THE
PLATFORM
WEEKLY
Automate the Difficult. Preserve the Human.

NOVEMBER 12, 2026
ISSUE 08

The Intelligence Between Things

Why the future of manufacturing may depend on understanding relationships, not just activities.



Automate the Difficult. Preserve the Human.

A WEEKLY EDITORIAL ON INTELLIGENT MANUFACTURING

The Intelligence Between Things

This week, a glimpse of a radically different manufacturing experience met the less visible work required to make it real. The result was not a finished future, but a clearer understanding of the one we are trying to build.

When the Future Became Visible

There are moments in a long undertaking when something familiar suddenly looks different. You have been discussing an idea for months, perhaps years, describing it in words, explaining the problems it might solve and asking others to imagine how the world could work. Then one day you see a small part of that possibility taking shape, and the conversation changes. This week, that moment came through an experiment in Lovable. On a phone screen, a jacket occupied the centre of a dark, almost cinematic environment. Around it were pathways into its definition, materials and the work involved in bringing it into existence. Scout was present within the experience, not hidden away in a separate application. The surfaces were translucent, the visual language was unusually fluid, and the product itself seemed to be the point of departure for everything else.

When Will saw it, he called it the *Minority Report* version. The comparison was immediate and memorable, but I don't think the interesting part was its resemblance to science fiction. We have all seen futuristic interfaces in films, and translucent panels alone will not improve the way a factory operates. What struck me was the change in perspective. For decades, enterprise software has been organized around the functions that companies perform. We open a purchasing module to find an order, a production module to inspect a schedule, a quality module to understand a failure and a communications tool to discover what people have said about it. We may have access to all the information, yet the responsibility for connecting it still rests with the person moving between those systems. The Lovable experiment invites us to imagine a different starting point: the thing we are trying to understand, with its relevant relationships available in context.

A jacket is a deceptively good example because it is not the same thing to everyone who touches it. To a designer, it begins as an intention, an expression of what a product should be. To a mill, it involves fibres, construction, colour, finishing and the practical knowledge required to reproduce a material. To a factory, it is a sequence of interdependent activities carried out by people and machines under real constraints. To a brand, it is also a commercial promise about quality, cost, timing and customer

expectation. These are not competing descriptions of different objects. They are different views of the same evolving reality. A material substitution may be sensible from a sourcing perspective and still threaten a previously approved sample. A small production change may protect a delivery date while creating a new quality risk. The intelligence we need lies in understanding how those perspectives affect one another.

Now imagine an environment that does not ask a person to reconstruct those relationships from a series of disconnected screens. A product manager opens the jacket and sees what has changed since the last review, which decisions require attention and why a particular issue matters. A mill technician sees the material's production history and the evidence behind its performance. A factory supervisor sees the operational consequences of a delay without being forced to examine unrelated work. A supplier sees the commitments and questions for which that supplier is responsible, but not information belonging to another organization. These are possibilities, not claims about what the current prototype can already do. Yet the possibility changes the ambition. The goal is no longer to make an old software structure look modern. It is to reconsider how people should encounter the complexity of manufacturing in the first place.

What makes this especially interesting is that the experiment was not confined to an elaborate desktop demonstration. We were exploring the experience on a mobile screen, where space is scarce and the cost of clutter is impossible to ignore. That constraint forces a useful question: if we cannot show everything, what truly deserves to be shown? Conventional systems often answer by shrinking the same menus and tables. A more intelligent environment might answer by understanding the user's role, the product's current circumstances and the decisions that cannot wait. The visual language becomes meaningful only when it serves that deeper purpose. The most advanced experience may be the one that allows people to think clearly without constantly thinking about the software they are using.

We should not confuse this glimpse with a completed product. The Lovable work is a limited experiment centred on one product, and much of the intelligence suggested by the experience remains to be engineered, connected and tested with customers. Its importance is not that we have already arrived. It is that we have found a more compelling way to ask where we are going. The interface makes an ambition visible that has sometimes been difficult to communicate: an operating environment in which sophisticated coordination happens beneath the surface, leaving people with a clearer view of what matters and the freedom to exercise their judgment. The next question is whether the engineering beneath that surface can support the promise above it.

