POSITION PAPER

The Human in Command

Why AI without human judgment is a mirage - and what operational leaders must do about it.

> **"AI gives superpowers to those who have mastered their craft.**
> **It puts an F1 car in the hands of someone who cannot stay on a go-kart track."**

Peter van Roosmalen
25+ years in Value Creation, manufacturing, logistics, and process industries
Value Creation · Augmented AI · Financial Value Translation
www.gembanomads.com

March 2026

OPENING

A confession, and a provocation

I get asked about AI constantly. How I use it. What works and what does not. Whether it will replace people like me. These are good questions, and I try to answer them honestly - because honest answers are in short supply in this conversation.

Most of what I hear in the market is a version of the same story: AI is a force multiplier. Deploy it, learn to prompt it, and watch the output scale. One experienced consultant becomes ten. Ten become a hundred. The knowledge is extracted, encoded, distributed. Problem solved.

I have been working in operational environments for 25 years. I have stood at machines at 6am trying to understand why a process fails. I have sat with leadership teams that were too close to their own problems to see them clearly. I have built improvement systems that lasted - and watched others collapse within months of the consultant leaving.

From that experience, I want to say something clearly: **AI without human judgment is not a shortcut. It is a mirage.**

> *A mirage is not an illusion in the sense of being imaginary. It is a real optical phenomenon. It looks convincing. It looks like water. It looks like the destination. The problem is that it disappears when you walk toward it - and in the desert, believing in it can kill you.*

AI-generated output has the same quality. It is coherent. It is confident. It is often impressively detailed. And it can be wrong in ways that are invisible to anyone who does not already possess the domain knowledge to question it. The danger is not that AI is unreliable. The danger is that it is unreliable in ways that look reliable - which is the most expensive kind of failure in any operational environment.

This paper is my honest account of what AI can and cannot do in operational environments. It is not a technology overview. It is a practitioner's argument for why the human in command is not a nice-to-have - it is the only thing that makes AI useful.

AI gives superpowers to those who have mastered their craft.

It puts an F1 car in the hands of someone who cannot stay on a go-kart track.

An F1 car does not become dangerous because it is poorly engineered. It becomes dangerous because it amplifies everything - including the mistakes of the person behind the wheel. At 300 km/h, there is no margin for imprecision. If the input is wrong, the outcome is not a small error. It is a crash with momentum.

PART 1

The promise that does not hold

There is a generation of voices telling organisations that AI democratises expertise. That the gap between a novice and an expert narrows when both have access to the same AI tools. That intelligence scales with the prompt.

This is partially true - and that partial truth is what makes it dangerous.

Yes, AI can help a less experienced analyst surface questions they would not have thought to ask. Yes, it can help a junior writer produce a first draft that is structurally sound. Yes, it can process data at a speed no human can match. These are real capabilities, and I use them daily.

But here is what the democratisation story leaves out: the ability to evaluate AI output requires exactly the expertise that AI is supposed to replace. When an AI produces a root cause hypothesis, you need domain knowledge to test it. When an AI generates a market analysis, you need context to spot the gaps. When an AI recommends an operational intervention, you need Gemba experience to know whether it will hold.

AI does not lower the bar for good thinking.

It raises the stakes for bad thinking. The output is more polished, more voluminous, more persuasive - which means errors travel further and faster than they did before.

I have seen this in practice. A team uses AI to build a business case. The numbers are clean, the structure is logical, the conclusion is confident. Three months into implementation, the root assumption was wrong - and no one in the room had the operational depth to question it at the time. The AI did not make them less intelligent. It made their blind spots more expensive.

PART 2

A pattern we already know

I want to draw a parallel that every operational leader will recognise - because I have lived it many times.

When lean manufacturing tools were packaged and sold as off-the-shelf programmes in the 1990s and 2000s, the promise was similar: implement the system, train the people, and the results follow. It is a pattern that is repeating itself today with AI. 5S. Visual management. Kaizen boards. TPM. These are powerful methods. Thousands of organisations implemented them. And the majority failed to sustain the results - not because the tools were wrong, but because the tools were treated as the answer.

The tool reveals the quality of the thinking behind it. It does not replace it.

The organisations that sustained improvement were not the ones with the best tools. They were the ones where leaders walked the Gemba daily, asked questions that required thinking, and activated the system through their own behaviour - not as something delivered by a consultant and then left to run itself.

AI is a more powerful version of the same dynamic. The speed is higher. The outputs are more impressive. The failure modes are less visible. And the implementation plateau - the moment when the novelty fades and the absence of a human leadership system becomes apparent - arrives just as reliably.

Phase	What it looks like	What is actually happening
Launch	Energy, visible output, enthusiasm. AI is everywhere.	Novelty effect. Tool does what it promises. No system required yet.
Plateau	Results stabilise. Critical errors surface. Fatigue sets in.	Human judgment gap becomes visible. No oversight system in place. AI output accepted without verification.
Decay	AI use becomes selective or performative. Trust erodes.	Without the human system to sustain it, AI reverts to a productivity trick - or gets abandoned entirely.

