IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 4, September 2026
Impact Factor: 8.2

Artificial Intelligence in Industry: A Case-Based Analysis of Performance Enhancement

Dushyant Bodkhey[1] and Dr. Shriram Dawkhar[2]
Assistant Professor, SBS Research Centre, Pune, India[1]
Director, SBS Research Centre, Pune, India[2]
dushyanta80@gmail.com

Abstract: *Artificial Intelligence (AI) has become a pivotal enabler of industrial transformation, improving operational performance across manufacturing, automotive, energy and logistics. This study presents a cross-sector, source-attributed case analysis of AI adoption in four documented real-world deployments: predictive maintenance at PepsiCo (Frito-Lay), computer-vision quality control at the BMW Group, digital-twin optimization in GE's Digital Wind Farm, and route-optimization automation in UPS's ORION system. Because the cases report performance in heterogeneous and largely company-disclosed metrics, the study makes two methodological contributions. First, it interprets the evidence through a dynamic-capabilities lens, treating AI as a sensing-and-reconfiguring capability rather than a stand-alone tool. Second, it introduces a normalized five-dimension "AI Industrial Impact Scorecard" (performance magnitude, evidence robustness, deployment maturity, breadth of impact and SME transferability) that renders heterogeneous cases comparable on a common ordinal scale. The analysis confirms substantial, verifiable performance gains, but reveals an inverse relationship between impact and transferability: the highest-impact, best-evidenced deployments (UPS, GE) are the least accessible to small and medium-sized enterprises (SMEs), whereas sensor-based predictive maintenance is the most transferable. The paper contributes case-grounded, critically appraised evidence on AI as a driver of industrial competitiveness, and derives a staged adoption path for resource-constrained firms.*

Keywords: Artificial Intelligence; Industrial Performance; Predictive Maintenance; Digital Twin; Computer Vision; Workflow Automation; Dynamic Capabilities; Case Study Analysis; Industry 4.0

I. INTRODUCTION

The fourth industrial revolution has reframed the factory, the power plant and the distribution network as cyber-physical systems in which sensing, computation and physical action are tightly coupled [1], [15]. Within this transition, Artificial Intelligence (AI) has become the layer that converts the resulting data streams into decisions. Industrial organizations are moving from reactive and manual operation towards predictive, automated and data-driven operation, applying predictive maintenance, intelligent robotics, computer vision, digital twins and real-time analytics to raise productivity, reduce operational cost and improve product quality [2], [12].

Yet the discourse around industrial AI remains curiously bifurcated. On one side, a large academic literature catalogues what AI could do, emphasising architectures, algorithms and theoretical benefits, but offering comparatively little measured evidence of realized performance in production settings. On the other, a vast body of industry reporting describes individual deployments in enthusiastic but fragmented and often promotional terms, with figures that are rarely placed side by side or appraised for reliability. Between the two, a decision-maker seeking to answer a simple question — how much does industrial AI actually improve performance, and under what conditions — is poorly served.

Two specific problems follow from this gap. First, few studies connect named, real-world deployments to the sourced performance figures they report, or scrutinise whether those figures are independently verified or merely asserted by the implementing firm. Second, where outcomes are reported, they appear in incompatible, sector-specific metrics —

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 4, September 2026 Impact Factor: 8.2

gallons of fuel, defect classes eliminated, percentage energy yield — that resist systematic cross-case comparison. The result is that the evidence base, although large, is difficult to aggregate into transferable insight.

This paper addresses both problems through a structured, source-attributed case analysis. Its contributions are threefold. (i) It analyses four documented deployments spanning manufacturing, automotive, energy and logistics, attributing every performance figure to a primary company source or an independent analysis. (ii) It frames AI-driven performance through the *dynamic-capabilities* perspective [17], which explains why identical technologies yield uneven returns across firms. (iii) It introduces and applies a normalized *AI Industrial Impact Scorecard* that makes heterogeneous cases comparable on a common scale and surfaces a previously under-examined trade-off between impact and transferability to SMEs.