The Complexity We Should No Longer Have to Carry

Manufacturing has always depended on people who can see beyond the immediate task. A production manager learns which delays will spread and which can be contained. A sourcing professional recognizes when an apparently harmless substitution will create a problem weeks later. A technician knows that two materials carrying the same specification may behave differently on a particular machine. This knowledge rarely lives in one place. It is distributed across organizations, conversations, documents and individual memories. It is also the reason manufacturing can continue to function despite systems that often understand transactions better than consequences. We have built powerful tools for recording what happened, but the relationships that explain why it matters are still too frequently reconstructed by people under pressure.

This week, a substantial phase of Process Management reached formal engineering acceptance and Product Owner approval, with a documented exception for some hosted verification that remains outstanding. That distinction matters because we should be as disciplined in describing progress as we are ambitious in describing the future. Yet the significance of the milestone goes well beyond a completed set of software functions. We have been forced to represent manufacturing processes as something richer than a sequence of boxes and arrows. Real work includes activities that proceed in parallel, conditions that change which steps are required, evidence that must exist before an approval, responsibilities that belong to different people and dependencies that can cross organizational boundaries. A process is not simply what happens next. It is a description of what must be true, who has the authority to act and what the consequences may be when circumstances change.

Consider the ordinary phrase "production is five days late." It sounds like a single fact, but its meaning depends on the relationships surrounding it. Perhaps cutting cannot begin until an approval arrives, while material procurement is already complete and can remain unaffected. Perhaps sewing and final quality inspection must move, but a separate design activity can continue. Perhaps the forecast changes without permission to alter the commitment made to a customer. In one verified engineering scenario, the system distinguished dependent activities from independent ones, preserved the original commitments and required authorization before applying changes. The achievement is not the arithmetic of moving dates. It is the discipline of preserving meaning while the situation changes. That is the kind of invisible work an intelligent environment must perform before it can offer a simple, trustworthy explanation to the person using it.

This is why the new visual experience and the Process Management milestone belong in the same story. One asks what manufacturing could feel like if technology understood the context of our work. The other begins establishing the governed relationships needed for that understanding to be reliable. If a

future product view is to tell someone that a material change threatens delivery, it must know which material was changed, which production activities depend on it, whose commitment is affected and who can authorize a response. A beautiful interface cannot invent those relationships. An eloquent AI model should not be allowed to guess them. The operating foundation must be able to distinguish evidence from inference, an expected date from a contractual promise, and a recommendation from an approved decision. Only then can apparent simplicity become more than a visual illusion.

There is a broader lesson here about the nature of intelligence. We tend to notice the visible moment when a machine answers a question, explains a problem or proposes a course of action. But the quality of that answer depends on a quieter achievement: whether the system has a faithful understanding of the world in which the question was asked. Manufacturing does not become intelligent merely because we attach a language model to a database. It becomes more capable when the relationships between products, processes, people and organizations are represented well enough that technology can reason about consequences without losing the context that makes those consequences important. That work is painstaking, and it is not always visually exciting. Yet it may be the most consequential part of the entire undertaking.

The Question Behind Scout

The same distinction surfaced in a conversation between Kevin and Will about Scout and the future role of intelligent agents. It would be easy to imagine Scout as an increasingly capable assistant that can answer questions about products, documents and schedules. But that description does not reach the harder architectural question. Where should intelligence actually live, and how much authority should it have? Some work needs predictable rules that can be tested, audited and consistently enforced. Other work involves ambiguity: interpreting a supplier's document, identifying a contradiction in a technical specification or investigating why an expected result did not occur. We should not force both kinds of work into the same mechanism simply because artificial intelligence has made new forms of automation available.

