

The Ultimate Guide to Industrial AI in Manufacturing

Everything You Need to Know About
Industrial AI in Manufacturing



Introduction

Artificial intelligence is no longer a future ambition for manufacturers – it is already reshaping how work gets done on the factory floor. From improving decision-making to accelerating problem-solving and enabling more resilient operations, Industrial AI is becoming a core pillar of modern manufacturing strategies with connected workers.

This guide provides a comprehensive deep dive into Industrial AI in manufacturing. It explores what Industrial AI is, how it has evolved, the use cases delivering value today and how emerging approaches like Agentic AI will shape the future of human-machine collaboration on the factory floor.

Table of Contents

1. What is Industrial AI?
2. What to Look for in an Industrial AI Solution
3. Top Use Cases of Industrial AI in Manufacturing
4. Selecting the Right Industrial AI Solution
5. Where to Start with Industrial AI
6. Next Steps in Your AI Journey



1. What Is Industrial AI?

Industrial AI refers to the application of artificial intelligence technologies – including machine learning, generative AI and autonomous AI agents – within industrial environments to augment human decision-making, optimize operations and improve factory-floor performance.

Unlike traditional automation, which focuses on executing predefined rules, Industrial AI focuses on intelligence. It analyzes context, learns from data, identifies patterns, and supports people in making better decisions in real time.

Crucially, Industrial AI is not about replacing workers. Manufacturing environments are too complex, variable and safety-critical for that. Instead, Industrial AI succeeds by augmenting frontline workers, supervisors and engineers, ensuring the right information and insights are available exactly when and where work happens.

Industrial AI vs. Traditional Factory Software

Manufacturers already rely on systems like MES, SCADA, ERP, CMMS and QMS. These systems of record are essential – but they were not designed to help connected workers reason through ambiguity, adapt to change or continuously improve.

Industrial AI complements these systems by closing the decision gap between data and action. It connects structured and unstructured information, embeds intelligence into workflows and supports human judgment at the Gemba.

The Evolution of AI in Manufacturing

- AI adoption in manufacturing has evolved through several phases:
- Rule-based automation and analytics, focused on monitoring and reporting
- Predictive AI, using historical data to forecast outcomes

- Generative AI, enabling natural language interaction and content creation
- Agentic AI, allowing autonomous AI agents to augment human performance

The Role of Generative AI in Manufacturing

Generative AI has been the most visible and widely adopted form of artificial intelligence to reach the manufacturing sector, marking a critical inflection point in how AI is used on the factory floor to connect workers. Unlike earlier AI approaches that required specialized expertise or complex configuration, Generative AI introduced intuitive, natural language interaction—making AI accessible to frontline workers for the first time.

At its core, Generative AI uses large language models to create, summarize, translate and contextualize information. In manufacturing environments, this capability has proven especially valuable for transforming dense, unstructured content into formats that are easier to consume and act on in the flow of work.

On the factory floor, Generative AI is already delivering tangible value by accelerating the creation and maintenance of digital work instructions, converting SOPs and manuals into step-by-step guidance and enabling intelligent search across operational knowledge. Instead of searching through binders, shared drives or multiple systems, connected workers can ask questions in natural language and receive clear, relevant answers in seconds.

Generative AI also plays a critical role in workforce enablement. By supporting multilingual translation and transcription, it helps break down language barriers in increasingly diverse manufacturing teams—improving comprehension, safety and quality without adding administrative burden.

Importantly, Generative AI operates as a supportive assistant, responding to user requests rather than acting independently. This makes it well suited for environments where trust and control are paramount.

However, because it is inherently reactive, Generative AI is best understood as a foundation rather than a final destination.

By reducing the effort required to create, access and maintain knowledge, Generative AI lays the groundwork for more advanced AI capabilities. It frees up human time, standardizes information and improves data quality—creating the conditions necessary for the next evolution of AI in manufacturing.

In this way, Generative AI is not just improving today's factory-floor workflows; it is enabling the transition toward more proactive, intelligent and agent-driven systems that will define the future of industrial operations.



The Emergence of Agentic AI: Everything You Need to Know About Augmenting Connected Workers

Agentic AI represents the next major evolution of artificial intelligence in manufacturing, moving beyond Generative AI toward autonomous, goal-driven systems that can actively support work on the factory floor. While Generative AI has proven its value by accelerating content creation, simplifying access to knowledge and enabling natural language interaction, it remains largely reactive – responding only when prompted by a user.

