AI & PLANNING SYSTEMS

The Supply Chain AI Debate Has Lost the Plot

Christian Kroschl · 17 April 2026

Article on kroschl.com · Read on LinkedIn

AI gets the attention. Ownership, foundations, and execution remain the real problem.

Coming back from LogiPharma in Vienna, one thing stood out.The supply chain AI debate has lost the plot.

For some, AI means ChatGPT or Claude. For others, it means copilots, digital twins, or another analytics layer with a new label. Very few ask the harder question:

Which decisions are we actually trying to improve? That is where the discussion usually breaks down.

I was reminded of this again recently in supply chain conversations shaped more by AI language than by decision logic. At the same time, I am reading *The Infinity Machine*, the story behind DeepMind. That contrast is useful.

AlphaFold did not succeed because people talked more about AI. It succeeded because the conditions were strong enough for AI to matter.

What made AlphaFold groundbreaking

AlphaFold was not a copilot story. It was not a story about putting a modern label on an old problem.

It was a serious scientific effort against a clearly defined challenge: predicting protein structure from amino acid sequences with useful accuracy.

That required several things to be true at the same time:

A sharply defined problem A meaningful target state Strong scientific data structures Clear evaluation logic A model architecture built for the problem Exceptional talent, research depth, compute, and funding

What made it groundbreaking was not just that it worked. It changed the level of what was considered possible.

For decades, protein folding had been one of the hardest problems in biology. AlphaFold showed that AI could reach a level of prediction quality that materially changed scientific work. It did not just produce interesting outputs. It shifted confidence, speed, and practical usefulness in a field where progress had been slow and difficult.

That is why AlphaFold matters. It was not another feature. It was a genuine step change. AlphaFold is not just an AI success story. It is a story about what must be true **before** AI can create breakthrough value.

Why the analogy matters for supply chains

That is exactly why AlphaFold is a useful analogy for supply chain.

Most companies are discussing AI as if the technology itself were the main question. It is not. The harder question is whether the underlying conditions are strong enough for advanced models to make sense at all.

The similarity is simple.

In both cases, value depends on more than algorithms. It depends on whether the problem is clearly framed, whether reality is represented with enough precision, whether success can be measured, and whether there is enough capability to build, train, and operationalize something serious.

This is where the gap begins. In AlphaFold, the problem was scientifically framed. In many supply chains, the problem is still politically fragmented. In AlphaFold, the target was clear. In many supply chains, even the decision to improve is not clearly owned.

In AlphaFold, the scientific environment was deep enough to support the work. In many companies, the data, process logic, and governance are not even stable enough to support a digital twin that deserves the name.

That is why the analogy matters. Not because proteins and supply chains are the same, but because both show that AI only becomes meaningful when the surrounding system is serious enough.

Why supply chains are different

Most supply chains do not fail because AI is immature. They fail because decision structures are weak. In many industries, supply chain can coordinate, but it cannot decide end to end.

Manufacturing owns capacity Commercial owns revenue Regions own local priorities Finance controls budget impact Technology clusters own systems Supply chain tries to connect the dots

That means one thing: **Responsibility is fragmented, while AI is discussed as if decisions were centralized.** But if no one really owns the decision, what exactly should AI optimize? This is also why many APS transformations stalled.

If APS is not trusted or not used in daily decision making, why should AI suddenly become relevant? And this is why the usual shortcut does not work either. Improve forecast accuracy, and all will be fine. That sounds good on a slide. In real supply chains, it is often false.

A 30 percent demand increase in the short term cannot simply be absorbed because a model saw it coming. Supply chain cannot just decide end to end and move. Regulatory constraints create delays and uncertainty. CMOs are often not flexible. Capacity, materials, release processes, and budget implications sit in different places.

So where exactly should AI optimize if the organization still cannot decide and act end to end?

If the digital twin is not really a twin because lead times, capacities, constraints, or execution realities are wrong, then AI is not solving the problem. It is scaling bad assumptions faster.

The real issue is not AI literacy

This is where I become skeptical when companies say they now need “AI literacy”. That sounds modern. In most cases, it is a convenient excuse. The issue is usually not that planners do not know enough AI vocabulary. The issue is that the company does not have:

Clear decision ownership Reliable operational data Strong process discipline The scientific capability to build serious models The budget and patience to industrialize them

DeepMind did not build AlphaFold with a few workshops on AI literacy. They built it with elite scientists, years of focused work, and major investment.

And that leads to an uncomfortable truth for most industries.

Most companies will never have that capability internally. They will not attract DeepMind level talent, they will not fund research like that, and they are usually not attractive enough as environments for that level of scientist.

A few technology firms and a very small number of research heavy companies may be exceptions. Most will not be.

That does not mean meaningful AI is impossible outside frontier research. It means the capability question has to be framed honestly.