The remainder of the paper is organised as follows. Section II develops the theoretical background and reviews the literature on the four technology families. Section III sets out the research design, case-selection criteria, data sources and the scoring framework. Section IV presents the four cases in detail. Section V reports the comparative analysis, including the normalized scores and two supporting figures. Section VI discusses benefits, theoretical and managerial implications, SME adoption and limitations. Sections VII and VIII conclude and outline future research.

II. THEORETICAL BACKGROUND AND LITERATURE REVIEW

A. AI and Industrial Performance: A Dynamic-Capabilities Lens

A central puzzle in the industrial-AI literature is variance: firms deploying nominally similar technologies report widely different returns. The dynamic-capabilities view of the firm offers a parsimonious explanation. Teece frames competitive advantage not as the possession of a resource but as the capacity to sense opportunities, seize them through timely resource commitment, and continuously reconfigure the asset base as conditions change [17]. Industrial AI maps cleanly onto this triad: sensing corresponds to data acquisition, condition monitoring and anomaly detection; seizing corresponds to acting on model outputs — scheduling a repair, rejecting a part, re-routing a vehicle, retuning a turbine; and reconfiguring corresponds to the continuous retraining, process redesign and organizational learning that AI makes possible.

The implication is that the algorithm is rarely the scarce factor. Value is unlocked by the complementary assets that surround it — instrumented equipment, clean and labelled data, integration with control systems, and the skills to interpret and act on predictions [12], [16]. This lens frames the empirical analysis that follows and, in particular, anticipates the impact–transferability trade-off examined in Section V: large incumbents are more likely to already possess these complementary capabilities than resource-constrained SMEs.

B. AI in Industrial Applications

Recent systematic reviews document AI's migration from pilot studies into series production across process optimization, quality assurance, supply-chain management and equipment maintenance [12], [16]. Gao et al. survey the state of the art in manufacturing AI and identify data quality, model interpretability and system integration as the recurring barriers between technical feasibility and realized value [12]. Peres et al. frame "Industrial Artificial Intelligence" as a systematic discipline spanning data, analytics, platform and operations layers, and argue that industrial adoption lags consumer AI precisely because the surrounding engineering and organizational layers are harder to assemble than the models themselves [16]. Lee et al. earlier articulated the same programme, positioning industrial AI as the bridge between Industry 4.0 architectures and measurable manufacturing outcomes [15].

C. Predictive Maintenance

Predictive maintenance (PdM) uses machine-learning models on sensor data — vibration, acoustic, thermal, current — to forecast equipment failure before it occurs, replacing both reactive repair and fixed-interval preventive maintenance [4], [14]. Carvalho et al.'s systematic review finds that supervised learning methods dominate, with random forests and neural networks most frequently reported [4]. Çınar et al. review ML for PdM in the context of sustainable smart manufacturing and report consistent downtime and cost benefits where the data pipeline is mature [14]. Industry benchmarks point in the same direction: McKinsey estimates 30–50% reductions in machine downtime and 20–40%

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 4, September 2026 Impact Factor: 8.2

extensions in machine life for well-run programmes [10], while Deloitte reports equipment-uptime improvements of roughly 10–20% alongside reduced maintenance cost [11]. These figures describe mature programmes, not first pilots, and a non-trivial false-positive rate can erode the gains [10].

D. Computer Vision in Quality Control

Deep-learning computer vision has improved the consistency of defect detection relative to manual inspection, reducing both missed defects and false rejections [13]. Saberironaghi et al. review defect-detection methods for industrial products and document the shift from hand-crafted features to convolutional and, more recently, transformer-based architectures capable of operating within production cycle times [13]. The practical value in high-volume manufacturing often lies as much in reducing false rejections — parts wrongly flagged as defective — as in catching true defects, because false rejections impose rework and inspection cost without quality benefit [7], [13].

E. Digital Twin Technology

A digital twin is a virtual, continuously updated representation of a physical asset, synchronised through sensor data, that supports simulation, monitoring and optimization in real time [3]. Tao and Zhang articulate the shop-floor digital-twin paradigm as the convergence of physical and virtual production [3]. In asset-intensive sectors such as energy, twins coupled with machine learning are used to tune operating parameters and extract additional output from existing hardware [8]. The same logic has been extended upward to the supply chain, where digital-twin models support disruption management and resilience [5].