Agentic AI changes that dynamic.

At a foundational level, AI agents are autonomous software systems capable of perceiving their environment, reasoning about objectives, coordinating tools and proposing or executing actions with minimal human intervention. These agents do not simply generate information; they understand context, plan next steps and learn over time, adapting their behavior as conditions change.

In a manufacturing environment – where variability, time pressure and safety constraints are the norm – this shift is especially significant. Agentic AI enables AI to function not just as a tool, but as a digital co-worker: one that operates alongside frontline workers to help troubleshoot issues, surface early warnings and recommend corrective actions before problems escalate.

From Assistants to Digital Co-Workers

Over the last few years, manufacturing has moved rapidly through the AI hype cycle and into a phase where tangible, shop-floor use cases matter more than experimentation. Generative AI applications, powered by large language models, are already delivering value by converting dense SOPs into digital work instructions, creating engaging training content and enabling smarter search across operational knowledge. These capabilities have reduced deployment time, lowered content creation costs and improved comprehension across multilingual workforces.

However, this is only the beginning.

Agentic AI builds on these foundations by enabling AI systems to take initiative. Rather than waiting for a worker to ask a question, an AI agent can recognize that help may be needed, interpret intent and proactively assist. In practice, this means AI evolving from a passive assistant into a semi-autonomous, goal-directed agent that supports workers in the flow of daily operations.

Importantly, this evolution does not diminish the role of human workers. On the contrary, Agentic AI is explicitly designed to augment, not replace, human expertise. By handling repetitive background tasks, coordinating information across systems and surfacing insights workers do not have time to find themselves, AI

agents free up human capacity to focus on higher-value, judgment-based work – improving both performance and job satisfaction.

Human-in-the-Loop: The Foundation of Trust

For Agentic AI to succeed on the factory floor, trust is non-negotiable. Manufacturing environments demand high reliability, safety and accountability – which is why the most effective Agentic AI implementations are built around a Human-in-the-Loop (HITL) model.

In this model, AI agents may be autonomous in how they analyze situations, coordinate tools, and propose actions, but humans always retain final authority. This distinction between autonomy and being fully autonomous is critical. AI agents can act, but they do not operate unchecked. Guardrails ensure that workers validate, approve or override AI-recommended actions before they are executed.

This approach does more than mitigate risk – it actively builds trust. By keeping humans involved in decision-making, Agentic AI becomes a collaborative partner rather than a black box. Over time, as workers see consistent, reliable support from AI agents, confidence grows and adoption accelerates.





An Evolutionary Path for Agentic AI in Manufacturing

The relationship between AI agents and frontline workers will not change overnight. Instead, it will evolve in clear phases:

Phase 1: Assistants and Automation (Today)

AI agents function primarily as conversational assistants and background automation tools. They support content creation, data retrieval, rule-based triggers and administrative tasks, reducing manual effort without altering core workflows.

Phase 2: Goal-Directed Autonomy (Near Future)

Reasoning agents emerge that can interpret user intent and coordinate multiple tools or sub-agents. For example, when an operator begins troubleshooting an issue, an orchestrator agent identifies the goal, searches relevant knowledge, prepares a draft issue report, and asks the user for approval before taking action.

Phase 3: Proactivity and Prognostics (Looking Ahead)

Agents evolve from reactive to proactive. They begin detecting early signals of breakdowns, deviations, or inefficiencies before they fully manifest. These agents not only predict what may happen, but recommend what actions to take – learning continuously and improving over time through feedback loops and operational outcomes.

Practical Agentic AI Use Cases on the Factory Floor

Agentic AI's value becomes most tangible through practical, worker-centric use cases:

AI-Assisted Troubleshooting: Orchestrator agents interpret intent, route requests to the right knowledge sources and prepare next steps – reducing downtime and cognitive load.

Early Warning Systems for Leaders: AI agents proactively surface emerging trends, such as recurring safety checklist issues, enabling intervention before incidents occur.

Shift Performance Prediction: By analyzing workforce skills, checklist completion and line conditions, agents can identify when a shift is at risk and recommend corrective actions to protect OEE and quality.

