# AI ACROSS THE INDUSTRIAL VALUE CHAIN

What we actually found

KRIYANTA.ai

Over the past year, we have built and tested AI engines across the industrial value chain: direct material cost, net working capital, manufacturing operations, sales and after-sales, engineering cost, IT service management, HR, and central functions.

The pattern was consistent.

AI is highly effective at finding value faster than traditional diagnostic work – comparing complex data, applying rules at scale, drafting outputs, routing exceptions, and triggering workflows. It can scan purchase orders, invoices, inventory lines, quotes, service tickets, CVs, bills of materials, and historical project records at a scale traditional diagnostics cannot match.

But industrial value is not created by AI output alone.

AI can identify a pricing outlier – but EBITDA moves only when the sourcing decision changes. AI can draft a quote – but margin improves only when pricing discipline and delivery assumptions are right. AI can flag an early-payment leak – but cash improves only when the payment is blocked or corrected.

This is the industrial AI execution gap: the distance between what AI can find, draft, recommend, route, or trigger – and what the business converts into measurable operational or financial impact.

The major advisory and technology firms have described the surrounding agenda well: data foundations, ROI logic, governance, enterprise reinvention, and fit-for-purpose adoption. All of that matters. But in practice, we kept seeing the same patterns – across domains, company sizes, and levels of AI maturity.

Not answers to one neat question. Six recurring conditions that determine whether AI output becomes operational and financial reality – or stays in a dashboard.

# An AI finding that stays in a dashboard is not EBITDA.

### 1. AI compresses diagnosis – but only if you know where to point it

AI is powerful immediately: it finds value faster than traditional diagnostic work.

In direct material cost – often the largest controllable EBITDA lever in industrial businesses – AI can scan raw purchase order data, segment spend, identify price outliers against physical performance drivers, cluster suppliers, flag high-cost-country dependency, and prepare negotiation logic. Work that previously took weeks can take hours.

In net working capital, it surfaces slow-moving inventory before it becomes an accounting issue, identifies payment-term drift, flags early-payment leakage, and detects lead-time inconsistencies that quietly tie up cash.

In sales and tendering, it compares a live opportunity against historical quotes, win rates, delivery performance, and realised margins – exposing pricing assumptions that look reasonable alone, but weak against comparable project history.

The diagnostic compression is real. In transformation contexts, speed matters commercially. But AI finds what you ask it to look for. The quality of the question determines the quality of the diagnosis.

**The better question is: where is value leaking, what moves EBITDA and cash, and what has to change operationally to capture it?**

# Work that previously took weeks can take hours. But only if you ask the right question first.

**2. Insight is not impact. Action is.**

Diagnosis is only the beginning. A finding does not move EBITDA. Action does.

A direct material cost engine may identify high-cost-country supplier dependency and sourcing potential based on part complexity, testing requirements, and supplier footprint. But the P&L changes only when a supplier is challenged, an RFQ is released, an alternative source is qualified, volume is shifted, and the saving is locked into the budget.

A net working capital engine may identify early-payment leakage or inventory buffers that no longer reflect actual lead times. But cash improves only when a payment is blocked before release, an ERP parameter is changed, or planning rules are reset.

An AI tendering engine may flag that a quote is underpriced against comparable historical projects. But margin improves only when the bid owner reprices the offer, adjusts scope, or escalates before submission.

Industrial value is created when a finding becomes an action, the action becomes a workflow, the workflow has an owner and deadline, and the result is visible in EBITDA or cash. That is the difference between a finding and a result.

Material cost is the primary EBITDA lever in most industrial businesses. Labour and functional cost are major structural levers. Net working capital is the cash lever – especially in PE-backed situations where liquidity discipline determines the trajectory of the business. AI must be connected to all three – and judged by whether it moves them, not by whether it finds them.

### 3. AI works best where the logic exists but execution is manual

The strongest AI applications are often not the most sophisticated ones. They sit in work that is rule-heavy, repetitive, fragmented, and too expensive to perform manually at scale.

This includes: CV screening against defined criteria; IT ticket classification and self-service routing; RFQ comparison against historical benchmarks; supplier shortlisting; payment term checks against contract logic; BOM extraction; working capital exception monitoring; and tender pricing validation.