A governed operating system can establish that a particular approval is required before production proceeds. An intelligent capability may help discover that the approval is missing, explain the likely consequences and assemble the evidence needed for a decision. Those are complementary responsibilities, not interchangeable ones. A system that can interpret uncertainty should not be allowed to quietly rewrite the rules under which a business operates. Equally, a rigid rules engine should not be mistaken for an adequate response to every unexpected situation. The future may depend on the cooperation between deterministic engineering, specialized intelligence and human authority, with each

contributing what it does best. The more capable the technology becomes, the more important those boundaries will be.

Kevin's thinking has begun to explore Scout responsibilities that extend beyond answering a question on demand, including document interpretation, onboarding, event-driven investigations and work that can be initiated when circumstances warrant attention. Will's contribution pushed further toward persistent investigations, specialized capabilities and the knowledge that belongs to individual organizations. A factory has its own vocabulary, operating habits, approval practices and accumulated experience. A mill may have learned how to produce a particular material consistently under conditions that would be unfamiliar elsewhere. An intelligent environment that ignores these differences risks offering confident but generic advice. One that can understand them, within appropriate permissions, could become considerably more useful over time.

That possibility introduces an uncomfortable but necessary distinction between remembering and knowing. A system might observe that a particular technician resolved a problem by adjusting a machine setting. It may be useful to preserve the observation and retrieve it when a similar situation arises. But a past action is not automatically an approved procedure, and a successful result in one context may be dangerous in another. Organizational knowledge needs provenance, evidence and authority. People must be able to see why something is being recommended and decide whether it should become part of the accepted way of working. If we get that relationship right, Scout could help preserve experience without turning yesterday's improvisation into tomorrow's unquestioned rule. If we get it wrong, we could automate the very misunderstandings we are trying to eliminate.

I find that distinction more important than the current fashion for autonomous agents. The objective is not to remove people from every decision. It is to remove unnecessary effort from the process of reaching a good decision. Imagine an investigation that quietly gathers the relevant documents, identifies which organizations are affected, distinguishes established facts from unresolved questions and brings the right people into the conversation before a small problem becomes a costly one. That would be a meaningful change in how manufacturing operates. It would also be a change rooted in human judgment, because the system's contribution is to make that judgment better informed, not to pretend it is no longer needed.

The Memory of Making

Will's work on the mill environment gave this discussion another dimension. He shared detailed reference frameworks for fabric production, garment manufacturing and colour management, along with earlier thinking about fabric libraries, recipes, production batches and test evidence. The frameworks describe seven fabric processes, nine manufacturing processes and fifteen colour-management processes. Those numbers are less important than the question they expose. A fabric specification tells us what a material is expected to be, but it does not necessarily tell us how that material was made, why it performed as it did or what a mill learned while producing it. A finished product can meet a specification while concealing an extraordinary amount of knowledge about the decisions and conditions that produced the result.

Will described a problem that becomes particularly apparent when brands source similar materials from different regions. A fabric may have the same nominal composition and construction yet behave differently because the machines, processes, environmental conditions or local production capabilities are not the same. In his earlier work, including conversations involving Gymshark, this was not an abstract possibility. Consistency between production runs and manufacturing locations was a practical concern. The difference between what was specified and what was actually made can affect quality, fit, performance and the ability to reproduce a successful result. In some manufacturing environments, valuable process knowledge has historically been maintained through handwritten notes and informal instructions. That is an extraordinary vulnerability for an industry whose success depends so heavily on repeatability and accumulated experience.

The proposed relationship between a fabric article, its recipe, Bill of Process, production lot, testing and supporting evidence begins to address that gap. Importantly, the mill's view of a material and the brand's view should not become two unrelated versions of reality. The mill needs to understand how the material was produced; the brand needs to understand how it was selected, approved and used. Each should see what is relevant and permitted, while the underlying relationships remain coherent. This is the same principle we encounter in Process Management, viewed through a different part of the manufacturing world. Knowledge becomes more useful when it can retain its history and meaning as it moves between contexts, without requiring every participant to surrender control over its own information.

