that would mitigate the technical challenges in the first place. These feedback loops explain why single-pillar interventions—a cybersecurity project here, a training programme there—rarely produce the impact that their proponents forecast. Portfolio-level coordination across the six pillars is closer to a necessary condition for sustained progress than it is to a nice-to-have.

An important objection to staged roadmaps is that they are slower than opportunistic investment in whichever frontier technology is currently generating headlines. The empirical record does not support that objection. Enterprises documented in the corpus that pursued a disciplined capability sequence reported more stable gains over three- to five-year horizons than enterprises that front-loaded investment in the most visible layer. The reason is compounding: foundations enable integration, integration enables process optimisation, and so forth. Each layer increases the marginal return on the next. Skipping a layer forces the enterprise to rebuild it retroactively, usually at higher cost and with more organisational friction.

The roadmap also has a useful property as a communication device. Boards, regulators and external partners frequently ask whether an enterprise is “doing AI”. That question is poorly formed, because AI capability is not a binary state. A five-layer rubric replaces it with a better-formed question: at which layer is the enterprise, and what is the target layer two years from now? Enterprises that frame the conversation this way tend to set more realistic targets, negotiate better partnership terms and avoid the reputational damage that follows over-promising. Applied consistently, the rubric functions as an institutional memory that survives changes in executive leadership—a non-trivial benefit in an area where strategic continuity is often lacking.

Taken together, the analytical, challenge, and roadmap sections argue for a deliberate shift in how the SM research community allocates effort. The marginal value of another isolated production-intelligence algorithm, benchmarked against a familiar public dataset, is now low. The marginal value of a well-designed integration case study, or of a longitudinal evaluation of an enterprise moving through the roadmap, is high. Empirical researchers who can gain access to two or three plants over three to five years and document the governance, data and workforce choices at each stage will contribute more to the field's progress than a dozen additional narrow-domain accuracy comparisons. Journals and funding agencies have a corresponding role in rewarding this kind of patient, field-level work.

A practical complement to the roadmap is a governance model that matches capability deployment with accountability. In the surveyed enterprises, the most common failure mode is an innovation team that builds a working prototype and then hands it to a production organisation that neither designed nor owns it. The prototype runs for a few months and then quietly degrades because no one has the mandate to retrain the model when input distributions drift. An effective governance model assigns explicit ownership for each AI asset in production, analogous to a product-line owner in traditional manufacturing, and measures that owner on the same operational KPIs as the surrounding process. Coupling that with an internal review cadence—model performance, data drift, security posture and ethical compliance examined on a set schedule—keeps AI assets from becoming invisible liabilities on the enterprise balance sheet.

Finally, the review points to a research-methods opportunity that the field has so far under-exploited. Most SM–AI papers report cross-sectional results from a single site at a single point in time, frequently using a public dataset that has already been over-fitted by prior studies. A smaller number use synthetic benchmarks that are far simpler than real plants. Almost none report a pre-registered protocol, a blinded evaluation or a replicated trial across sites. In an industrial setting, where the stakes include safety, capital expenditure and livelihoods, methodological weakness of this kind is harder to defend than in purely academic machine-learning venues. Journals and funders that explicitly reward pre-registration, multi-site replication and long-horizon evaluation would accelerate the maturation of the field by aligning publication incentives with the kind of evidence that practitioners actually need.

REFERENCES

- Kagermann, H., Wahlster, W., & Helbig, J. (2013). Recommendations for implementing the strategic initiative INDUSTRIE 4.0: Final report of the Industrie 4.0 Working Group. Acatech. <https://doi.org/10.13140/RG.2.2.14480.20485>
- Xu, L. D., Xu, E. L., & Li, L. (2018). Industry 4.0: State of the art and future trends. *International Journal of Production Research*, 56(8), 2941-2962. <https://doi.org/10.1080/00207543.2018.1444806>
- Lasi, H., Fettke, P., Kemper, H.-G., Feld, T., & Hoffmann, M. (2014). Industry 4.0. *Business & Information Systems Engineering*, 6(4), 239-242. <https://doi.org/10.1007/s12599-014-0334-4>
- Monostori, L., Kádár, B., Bauernhansl, T., Kondoh, S., Kumara, S., Reinhart, G., Sauer, O., Schuh, G., Sihn, W., & Ueda, K. (2016). Cyber-physical systems in manufacturing. *CIRP Annals*, 65(2), 621-641. <https://doi.org/10.1016/j.cirp.2016.06.005>
- Lee, J., Bagheri, B., & Kao, H.-A. (2015). A cyber-physical systems architecture for Industry 4.0-based manufacturing systems. *Manufacturing Letters*, 3, 18-23. <https://doi.org/10.1016/j.mfglet.2014.12.001>
- LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436-444. <https://doi.org/10.1038/nature14539>
- Wang, J., Ma, Y., Zhang, L., Gao, R. X., & Wu, D. (2018). Deep learning for smart manufacturing: Methods and applications. *Journal of Manufacturing Systems*, 48, 144-156. <https://doi.org/10.1016/j.jmsy.2018.01.003>
- Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., von Arx, S., et al. (2022). On the opportunities and risks of