None of this requires senior judgement at every step. It requires consistent application of known logic across large volumes of data and decisions.

The rules often already exist – buried in spreadsheets, ERP fields, local practices, old contracts, and employee experience. AI creates value when it applies those rules consistently, quickly, and across the full dataset.

## The logic exists. AI applies it at a speed and scale that manual execution never could.

**4. AI does not fix broken processes. It makes them impossible to ignore.**

When the process is broken, AI surfaces the mess underneath before it finds the value opportunity.

Inconsistent master data. Unclear ownership. Local workarounds. Missing approval rights. Drifting ERP parameters. No clean baseline. No measure owner.

This is not a reason to avoid AI. It is a reason to sequence it correctly:

If the work is unnecessary, eliminate it.

If the process is broken, redesign it.

If the process is valid but manual, apply AI.

If the legacy system cannot be replaced quickly, wrap it with an AI execution layer.

If the decision involves risk, trust, labour relations, or technical judgement, keep humans in the approval loop.

## The mistake is not applying AI in complex industrial environments. The mistake is automating the mess.

**5. The real adoption barrier is trust, not technology**

Even when the process is sound and the AI output is useful, value can still stall. The harder problem is often operational trust.

In one HR service management context, the opportunity was clear: move from a human-centric HR support model to an AI-guided service model. The harder part was not the AI. It was changing behaviour. Employees were used to submitting a request and waiting for HR.

Until they trusted the AI-guided portal enough to resolve simple questions there first, request volume stayed with human agents, workload did not shift, and the financial case remained theoretical.

Adoption is not a soft metric. It is an economic lever. If users bypass AI, the workload stays in the old model and the savings stay on a slide.

The same applies across procurement, finance, IT, and sales. A buyer must trust the sourcing logic before challenging a supplier. Finance must trust the exception logic before blocking a payment. IT users must trust the AI-guided portal before submitting fewer tickets to human support.

In Europe, trust also has a legal dimension. AI deployments that change how work is done, who does it, or how performance is monitored require Works Council engagement from the design phase. Effective industrial AI must be designed for trust from the outset – especially in co-determined environments.

The technology timeline and the social license timeline are not the same. Treat Mitbestimmung as an execution variable, not a legal hurdle, and the question shifts from whether to how.

## Trust must be built into the deployment model – not added as a communication plan after go-live.

### 6. The architecture that matters most is decision architecture

Data architecture matters. Technology architecture matters. Governance matters.

But financial impact depends on something more practical: decision architecture.

That means being clear on five things: what has to change, who owns it, what data is needed, what workflow triggers the action, and how EBITDA or cash impact is verified.

This is where many AI programmes are under-designed – not at the model layer, but at the handover from recommendation to action.

Without a clear owner, AI output becomes a report. With one, it becomes a trigger.

Human approval gates are not a weakness. In industrial environments, they are often what makes AI deployable. The goal is not full autonomy. The goal is structured, auditable action with a clear owner at every consequential decision point.

A simple workflow can create more value than an impressive AI platform. An agent that flags an early-payment exception, routes it to the right approver, blocks release in SAP, and records the avoided leakage may not look spectacular in a demo. But it moves cash – and every step is traceable, reversible, and owned.

AI should not be judged by how impressive the output looks. It should be judged by whether it changes the actions that drive the economics.

# Output without ownership is just a report.

### Where to start

The next phase of industrial AI will not be won by the company with 10,000 pilots, the longest workflow catalogue, or the most exhaustive diagnostic phase. It will be won by the company that connects AI to operational levers, workflow owners, and verified financial impact.

The starting questions are not about technology. They are about value: where is margin being lost silently? Where is cash trapped? Where does the organisation know what to do, but fail to do it consistently at scale?

An AI Roadmap Sprint starts with one practical question: where in the value chain can AI create structural advantage, and what must happen for that advantage to become EBITDA or cash?

The output is a prioritised execution plan: value lever, data requirement, workflow logic, owner, human approval points, implementation sequence, and financial impact path.

AI finds the value. Execution captures it.