In each case, AI acts as a supportive partner, providing a “helping hand” rather than dictating outcomes.

Building the Future, Hand in Hand

Research suggests manufacturing leaders are increasingly open to AI-driven decision support, even in critical functions. Trust in AI is already higher in manufacturing than in many other industries – creating a strong foundation for Agentic AI adoption. Still, questions remain around cultural readiness, governance and the right balance between efficiency and oversight.

What is clear is this: the true power of Agentic AI to augment connected workers lies in the partnership between human expertise and machine proactivity. Manufacturers that remain curious, open and engaged with this technology will be best positioned to unlock productivity gains and build resilient, future-ready operations.

The future of manufacturing will not be defined by machines replacing people – but by people empowered by intelligent, proactive digital co-workers.

Benefits of Industrial AI on the Factory Floor

When implemented effectively, Industrial AI delivers strong outcomes for an augmented connected workforce, such as:

- Increased productivity and OEE
- Faster time to value
- Improved safety and compliance
- Better knowledge capture and retention
- A more empowered, engaged workforce



2. What to Look for in an Industrial AI Solution

Industrial AI has enormous potential to improve manufacturing performance, but value is not guaranteed simply by adding AI to existing systems. In practice, many AI initiatives fail to scale because they are disconnected from real work, difficult to trust or misaligned with the realities of the factory floor.

To deliver sustainable value, an Industrial AI solution must be designed specifically for manufacturing environments and frontline users. The right solution should meet the following criteria.

Human-Centric, with Explainability and Trust

At its core, Industrial AI must be built to augment human expertise, not override it. Manufacturing is a domain where context, judgment and experience matter deeply, and frontline workers need to understand why an AI system is making a recommendation.

Human-centric Industrial AI provides transparency into how insights are generated and allows workers to validate, approve or override AI-suggested actions. This explainability is essential for building trust, ensuring safety and encouraging adoption. Solutions that operate as black boxes may look impressive in demos, but they rarely succeed on the factory floor.

Integrated, Not a Disconnected AI Layer

Industrial AI should not exist as a standalone analytics tool or an isolated AI overlay. Instead, it must be deeply integrated across workflows, data sources and applications.

The most effective solutions connect AI across knowledge management, tasks, issues, training and analytics, so insights flow naturally from one activity to the next. When AI is fragmented across multiple tools, workers are forced to switch contexts, duplicate effort and question which system to trust—undermining the very efficiency AI is meant to deliver.

Embedded in Daily Operations

AI creates value only when it is available at the moment of work. Industrial AI should be embedded directly into daily factory-floor activities such as shift handovers, inspections, troubleshooting, maintenance rounds and continuous improvement routines.

If AI insights live exclusively in dashboards or reports that are reviewed after the fact, they will have limited impact. Embedded AI supports real-time decision-making at the Gemba, where issues occur and value is created.

Easy to Use for Non-Technical Workers

Frontline adoption depends on simplicity. An Industrial AI solution must be intuitive enough for operators, technicians and supervisors—many of whom are not data scientists or IT specialists.

Natural language interaction, clear prompts, minimal configuration and mobile-first design all contribute to ease

of use. If a solution requires extensive training or technical expertise, adoption will stall and value realization will be delayed.

Built for Scale, Security, and Compliance

Manufacturing organizations operate at scale, often across multiple plants, regions and regulatory environments. Industrial AI must be enterprise-ready from day one.

This includes robust cybersecurity, data privacy controls, role-based access, auditability and compliance with industry regulations. Scalability also means supporting multilingual workforces, sharing standards across sites and maintaining consistent performance as usage grows.

Manufacturing-First, Not Retrofitted from IT

Finally, Industrial AI must be designed specifically for manufacturing—not adapted from generic enterprise or IT use cases. Manufacturing-first solutions reflect the structure of factory operations, align with continuous improvement methodologies and support the unique rhythms of shift-based work.

AI that is retrofitted from office or IT environments often fails to account for the complexity, variability and physical realities of the factory floor. Manufacturing-first design is the difference between AI that looks promising in theory and AI that delivers results in practice.





3. Top Use Cases of Industrial AI in Manufacturing

Industrial AI creates value across every function on the factory floor. The most successful manufacturers focus on use cases that augment human performance, reduce friction, and improve consistency.