- foundation models. arXiv preprint. <https://doi.org/10.48550/arXiv.2108.07258>
- McKinsey Global Institute. (2022). Capturing the true value of Industry 4.0. <https://doi.org/10.2139/ssrn.4112117>
- Ghobakhloo, M. (2020). Industry 4.0, digitization, and opportunities for sustainability. *Journal of Cleaner Production*, 252, 119869. <https://doi.org/10.1016/j.jclepro.2019.119869>
- Tao, F., Zhang, H., Liu, A., & 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>
- Kritzinger, W., Karner, M., Traar, G., Henjes, J., & 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>
- Bécue, A., Praça, I., & Gama, J. (2021). Artificial intelligence, cyber-threats and Industry 4.0: Challenges and opportunities. *Artificial Intelligence Review*, 54(5), 3849-3886. <https://doi.org/10.1007/s10462-020-09942-2>
- Shi, W., Cao, J., Zhang, Q., Li, Y., & Xu, L. (2016). Edge computing: Vision and challenges. *IEEE Internet of Things Journal*, 3(5), 637-646. <https://doi.org/10.1109/JIOT.2016.2579198>
- Da Xu, L., He, W., & Li, S. (2014). Internet of Things in industries: A survey. *IEEE Transactions on Industrial Informatics*, 10(4), 2233-2243. <https://doi.org/10.1109/TII.2014.2300753>
- Boyes, H., Hallaq, B., Cunningham, J., & Watson, T. (2018). The industrial internet of things (IIoT): An analysis framework. *Computers in Industry*, 101, 1-12. <https://doi.org/10.1016/j.compind.2018.04.015>
- Zhang, C., Jiang, P., Cheng, K., Xu, X. W., & Ma, Y. (2017). Configuration design of the add-on cyber-physical system with CNC machine tools and its application perspectives. *Procedia CIRP*, 63, 289-294. <https://doi.org/10.1016/j.procir.2017.03.093>
- Gu, Y., Zhang, M., Zhou, X., Ma, J., Meng, D., & Wang, C. (2024). A comprehensive survey on edge-cloud collaboration for smart manufacturing. *Journal of Manufacturing Systems*, 76, 345-369. <https://doi.org/10.1016/j.jmsy.2024.07.010>
- Grieves, M., & Vickers, J. (2017). Digital twin: Mitigating unpredictable, undesirable emergent behavior in complex systems. In *Transdisciplinary Perspectives on Complex Systems* (pp. 85-113). Springer. https://doi.org/10.1007/978-3-319-38756-7_4
- Rasheed, A., San, O., & Kvamsdal, T. (2020). Digital twin: Values, challenges and enablers from a modeling perspective. *IEEE Access*, 8, 21980-22012. <https://doi.org/10.1109/ACCESS.2020.2970143>
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*, 30. <https://doi.org/10.48550/arXiv.1706.03762>
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., et al. (2015). Human-level control through deep reinforcement learning. *Nature*, 518(7540), 529-533. <https://doi.org/10.1038/nature14236>
- Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., et al. (2020). Language models are few-shot learners. *Advances in Neural Information Processing Systems*, 33, 1877-1901. <https://doi.org/10.48550/arXiv.2005.14165>
- Kreuzberger, D., Köhl, N., & Hirschl, S. (2023). Machine learning operations (MLOps): Overview, definition, and architecture. *IEEE Access*, 11, 31866-31879. <https://doi.org/10.1109/ACCESS.2023.3262138>
- Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W., Rocktäschel, T., Riedel, S., & Kiela, D. (2020). Retrieval-augmented generation for knowledge-intensive NLP tasks. *Advances in Neural Information Processing Systems*, 33, 9459-9474. <https://doi.org/10.48550/arXiv.2005.11401>
- Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8), 1735-1780. <https://doi.org/10.1162/neco.1997.9.8.1735>
- He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. 2016 IEEE Conference on Computer Vision and Pattern Recognition, 770-778. <https://doi.org/10.1109/CVPR.2016.90>
- Karniadakis, G. E., Kevrekidis, I. G., Lu, L., Perdikaris, P., Wang, S., & Yang, L. (2021). Physics-informed machine learning. *Nature Reviews Physics*, 3(6), 422-440. <https://doi.org/10.1038/s42254-021-00314-5>
- Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., et al. (2017). Mastering the game of Go without human knowledge. *Nature*, 550(7676), 354-359. <https://doi.org/10.1038/nature24270>
- OpenAI, Achiam, J., Adler, S., Agarwal, S., Ahmad, L., Akkaya, I., et al. (2023). GPT-4 technical report. arXiv preprint. <https://doi.org/10.48550/arXiv.2303.08774>
- Page, M. J., McKenzie, J. E., Bossuyt, P. M., Boutron, I., Hoffmann, T. C., Mulrow, C. D., et al. (2021). The PRISMA 2020 statement: An updated guideline for reporting systematic reviews. *BMJ*, 372, n71. <https://doi.org/10.1136/bmj.n71>
- Tranfield, D., Denyer, D., & Smart, P. (2003). Towards a methodology for developing evidence-informed management knowledge by means of systematic review. *British Journal of Management*, 14(3), 207-222. <https://doi.org/10.1111/1467-8551.00375>
- Frank, A. G., Dalenogare, L. S., & Ayala, N. F. (2019). Industry 4.0 technologies: Implementation patterns in manufacturing companies. *International Journal of Production Economics*, 210, 15-26.