F. Workflow Automation and Optimization

In logistics and distribution, AI-driven optimization and automation reduce cost and environmental impact at network scale by solving large routing and scheduling problems continuously against live data on traffic, demand and capacity [9]. Unlike the preceding technology families, the binding constraint here is less sensing hardware than proprietary data assets and operations-research competence; the optimization is only as good as the data and the modelling capability behind it [16], [17].

G. Measurement and the Research Gap

Across these literatures, three measurement problems recur. Outcomes are reported in incompatible units; many figures originate from the implementing firm or its technology vendor and are not independently audited; and unsuccessful deployments are systematically under-reported, biasing the visible evidence upward. Prior work rarely connects named deployments to sourced, critically appraised figures, and provides no common scale for comparing outcomes across sectors. This study addresses both the attribution gap and the comparability gap — the first through source attribution, the second through the normalized scoring framework introduced in Section III.

III. RESEARCH METHODOLOGY

A. Research Design

The study adopts a qualitative multiple-case design, which is appropriate for examining a contemporary phenomenon in its real-world context where the investigator cannot control events and where rich, situated understanding is the goal. It is explicitly a desk-based synthesis of documented deployments rather than primary field measurement; this choice trades the control of first-hand data collection for breadth across sectors and access to large-scale commercial deployments that researchers could not themselves instrument. The design is framed, and its limitations discussed, accordingly.

B. Case Selection

Four cases were selected against four criteria, applied as filters to a longer list of candidate deployments: (i) the deployment is in commercial series operation rather than a laboratory or proof-of-concept trial; (ii) it represents a distinct AI technology family — predictive maintenance, computer vision, digital twin, or workflow automation — so that the set spans the main modes of industrial AI; (iii) each case sits in a different industrial sector, maximising cross-sector variance; and (iv) performance outcomes are available from a primary company source or an independent

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 4, September 2026 Impact Factor: 8.2

analysis rather than from secondary aggregation alone. The resulting set — PepsiCo (Frito-Lay), BMW Group, GE and UPS — satisfies all four criteria and offers maximal variation in sector, technology and reporting source.

C. Data Sources

Evidence was drawn from company press releases and innovation disclosures, an INFORMS operations-research case record, and peer-reviewed review articles. Each performance figure reported in Section V is attributed to its originating source. Where only industry-level benchmarks are available — for example, aggregate predictive-maintenance statistics — these are labelled explicitly as benchmarks rather than as case results, and are used only to contextualise, not to substitute for, case evidence.

D. The AI Industrial Impact Scorecard

To compare cases reported in incompatible metrics, each deployment is scored 1–5 on five dimensions and summarised by an unweighted composite (the arithmetic mean of the five). The dimensions capture not only how large a reported gain is but how trustworthy and how transferable it is — properties a raw percentage conceals. The rubric (Table II) was defined a priori, before scoring, to reduce the risk of fitting the scale to a desired ranking. Scoring is ordinal and interpretive: it is intended to support structured comparison and to make reasoning explicit, not to claim cardinal precision.

TABLE II. AI Industrial Impact Scorecard — scoring rubric

Dimension	What it measures	Anchors (1 = low, 5 = high)
D1 Performance magnitude	Size of the reported improvement over the pre-AI baseline	1 = marginal; 3 = moderate/bounded (~10–20%); 5 = transformational (>30% or a defect class eliminated)
D2 Evidence robustness	Independence and verifiability of the reported figures	1 = single vendor claim; 3 = primary company disclosure; 5 = independent/third-party or peer-reviewed
D3 Deployment maturity	Operational stage and scale of the deployment	1 = pilot/PoC; 3 = multi-site; 5 = full commercial-scale rollout
D4 Breadth of impact	Scope of operations affected	1 = single task; 3 = a line/process; 5 = enterprise- or network-wide
D5 SME transferability	Ease of adoption by resource-constrained firms	1 = large-firm only (capital/data-intensive); 3 = moderate; 5 = readily adoptable (off-the-shelf, low data barrier)

E. Analytical Approach

The analysis proceeds in three steps. First, each deployment's reported outcome is tabulated against its pre-AI baseline (Table I). Second, the scorecard is applied, with the rationale for each score stated explicitly (Section V.B) so that readers can contest individual judgements; the results are presented numerically (Table III) and visually (Figures 1 and 2). Third, the cases are synthesised thematically into common benefits, challenges, and theoretical and managerial implications. Because several figures are company- or vendor-reported, they are treated throughout as claims to be compared rather than as independently audited results.