Operations Use Cases: Improve Productivity and OEE

Operations teams face constant pressure to maintain output, quality and uptime – often with limited visibility and fragmented information.

Instead of searching through systems

Industrial AI supports operations by:

- Assisting operators during troubleshooting
- Predicting shift-level performance risks
- Highlighting deviations before they escalate
- Summarizing operational data into actionable insights

or waiting for reports, workers receive contextual guidance in real time, enabling faster decisions and fewer disruptions.

Health & Safety Use Cases: Reduce Risk and Improve Compliance

Safety performance depends not only on procedures, but on awareness, consistency and proactive intervention.

Industrial AI enhances safety by:

- Identifying patterns in near-misses and observations
- Flagging emerging risks before incidents occur
- Supporting standardized inspections and audits
- Reinforcing safety culture through real-time feedback

AI does not replace safety professionals – it amplifies their ability to see trends and act earlier.

Maintenance Use Cases: Improve Asset Availability

Industrial AI extends maintenance effectiveness by:

- Detecting early warning signs of failure
- Supporting autonomous maintenance by operators
- Improving communication between production and maintenance
- Reducing unplanned downtime

Rather than replacing CMMS or EAM systems, Industrial AI connects maintenance intelligence directly to frontline workflows.

Continuous Improvement Use Cases: Accelerate Learning

Continuous improvement depends on identifying problems, learning from them, and standardizing better ways of working.

Industrial AI supports CI by:

- Clustering issues and recurring deviations
- Helping to find likely root causes
- Supporting PDCA cycles with data-driven insights
- Capturing improvements as institutional knowledge

This transforms CI from a reactive process into a continuous learning loop.

Learning & Development Use Cases: Build Workforce Capability

With high turnover and skills gaps, training must be faster, more relevant and more accessible.

Industrial AI enables:

- Faster creation of digital work instructions
- Personalized, role-specific learning paths
- Skills gap identification and tracking
- On-the-job learning at the moment of need



4. Selecting the Right Industrial AI Solution

Selecting the right Industrial AI solution is a strategic decision that will influence not only operational performance, but also how artificial intelligence is perceived, trusted and adopted across the organization. A well-chosen solution can accelerate improvement and empower frontline teams, while a poor fit can reinforce skepticism and stall future innovation.

To make an informed choice, manufacturers should take a structured approach that balances business needs, workforce realities and long-term scalability.

Conduct a Needs Assessment

The selection process should begin with a clear understanding of where AI can deliver the most value. Rather than starting with technology capabilities, focus on operational challenges and decision points.

Key questions to ask include:

- Where do frontline workers lack timely or accurate information?
- Where does process variability lead to quality issues, downtime or safety risk?
- Where do the same issues resurface without lasting resolution?

These problem areas often represent decision gaps where human judgment is stretched by complexity, time pressure or fragmented data. Industrial AI is most effective when applied to these moments – augmenting human decision-making rather than adding another layer of complexity or automation.

Vendor Evaluation Criteria

Once needs are defined, manufacturers should evaluate vendors through a manufacturing-first lens. Not all AI providers are equally suited to industrial environments.

Strong Industrial AI vendors:

- Specialize in manufacturing and understand factory-floor realities
- Have proven deployments in live production environments, not just pilots
- Demonstrate a clear commitment to human-in-the-loop design
- Embed AI directly into workflows where work is performed, rather than isolating it in dashboards or reports

Experience matters. Vendors with real-world manufacturing use cases are better equipped to navigate variability, safety requirements and change management challenges.

Security, Governance, and Trust

AI adoption hinges on trust. Manufacturers must have confidence that AI systems are secure, transparent and governed appropriately.

Evaluation should include:

- Robust data security and privacy controls
- Clear explainability around how AI insights and recommendations are generated
- Defined governance mechanisms that determine when AI can act autonomously and when human approval is required

These guardrails protect the organization while reinforcing workforce confidence in AI-supported decisions.

Built for Enterprise Scale

Finally, the selected solution must be able to scale across the enterprise. Manufacturing organizations often operate multiple plants across regions, each with local nuances but shared standards.

Enterprise-ready Industrial AI should support:

- Multi-plant and multi-region deployments
- Multilingual workforces and diverse skill levels
- Global standardization with flexibility for local adaptation

By choosing a solution that aligns with both immediate needs and long-term growth, manufacturers can ensure their Industrial AI investment delivers sustained value and becomes a trusted part of daily operations.