There is a philosophical possibility hidden inside this practical work. Manufacturing has always been a learning activity, but its memory is fragmented. Every successful production run, every failed test and every adjustment made by an experienced person contains a potential lesson. Most of those lessons never become available to the next person facing a similar problem, let alone to another organization working

on a related product. If an operating environment could preserve the connection between intention, execution, evidence and outcome, it might eventually help people recognize not only that something went wrong, but why similar problems recur. It might help a business learn from its own history without forcing its people to remember everything. That is a much more compelling use of artificial intelligence than generating another report after the fact.

The Future Must Survive Contact with Reality

These developments also reveal why the relationship between our MVP and the longer-term vision is changing. For some time, they have appeared to be parallel tracks. Kevin's engineering has had to answer to the practical demands of building a dependable product, while the broader exploration has asked what an intelligent manufacturing environment might eventually become. There is value in that separation. The MVP needs disciplined scope, a coherent architecture and the ability to earn the trust of real customers. A vision needs room to challenge the assumptions that have shaped enterprise software for decades. But treating them as separate destinations would be a mistake. We are not building one conventional product now and another futuristic product later. We are trying to build one operating model that can grow in capability without losing its integrity.

The immediate priority is unambiguous. The MVP must work in real manufacturing environments and demonstrate value to people willing to rely on it. Process Management, permissions, product relationships, reliable execution and the underlying data model matter because they support that objective. The Lovable experience does not outrank those requirements, nor does a speculative architecture for intelligent agents. Yet the future-facing experiments have begun to influence the questions we ask of the MVP. If the eventual experience is to be organized around a product rather than a collection of modules, then the product's relationships must be coherent across functions and organizations. If Scout is to explain consequences, then those consequences must be grounded in governed operating facts. If manufacturing knowledge is to accumulate over time, then evidence and history cannot be treated as disposable attachments. The vision is not another development program competing for resources. It is a discipline for recognizing which decisions today will matter tomorrow.

This week offered a reminder that such discipline also means resisting the temptation to confuse activity with evidence. Our separate investigation into the technical history of The Platform and its predecessors has been confronting the difference between counting software changes and understanding what those changes actually demonstrate. A large number of commits can reveal an enormous amount of work, but it cannot by itself establish the nature of every technological uncertainty, experiment or discovery. That lesson applies just as strongly to our own product narrative. We should not mistake an approved engineering milestone for proof that every customer scenario has been validated, nor a compelling visual

experiment for evidence that a new interaction model has already succeeded in practice. Ambition becomes more credible when we are precise about what has been demonstrated, what has been proposed and what we still need to learn.

This is where I find the relationship between vision and execution most encouraging. The engineering is beginning to reveal where the real complexity lies, and the vision is beginning to give that engineering a clearer purpose. Each informs the other without requiring them to advance at the same pace. A difficult dependency uncovered in the MVP may change how we imagine the future experience. A new way of looking at a product may reveal an architectural relationship that needs to be preserved today. The process is not always orderly, and it will not be free of disagreement. But perhaps the strongest ideas are not the ones that remain untouched by practical constraints. They are the ones that become more coherent as those constraints are understood.

The Value of a Difficult Conversation

The week's discussions with Dmitrii Tikhonov added an important outside perspective. Following a conversation with Will and Chris, he challenged the distinction between genuine intelligence and predefined rules, the difficulty of integrating with existing enterprise systems, the likely effort involved in customer adoption and the commercial value of what we are building. He also asked why additional capital should be pursued before securing a first paying customer. These are not minor objections, and they should not be dismissed simply because they come from someone who has not lived inside the initiative. They are precisely the questions that force an ambitious project to explain the difference between an attractive idea and a credible business.

His challenge about intelligence is especially timely. We can point to meaningful progress in governed processes and explain how Scout might eventually investigate more complex situations. But we should not pretend that deterministic process logic is the same as a fully realized intelligent operating environment. Nor should we undervalue it simply because it does not resemble the dramatic demonstrations associated with modern AI. Reliable relationships, evidence and authority may be what make advanced intelligence useful in the first place. The question is not whether every part of The Platform uses an AI model. It is whether the operating environment can help people understand, coordinate and decide more effectively than the fragmented tools available to them today. That claim will ultimately need to be tested in the lives of customers, not settled in an architectural debate.