F. Threats to Validity

Four standard threats to validity in case-based research are addressed explicitly. Construct validity — whether the study measures what it claims — is supported by defining each scorecard dimension a priori with explicit anchors (Table II) and by attributing every performance figure to a named source. Internal validity, concerning causal inference, is deliberately bounded: the study reports associations between AI adoption and performance outcomes as disclosed, and does not claim to isolate AI as the sole cause of those outcomes against confounders such as concurrent process improvement. External validity is limited by the choice of four large firms and is discussed as the SME-generalisability

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 4, September 2026 Impact Factor: 8.2

boundary, which the transferability dimension renders visible rather than hides. Reliability — whether another analyst would reach the same scores — is supported by publishing the rubric and the per-case rationale (Section V.B) so that each ordinal judgement can be independently reviewed and contested; the scores are presented as reasoned, reproducible judgements rather than measurements.

IV. INDUSTRIAL CASE STUDIES

A. Manufacturing — Predictive Maintenance at PepsiCo (Frito-Lay)

Context. PepsiCo's Frito-Lay division operates high-throughput snack-production lines in which unplanned stoppages are costly: a failed motor or bearing can idle an entire line and spoil in-process product. Traditional maintenance alternated between running equipment to failure and servicing it on fixed schedules, neither of which matches maintenance to actual machine condition.

AI solution and mechanism. Working with machine-health firm Augury, Frito-Lay fitted production equipment with vibration and ultrasonic sensors that stream continuously to machine-learning models trained to recognise the signatures of developing faults — imbalance, misalignment, bearing wear — well before failure. In dynamic-capabilities terms, the sensors and models supply the sensing layer, and the maintenance-scheduling response supplies the seizing layer [17].

Reported results. PepsiCo Labs reported that, after one year, its Frito-Lay pilot plants recorded zero breakdowns, interruptions or incremental costs attributable to the monitored equipment, and the company subsequently moved to scale the solution to nearly all of its manufacturing plants [6]. The direction and magnitude are consistent with independent PdM benchmarks of 30–50% downtime reduction and 10–20% uptime improvement for mature programmes [10], [11], [14]. The figures are company-disclosed and qualitative at the headline level, and are therefore read here as directional.

B. Automotive — Computer Vision at the BMW Group

Context. Automotive production demands consistent inspection at line speed across stamping, assembly and final checks. Conventional rule-based camera systems are prone to false rejections — flagging dust, oil films or lighting artefacts as defects — which trigger unnecessary manual re-inspection and slow the line.

AI solution and mechanism. BMW has used AI image recognition in series production since 2018 [7]. At the Dingolfing press shop, a neural network trained on around 100 reference images per feature compares each component against known-good examples in milliseconds, distinguishing genuine fine cracks from dust or oil that earlier systems misread — so-called pseudo-defects. At the Steyr engine plant, trained analysis of cold-test data separates insignificant torque-measurement irregularities from real faults, removing unnecessary manual re-tests.

Reported results. BMW reports that press-shop pseudo-defects “no longer occur,” that redundant manual double-checks are avoided, and that reliability reaches close to 100% after a short tuning run [7]. The value here is the elimination of a false-rejection class and the associated inspection labour rather than a single headline percentage; the review literature corroborates that deep-learning vision now meets production cycle-time and consistency requirements [13]. Evidence is primary-source and largely qualitative.

C. Energy — Digital Twin in GE's Digital Wind Farm

Context. A wind farm's output depends on how each turbine is configured and operated against site-specific wind conditions; small, continuous optimizations compound across dozens of turbines and decades of operation.

AI solution and mechanism. GE's Digital Wind Farm builds a cloud-based digital twin of each site on the Predix platform. Engineers select among turbine configurations at design time, and in operation dozens of per-turbine sensors feed a machine-learning analytics engine that recommends continuous adjustments — for example, trading rotor speed against noise limits or yaw against wind direction. The twin is the sensing-and-simulation layer; operational retuning is the reconfiguring layer [17].