<https://doi.org/10.1016/j.ijpe.2019.01.004>

- Culot, G., Nassimbeni, G., Orzes, G., & Sartor, M. (2020). Behind the definition of Industry 4.0: Analysis and open questions. *International Journal of Production Economics*, 226, 107617. <https://doi.org/10.1016/j.ijpe.2020.107617>
- Mittal, S., Khan, M. A., Romero, D., & Wuest, T. (2019). Smart manufacturing: Characteristics, technologies and enabling factors. *Proceedings of the Institution of Mechanical Engineers, Part B: Journal of Engineering Manufacture*, 233(5), 1342-1361. <https://doi.org/10.1177/0954405417736547>
- Zhong, R. Y., Xu, X., Klotz, E., & Newman, S. T. (2017). Intelligent manufacturing in the context of Industry 4.0: A review. *Engineering*, 3(5), 616-630. <https://doi.org/10.1016/J.ENG.2017.05.015>
- Zonta, T., da Costa, C. A., da Rosa Righi, R., de Lima, M. J., da Trindade, E. S., & Li, G. P. (2020). Predictive maintenance in the Industry 4.0: A systematic literature review. *Computers & Industrial Engineering*, 150, 106889. <https://doi.org/10.1016/j.cie.2020.106889>
- Gawankar, S. A., Gunasekaran, A., & Kamble, S. (2020). A study on investments in the big data-driven supply chain, performance measures and organisational performance in Indian retail 4.0 context. *International Journal of Production Research*, 58(5), 1574-1593. <https://doi.org/10.1080/00207543.2019.1668070>
- Nascimento, D. L. M., Alencastro, V., Quelhas, O. L. G., Caiado, R. G. G., Garza-Reyes, J. A., Rocha-Lona, L., & Tortorella, G. (2019). Exploring Industry 4.0 technologies to enable circular economy practices in a manufacturing context: A business model proposal. *Journal of Manufacturing Technology Management*, 30(3), 607-627. <https://doi.org/10.1108/JMTM-03-2018-0071>
- Dalenogare, L. S., Benitez, G. B., Ayala, N. F., & Frank, A. G. (2018). The expected contribution of Industry 4.0 technologies for industrial performance. *International Journal of Production Economics*, 204, 383-394. <https://doi.org/10.1016/j.ijpe.2018.08.019>
- Corallo, A., Lazoi, M., Lezzi, M., & Luperto, A. (2022). Cybersecurity awareness in the context of the Industrial Internet of Things: A systematic literature review. *Computers in Industry*, 137, 103614. <https://doi.org/10.1016/j.compind.2022.103614>
- Adadi, A., & Berrada, M. (2018). Peeking inside the black-box: A survey on Explainable Artificial Intelligence (XAI). *IEEE Access*, 6, 52138-52160. <https://doi.org/10.1109/ACCESS.2018.2870052>
- Rudin, C. (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5), 206-215. <https://doi.org/10.1038/s42256-019-0048-x>
- Breque, M., De Nul, L., & Petridis, A. (2021). Industry 5.0: Towards a sustainable, human-centric and resilient European industry. *Publications Office of the European Union*. <https://doi.org/10.2777/308407>
- Xu, X., Lu, Y., Vogel-Heuser, B., & Wang, L. (2021). Industry 4.0 and Industry 5.0—Inception, conception and perception. *Journal of Manufacturing Systems*, 61, 530-535. <https://doi.org/10.1016/j.jmsy.2021.10.006>
- Jaiswal, A., Arun, C. J., & Varma, A. (2022). Rebooting employees: Upskilling for artificial intelligence in multinational corporations. *The International Journal of Human Resource Management*, 33(6), 1179-1208. <https://doi.org/10.1080/09585192.2021.1891114>
- Villani, V., Pini, F., Leali, F., & Secchi, C. (2018). Survey on human–robot collaboration in industrial settings: Safety, intuitive interfaces and applications. *Mechatronics*, 55, 248-266. <https://doi.org/10.1016/j.mechatronics.2018.02.009>
- Hogan, A., Blomqvist, E., Cochez, M., d'Amato, C., de Melo, G., Gutierrez, C., et al. (2022). Knowledge graphs. *ACM Computing Surveys*, 54(4), Article 71. <https://doi.org/10.1145/3447772>
- Leng, J., Sha, W., Wang, B., Zheng, P., Zhuang, C., Liu, Q., Wuest, T., Mourtzis, D., & Wang, L. (2022). Industry 5.0: Prospect and retrospect. *Journal of Manufacturing Systems*, 65, 279-295. <https://doi.org/10.1016/j.jmsy.2022.09.017>
- Nahavandi, S. (2019). Industry 5.0 — A human-centric solution. *Sustainability*, 11(16), 4371. <https://doi.org/10.3390/su11164371>