5. Where to Start with Industrial AI

One of the most common pitfalls manufacturers encounter with Industrial AI is attempting to apply it everywhere at once. While the technology is powerful, unfocused deployments often lead to stalled initiatives, limited adoption and skepticism on the factory floor. A successful Industrial AI journey starts with focus, clarity and a strong foundation of trust.



Begin with a Business Value Assessment

The first step is to conduct a business value assessment that identifies where Industrial AI can deliver meaningful impact quickly. Rather than pursuing complex, high-risk applications out of the gate, manufacturers should prioritize high-impact, low-risk use cases that solve real operational problems.

Common starting points include:

- Knowledge conversion, such as transforming SOPs, checklists and manuals into accessible digital work instructions
- Troubleshooting assistance, where AI helps workers find relevant information, past solutions or next steps more quickly
- Reporting automation, reducing the time spent compiling shift reports, audit documentation or performance summaries

These use cases deliver quick wins, reduce manual effort, and demonstrate tangible value—helping to build confidence among both frontline teams and leadership.

Assess Organizational Readiness

Technology alone does not determine success. Industrial AI adoption depends just as much on organizational readiness as it does on technical capability.

Manufacturers should assess:

Workforce trust in AI, including perceptions around transparency, fairness and job impact

Change management readiness, ensuring leaders are prepared to communicate purpose and support adoption

Data availability and quality, as AI relies on accurate, contextual information to deliver reliable insights

Identifying gaps early allows organizations to address them proactively rather than encountering resistance later.

Steps to Implement Industrial AI

- With priorities and readiness established, implementation should follow a structured, iterative approach:
- Define clear objectives tied to specific operational outcomes, such as reduced downtime or faster onboarding

- Select focused initial use cases aligned with the business value assessment
- Engage frontline stakeholders early, incorporating their feedback into design and rollout
- Monitor results and iterate, refining use cases based on real-world performance

This incremental approach minimizes risk while accelerating learning.

Tips for Successful Adoption

Finally, adoption hinges on how Industrial AI is positioned and introduced. Clear communication is critical. Manufacturers should emphasize augmentation, not replacement, reinforcing that AI exists to support connected workers, not displace them.

Starting focused on key objectives and celebrating early successes helps build momentum and trust. Over time, these foundations enable Industrial AI to become a natural, valued part of daily factory-floor operations.



6. Conclusion: Next Steps in Your Industrial AI Journey

Industrial AI is becoming essential for manufacturers navigating Industry 5.0. The future belongs to organizations that pair human expertise with intelligent digital co-workers, creating safer, more productive, and more resilient operations.

By starting with the right use cases and the right mindset, manufacturers can build trust, unlock value, and lead the next era of industrial performance

Achieve an AI Advantage in Manufacturing

Watch this on-demand webinar to explore how the next generation of AI-augmented connected workers will collaborate seamlessly with intelligent assistants and autonomous agents – tools capable of understanding context, learning from operations, and taking action alongside frontline teams.

Watch Now ▶



Want to Bring AI to Your Factory Floor?

Book a Demo

About IFS



IFS is the world's leading provider of Industrial AI and enterprise software for hardcore businesses that make, service, and power our planet. Our technology enables businesses which manufacture goods, maintain complex assets, and manage service-focused operations to unlock the transformative power of Industrial AI™ to enhance productivity, efficiency, and sustainability.

IFS Cloud is a fully composable AI-powered platform, designed for ultimate flexibility and adaptability to our customers' specific requirements and business evolution. It spans the needs of Enterprise Resource Planning (ERP), Enterprise Asset Management (EAM), Supply Chain Management (SCM), and Field Service Management (FSM). IFS technology leverages AI, machine learning, real-time data and analytics to empower our customers to make informed strategic decisions and excel at their Moment of Service™.

IFS was founded in 1983 by five university friends who pitched a tent outside our first customer's site to ensure they would be available 24/7 and the needs of the customer would come first. Since then, IFS has grown into a global leader with over 7,000 employees in 80 countries. Driven by those foundational values of agility, customer-centricity, and trust, IFS is recognized worldwide for delivering value and supporting strategic transformations. We are the most recommended supplier in our sector. Visit ifs.com to learn why.