Reported results. GE estimated that the approach could increase a wind farm's energy production by up to 20% and generate approximately US$100 million in additional value over the lifetime of a 100 MW farm [8]. The striking

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 4, September 2026
Impact Factor: 8.2

feature is that the gain is extracted from existing physical hardware purely through software optimization. The figure is a company estimate expressed as an upper bound ("up to") and is therefore treated as a strong but vendor-sourced, projective claim.

D. Logistics — Route Optimization in UPS ORION

Context. UPS operates one of the world's largest delivery networks, where a small reduction in distance per route, multiplied across tens of thousands of drivers and millions of daily stops, produces very large absolute savings in fuel, time and emissions.

AI solution and mechanism. The On-Road Integrated Optimization and Navigation (ORION) system applies advanced operations-research and optimization algorithms to design delivery routes from map data, package information and real-time inputs, continuously balancing distance, delivery commitments and capacity. UPS's proprietary package and telematics data are the complementary asset that makes the optimization effective — precisely the kind of asset SMEs lack [16], [17].

Reported results. INFORMS records that, at full deployment, ORION was expected to cut roughly 100 million miles driven and save about 10 million gallons of fuel per year, reduce CO_2 emissions by roughly 100,000 metric tons annually, and deliver US$300–400 million in annual savings, with more than US$320 million already achieved by December 2015 [9]. These are among the best-documented figures in the set, recorded by an independent professional body rather than asserted solely by the firm.

V. COMPARATIVE ANALYSIS

A. Reported Outcomes

Table I consolidates the four deployments, each figure attributed to its source. The metrics are deliberately left in their native units to avoid false precision; comparison across them is handled by the scorecard that follows.

TABLE I. Reported performance outcomes across four AI deployments

Case / Sector	AI technology	Reported outcome (vs. pre-AI baseline)	Source
PepsiCo – Frito-Lay (food mfg.)	Predictive maintenance	Pilot plants reported zero breakdowns, interruptions or incremental costs over one year; scaled to nearly all plants	PepsiCo / Augury [6]
BMW Group – Dingolfing, Steyr (automotive)	Deep-learning computer vision	Press-shop pseudo-defects eliminated; redundant manual re-inspections avoided; ~100% reliability after tuning	BMW Group [7]
GE – Digital Wind Farm (energy)	Digital twin (Predix)	Up to 20% increase in energy production; ~US$100M added value over a 100 MW farm's life	GE [8]
UPS – ORION (logistics)	Route optimization / automation	≈100M fewer miles and ≈10M gallons of fuel per year; US$300–400M annual savings at full deployment	INFORMS / UPS [9]

B. Normalized Scoring and Rationale

Applying the scorecard renders the four cases comparable (Table III). The rationale for each row is stated so that individual judgements can be contested.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 4, September 2026
Impact Factor: 8.2

TABLE III. AI Industrial Impact Scorecard — case scores

Case	D1	D2	D3	D4	D5	Composite /5
PepsiCo – Frito-Lay	4	3	4	4	4	3.8
BMW Group	4	3	5	3	3	3.6
GE Digital Wind Farm	5	3	4	4	2	3.6
UPS ORION	5	5	5	5	2	4.4

PepsiCo – Frito-Lay (3.8). High on performance (a reported elimination of monitored-equipment breakdowns, D1 = 4) and, unusually, high on transferability (D5 = 4): sensor-retrofit PdM delivered as a managed service has a low data barrier. Evidence is company-disclosed (D2 = 3) and the rollout is multi-site and scaling (D3 = 4, D4 = 4).