Frontier AI needs extraordinary science. Industrial AI still needs serious problem framing, data engineering, model engineering, domain expertise, and ownership.

That is why partnerships matter. That is why selective investment matters. And that is why the first question is not “How do we use AI?” but:

Where does advanced modeling really make sense, and where should we stop the bullshit?

The market is diluting the term

The current AI debate in supply chain is not just overhyped. It is badly categorized.

At conferences, in vendor pitches, and in consulting narratives, almost everything is now called AI. That creates visibility, but it destroys precision.

Statistics becomes AI. Optimization becomes AI. Rule based automation becomes AI. A dashboard with a chatbot becomes AI. GenAI is treated as if it were the whole discipline.

The result is predictable. The label gets stronger while the substance gets weaker.

Many people speaking about AI in supply chain are in fact speaking about LLMs. Some barely distinguish between GenAI, machine learning, optimization, or reinforcement learning.

Once the language is imprecise, decisions become imprecise too.

Companies start buying “AI” without knowing whether they need pattern recognition, optimization, sequential policy learning, or simply better process discipline.

That is not transformation. That is market driven category confusion.

There is more than AI buzz

One reason the debate stays weak is that too many different methods are collapsed into one label.

Machine learning helps where patterns can be learned, for example in forecasting, lead time prediction, ETA prediction, or anomaly detection **Optimization / OR** helps where decisions are structured and constraints are explicit **Reinforcement learning** becomes relevant where supply chains face sequential control problems such as dispatching, dynamic allocation, or adaptive replenishment **Deep learning** matters where perception, high dimensional patterns, or complex sequence behavior are involved, for example in vision based quality inspection or sensor rich production environments **Pretrained transformers and GPTs** are useful for language, interaction, reasoning support, and unstructured knowledge, not as a substitute for planning logic or control models

It also helps to be specific.

Convolutional neural networks matter in visual tasks such as quality inspection, warehouse image recognition, label reading, packaging checks, or asset condition monitoring.

Recurrent neural networks historically mattered for time series and sequence tasks, but in many modern settings they are no longer the main story. They have often been replaced or complemented by newer sequence models, including transformer based approaches.

Reinforcement learning is more relevant for real execution problems than many conference discussions suggest, because supply chains are full of sequential decisions under uncertainty. But it is also much harder to implement well. It requires feedback loops, reward logic, simulation environments, and guardrails.

In some advanced settings, deep learning and reinforcement learning can be combined. Deep learning helps represent a complex environment. Reinforcement learning then learns which action is sensible in that environment.

LLMs sit elsewhere. They are strong as an interface and reasoning layer. They can support planners, summarize issues, classify documents, structure unstructured information, or improve access to knowledge. What they are not is a replacement for optimization logic, planning policy, or execution control.

The method has to follow the decision. Not the other way around.

The productivity myth

The next weak claim is that AI will suddenly cut 25 to 30 percent of effort in SC&O transformations. That sounds attractive. It is also highly misleading.

AI may help with selected tasks such as documentation, analysis support, or test case preparation. But real transformation effort does not disappear through better text generation or faster slide production.

The heavy part of transformation sits elsewhere: decision rights, data correction, process alignment, template discipline, stakeholder conflict, adoption, and governance.

AI can support parts of the journey. It does not carry the weight of the journey. AI does not cut transformation effort by 30 percent when most of the work is still structural, political, and organizational.

The market is reinforcing its own fiction

This problem is now being fed from both sides.

Consultants push inflated AI efficiency stories to stay commercially relevant. Procurement teams then feed the same logic back into sourcing processes, expecting lower effort, fewer people, and faster delivery.

At that point, the claim stops being a weak idea and becomes a market expectation.

Inflated AI assumptions are no longer just bad thinking. They start shaping budgets, delivery models, and transformation decisions.

A weak AI claim becomes dangerous the moment procurement turns it into a pricing assumption. From there, fiction starts shaping delivery.

What supply chains can learn from AlphaFold

The right lesson is not to copy DeepMind. The right lesson is to copy the discipline behind it. Before talking about AI in supply chain, companies should ask:

What is the decision? Who owns it? What data defines reality? How is success measured? What feedback updates the model or parameter? Do we have the capability to build this, or do we need a partner?

Without those answers, AI remains a slideware discussion. With them, it can become useful. Without them, it becomes expensive noise.

The bottom line

AI will not fix a fragmented supply chain organization. It will expose it.

AlphaFold worked because the problem, data, and scientific structure were strong enough for AI to matter. Supply chains need the same discipline, just in their own form: decision ownership, reliable structures, measurable outcomes, and the honesty to know when internal capability is not enough.

Until then, most AI talk in supply chain will remain what it often is today: **too broad, too vague, and too detached from the real decisions that run the business.**