This is not speculation. It is the operational improvement cycle, running at a faster cadence with higher stakes. The organisations that avoid it are not the ones that restrict AI - they are the ones that build the human system alongside it.

PART 3

What Augmented AI actually means

I use the term Augmented AI deliberately, and I use it differently from the way it circulates in marketing materials. For me it is not a product category. It is a set of non-negotiable operating conditions.

Augmented AI is not AI that helps humans. That is too vague to be useful, and it includes every AI product ever sold. Augmented AI is a specific relationship between human judgment and machine output - one where the human remains structurally in command at every critical point.

	Condition	What it means in practice
01	**The human sets the question**	AI does not define the problem. The practitioner does. AI explores within the boundaries of that question - it does not reframe it without explicit human direction.
02	**The human validates the reasoning**	Every AI-generated output is a hypothesis, not a conclusion. The practitioner tests it against operational reality, historical data, and causal logic before acting on it.
03	**The human owns the decision**	AI recommends. The practitioner decides. This is not a formality - it is the structural safeguard that keeps accountability, learning, and improvement where they belong: with the person who understands the context.

Remove any one of these conditions and you no longer have augmented intelligence. You have automation with a human signature on top - and the appearance of judgment without the substance of it.

In my own work, I apply this through a method I call 5WAI - a structured approach to root cause analysis that uses AI coaching to challenge the practitioner's thinking at each step. The AI does not produce the root cause. It asks better questions, surfaces assumptions, and flags logical gaps. The practitioner walks away with a sharper analysis and a deeper understanding of the reasoning chain - which is precisely what makes the improvement sustainable.

The name 5WAI reflects the principle: five levels of Why, augmented by AI, driven by human inquiry. It is not a tool that replaces thinking. It is a discipline that deepens it.

PART 4

What this asks of operational leaders

If the human must remain in command, the human must be capable of command. That is a more demanding standard than most AI adoption programmes acknowledge - and most do not acknowledge it at all.

The gap I see most consistently in organisations deploying AI is not technical. Leaders do not need to understand transformer architectures. They need three specific capabilities that current AI training programmes almost universally fail to develop:

- Critical evaluation of AI output - the ability to question AI-generated conclusions against real operational knowledge. Not rejection, not blind acceptance. Structured scepticism: what is the evidence for this? What is the assumption I cannot see? What does the Gemba say?
- Structured oversight systems - governance that makes human review routine, not exceptional. Who checks which AI outputs, at what frequency, with what authority to intervene. Without this system, oversight depends on individual discipline under pressure - which is exactly when it disappears.
- Decision quality awareness - the recognition that stress, cognitive overload, and time pressure are precisely the conditions under which people defer most readily to AI output. A leader who does not know their own judgment risk profile will be most vulnerable at the worst possible moment.

This is not a technology challenge. It is a leadership development challenge.

The same disciplines that sustain Value Creation - daily activation, structured questioning, accountability for outcomes - are the disciplines that make AI safe and valuable in operational environments.

The methodology is familiar. The context is new. The urgency is real.

The EU AI Act, which entered into force in August 2024 and has been binding since February 2025, formalises this in Article 4: organisations deploying AI systems have a legal obligation to ensure adequate AI literacy in their staff. This is not a future requirement. It is in force now. Most organisations are not compliant - not because they lack tools, but because they have treated AI adoption as a technology project rather than a leadership development challenge.

CLOSING

The honest answer

Will AI replace operational leaders and consultants? Some of what we do, yes - faster than most expected. The work that is purely analytical, purely generative, purely pattern-matching: AI does this well and will only improve.

But the work that matters most in operational improvement has never been purely analytical. It is the work of standing at a machine and understanding why the operator does what they do. It is the work of helping a leadership team see a pattern they are too close to recognise. It is the work of maintaining trust in a room when the news is difficult.

That work requires presence, judgment, and the willingness to be wrong in public and learn from it. AI augments all three. It does not replace any of them.

The mirage is convincing because it is partly real. AI is genuinely powerful. But the organisations that will gain lasting competitive advantage from it are not the ones that deploy the most tools. **They are the ones that build the human system capable of using those tools with judgment, discipline, and accountability.**

That is what I mean by the human in command. Not as a slogan. As an operating principle.

Augmented AI is not a technology strategy.

It is a leadership strategy for an age in which every tool

is more powerful than the last -

and the human in command must be equally capable.

Peter van Roosmalen
25+ years in Value Creation, manufacturing, logistics, and process industries
Value Creation · Augmented AI · Financial Value Translation
www.gembanomads.com

To discuss how these principles apply to your organisation, reach out directly.