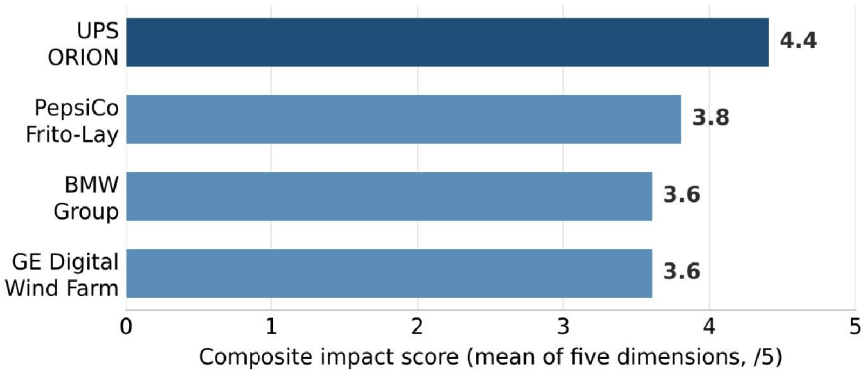
BMW Group (3.6). Strong on maturity (series production since 2018 across named plants, D3 = 5) and performance (a false-rejection class eliminated, D1 = 4), but narrower in scope (specific inspection stations, D4 = 3) and moderately transferable (D5 = 3), as vision QC requires data and engineering effort. Evidence is primary but qualitative (D2 = 3).

GE Digital Wind Farm (3.6). Highest on raw magnitude (up to 20% output gain, D1 = 5) and fleet-wide in breadth (D4 = 4), but the figure is a vendor estimate (D2 = 3) and the approach is capital- and data-intensive, hence poorly transferable to smaller asset owners (D5 = 2).

UPS ORION (4.4). The strongest overall: transformational, independently documented, network-wide impact at full commercial scale (D1–D4 = 5). Its only low score is transferability (D5 = 2), because the result depends on proprietary data assets and operations-research capability beyond the reach of most firms.

Figure 1 plots the composite scores; Figure 2 shows the full dimension-by-case matrix, making the patterns visible at a glance.

Figure 1. Composite AI Industrial Impact scores by case



IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 4, September 2026
Impact Factor: 8.2

Figure 2. Per-dimension scores across cases (1 = low, 5 = high)



C. Cross-Case Patterns

Two patterns stand out. First, every case improves against its pre-AI baseline, but the strength of the *evidence* varies more than the strength of the *claim*: the independently documented cases (UPS, and to a lesser extent GE's publicly scrutinised figures) should carry more weight than purely internal disclosures. Second, and more importantly, the data reveal an **inverse relationship between impact and SME transferability**. The two highest-impact cases — UPS and GE — score lowest on transferability (D5 = 2), precisely because their returns depend on proprietary data and specialist capability. Conversely, the most transferable deployment, sensor-based predictive maintenance at PepsiCo, is not the highest-impact but is the most readily adoptable by firms without deep data infrastructure. This trade-off, rather than any single headline figure, is the study's central empirical finding, and it follows directly from the dynamic-capabilities logic of Section II.

D. Synthesis by Technology Family

Reading the cases by technology family rather than by firm clarifies why the transferability scores fall as they do. Table IV summarises, for each family, the sensing requirement, the capability that most often constrains realized value, and the resulting barrier to adoption by resource-constrained firms. The pattern is consistent with the dynamic-capabilities account: families whose bottleneck is a widely available, retrofittable sensing service (predictive maintenance) are broadly adoptable, whereas families whose bottleneck is proprietary data scale or a high-fidelity model (optimization, digital twins) concentrate their returns among firms that already hold those assets.

TABLE IV. Industrial AI technology families — capability requirements and adoption barriers

Technology family	Primary sensing requirement	Capability bottleneck	SME barrier	Representative case
Predictive maintenance	Condition sensors (vibration, acoustic, thermal)	Labelling of failure events; maintenance-workflow integration	Low–moderate (available as managed service)	PepsiCo – Frito-Lay
Computer vision QC	High-resolution imaging at line speed	Labelled image sets per feature; line integration	Moderate	BMW Group
Digital twin	Dense, continuous asset telemetry	High-fidelity physics/ML model and compute	High (capital- and data-intensive)	GE Digital Wind Farm
Workflow optimization	Operational and transactional data at scale	Proprietary data scale and operations-research skill	High	UPS ORION

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 4, September 2026
Impact Factor: 8.2

VI. DISCUSSION

A. Benefits of AI in Industry

Across the four cases, the benefits of industrial AI cluster into five recurring themes: enhanced operational efficiency and asset utilization; reduced unplanned downtime and maintenance cost; more consistent product quality with fewer false rejections; measurable fuel, energy and emissions savings; and a broad shift from reactive to predictive, data-driven decision-making. Notably, several of these benefits are extracted from existing physical assets — GE lifts output from installed turbines, UPS from the existing fleet — which makes AI an unusually capital-efficient route to performance gains where the complementary data assets already exist.

B. Theoretical Implications

The cases support a dynamic-capabilities reading of industrial AI [17]. In every instance, value materialised where the firm had already built the sensing layer and the organizational capacity to act on and reconfigure around AI outputs; the model itself was rarely the binding constraint, whereas data infrastructure, system integration and skills repeatedly were [12], [16]. This both explains the observed variance in returns and predicts the impact–transferability trade-off: the complementary capabilities that make high-impact AI viable are exactly those that large incumbents possess and SMEs typically lack. Industrial AI, on this reading, tends to widen rather than narrow the capability gap between large and small firms unless that asymmetry is deliberately addressed.

C. Managerial Implications

For managers, the analysis suggests that technology selection should be subordinate to capability assessment. Before choosing an AI application, a firm should ask which of the sensing, seizing and reconfiguring capabilities it already possesses and which it must build. The scorecard offers a practical screening tool: applications that score high on transferability relative to the firm's current capability are the rational entry points, with higher-impact but capability-hungry applications sequenced later. The persistent lesson across cases is to invest in data quality, integration and skills in parallel with the model, not after it.

D. Challenges

The same cases expose recurring obstacles that temper the headline gains:

- Integration with legacy equipment and heterogeneous control systems
- High initial capital and instrumentation cost, concentrated before any return
- Dependence on data quality, labelling and availability
- Cybersecurity exposure of newly connected industrial assets
- Workforce skill gaps and the change-management burden of new workflows

E. Implications for SMEs

The impact–transferability trade-off has a direct, constructive corollary for resource-constrained firms. Rather than emulating the flagship deployments of large incumbents, SMEs are better served by a staged path: begin with high-transferability, sensor-based predictive maintenance offered as a managed service — low data barrier, retrofittable, fast time-to-value — and use it to build the data discipline and skills base; progress next to vision-based quality control; and only then attempt capital- and data-intensive applications such as digital twins or network-scale optimization. At the ecosystem level, cluster- or association-level data-sharing and shared analytics services could lower the transferability barrier that currently concentrates the largest gains among the largest firms.

F. Limitations

Several limitations bound these conclusions. First, several figures originate from the implementing firms or their technology vendors and have not been independently audited; promotional incentives may bias disclosure, and unsuccessful deployments are under-reported, so the visible evidence likely overstates typical returns. Second, the scorecard is ordinal and interpretive, and its unweighted composite treats all five dimensions as equally important, which may not hold for every decision context; a weighting scheme would change the composites though not the central trade-off. Third, the four cases are large, well-resourced firms, so the results may not generalise to SMEs — a

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 4, September 2026
Impact Factor: 8.2

boundary the transferability dimension makes explicit rather than resolves. Fourth, with four cases the analysis is illustrative rather than statistically representative. These limitations motivate the future work set out in Section VIII.

VII. CONCLUSION

Through four named, source-attributed cases, this study shows that AI can significantly enhance industrial performance across manufacturing, automotive, energy and logistics, with documented improvements in reliability, fuel and cost efficiency, defect reduction and energy yield. By framing the evidence through dynamic capabilities and comparing it through a normalized scorecard, the paper moves beyond the undifferentiated claim that "AI helps" to a sharper, conditional conclusion: realized value depends on complementary data and organizational capabilities, and the most powerful deployments are, for that very reason, the least transferable to smaller firms. For practitioners, the actionable message is to match AI applications to existing capabilities and to sequence adoption from transferable to capability-hungry. For researchers, the claimed gains — being largely company-reported — should be interpreted critically and validated with independent measurement wherever possible.

VIII. FUTURE RESEARCH DIRECTIONS

1. Independent, audited before/after measurement of AI deployments to counter disclosure bias
2. Weighted and empirically validated versions of the impact scorecard across larger case samples
3. Explainable AI (XAI) for safety-critical industrial decisions
4. Edge AI for low-latency, real-time industrial control
5. SME-focused adoption models and cluster-level data-sharing mechanisms
6. Standardized, cross-sector frameworks for measuring industrial AI performance

REFERENCES

[1] J. Lee, B. Bagheri, and H.-A. Kao, "A cyber-physical systems architecture for Industry 4.0-based manufacturing systems," Manufacturing Letters, vol. 3, pp. 18–23, 2015.
[2] T. Wuest, D. Weimer, C. Irgens, and K.-D. Thoben, "Machine learning in manufacturing: advantages, challenges and applications," Production & Manufacturing Research, vol. 4, no. 1, pp. 23–45, 2016.
[3] F. Tao and M. Zhang, "Digital twin shop-floor: a new shop-floor paradigm towards smart manufacturing," IEEE Access, vol. 5, pp. 20418–20427, 2017.
[4] T. P. Carvalho et al., "A systematic literature review of machine learning methods applied to predictive maintenance," Computers & Industrial Engineering, vol. 137, 106024, 2019.
[5] D. Ivanov and A. Dolgui, "A digital supply chain twin for managing the disruption risks and resilience in the era of Industry 4.0," Production Planning & Control, vol. 32, no. 9, pp. 775–788, 2021.
[6] Augury, "PepsiCo Labs and Augury prevent snack-plant breakdowns together" (Frito-Lay machine-health deployment), case materials. Available: augury.com (accessed Nov. 2026).
[7] BMW Group, "Fast, efficient, reliable: Artificial intelligence in BMW Group Production," press release, Munich, 15 Jul. 2019. Available: press.bmwgroup.com (accessed Nov. 2026).
[8] General Electric, "Digital Wind Farm" (Predix digital-twin platform), 2015; reported in The Maritime Executive, "Digital wind farms promise 20 percent productivity gain," 2015. Available: maritime-executive.com (accessed Nov. 2026).
[9] INFORMS, "UPS On-Road Integrated Optimization and Navigation (ORION) project," O.R. & Analytics Success Stories, 2016. Available: informs.org (accessed Nov. 2026).
[10] McKinsey & Company, "Smartening up with artificial intelligence (AI): what's in it for Germany and its industrial sector?" industry report, 2017.
[11] Deloitte, "Predictive maintenance and the smart factory," Deloitte Analytics Institute, 2017.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 4, September 2026 Impact Factor: 8.2

[12] R. X. Gao, J. Krüger, M. Merklein, H.-C. Möhring, and J. Váncza, "Artificial Intelligence in manufacturing: state of the art, perspectives, and future directions," CIRP Annals – Manufacturing Technology, 2024, doi: 10.1016/j.cirp.2024.04.101.

[13] A. Saberironaghi, J. Ren, and M. El-Gindy, "Defect detection methods for industrial products using deep learning techniques: a review," Algorithms, vol. 16, no. 2, art. 95, 2023, doi: 10.3390/a16020095.

[14] Z. M. Çınar, A. Abdussalam Nuhu, Q. Zeeshan, O. Korhan, M. Asmael, and B. Safaei, "Machine learning in predictive maintenance towards sustainable smart manufacturing in Industry 4.0," Sustainability, vol. 12, no. 19, art. 8211, 2020, doi: 10.3390/su12198211.

[15] J. Lee, H. Davari, J. Singh, and V. Pandhare, "Industrial Artificial Intelligence for Industry 4.0-based manufacturing systems," Manufacturing Letters, vol. 18, pp. 20–23, 2018, doi: 10.1016/j.mfglet.2018.09.002.

[16] R. S. Peres, X. Jia, J. Lee, K. Sun, A. W. Colombo, and J. Barata, "Industrial Artificial Intelligence in Industry 4.0 – systematic review, challenges and outlook," IEEE Access, vol. 8, pp. 220121–220139, 2020, doi: 10.1109/ACCESS.2020.3042874.

[17] D. J. Teece, "Explicating dynamic capabilities: the nature and microfoundations of (sustainable) enterprise performance," Strategic Management Journal, vol. 28, no. 13, pp. 1319–1350, 2007, doi: 10.1002/smj.640.