Dmitrii's commercial questions are equally valuable, although the answer is not as simple as choosing between funding and a first customer. Customer validation is essential, and a paying customer can teach us more about adoption, workflow friction and value than another month of internal demonstrations.

At the same time, artificial intelligence's ability to accelerate software development does not eliminate the cost of the people responsible for making difficult decisions. The work draws on years of manufacturing knowledge and engineering experience, and those resources must be sustained if the product is to reach the point where customers can rely on it. The purpose of near-term funding is not to finance an endless journey toward a speculative future. It is to preserve the capacity to complete, validate and commercialize the foundations already being built. That is a business question as much as a technological one.

I also find it encouraging that the conversation has not ended with criticism. Dmitrii has expressed interest in examining the technology more closely, and a technical discussion with Kevin has been under consideration. We should not turn that interest into a promise of investment, a partnership or a new development team. It is an opportunity for another informed perspective to test our assumptions. The broader advisory and corporate conversations are part of the same maturation. As the undertaking becomes more tangible, questions of ownership, responsibility, intellectual property, commercial relationships and capital become harder to postpone. Those questions do not sit outside the vision. They determine whether the organization we build will be capable of sustaining it.

What Becomes Possible When the System Understands

I keep returning to the image of that jacket on the phone screen. Not because it is the most technically sophisticated thing we have built, and not because a visual experiment should determine the product roadmap. I return to it because it makes a difficult idea easier to imagine. For much of the history of enterprise software, people have been asked to enter the system's world, learn its categories and reconstruct reality from the fragments it presents. We are beginning to ask whether technology could instead enter the world of the person doing the work, understand the product and its circumstances, and bring forward the information that helps that person act wisely. That would be a different relationship with software, and potentially a different relationship between the organizations that make manufacturing possible.

The implications extend beyond one company. A brand, mill, factory and supplier will never share identical interests, responsibilities or knowledge. Nor should they be expected to. An intelligent manufacturing network would have to respect those boundaries while making the relationships between them more understandable. It would need to preserve permissions, evidence and accountability even as intelligent capabilities become more proactive. It might eventually allow specialized systems to investigate problems, coordinate permitted activities and help people recognize consequences that are currently discovered too late. But the measure of progress would not be how much autonomy we could demonstrate. It would be whether the network became more trustworthy, more responsive and more capable of learning without asking people to surrender the judgment that gives their work meaning.

There is a temptation to describe the future of artificial intelligence as a contest between what machines can do and what humans will no longer need to do. I find that framing too small. Manufacturing is full of work that people should never have been required to carry alone: searching for information that already exists, reconciling conflicting records, chasing explanations across organizational boundaries and rediscovering lessons that someone learned years earlier. If technology can take on more of that burden, people may have more time for the parts of manufacturing that demand imagination, expertise, responsibility and care. A technician could concentrate on improving a process rather than reconstructing its history. A production manager could spend more time solving a problem and less time discovering that it exists. A product team could make better decisions because the consequences of those decisions were visible before they became expensive.

We are still a long way from that future, and this week's work should not be presented as proof that we have arrived. The MVP must earn its place in the market, the visual experience must be tested against real work, the proposed intelligence must be grounded in trustworthy operating knowledge, and the commercial model must survive the scrutiny of customers and partners. Yet there is a difference between a distant vision that floats above reality and one that begins to shape the foundations being built beneath it. This week, through a futuristic visual experiment, a demanding engineering milestone, deeper manufacturing knowledge and a set of difficult questions, that connection became clearer. Perhaps the most important technology will not be the technology that announces its intelligence most loudly. It will be the technology that understands enough of the complexity to let people work with greater clarity, confidence and freedom.

Automate the Difficult. Preserve the Human.

The Platform Weekly reflects on the ideas, discoveries and questions shaping the development of one Intelligent Manufacturing operating environment. It is an exploration of the future we are trying to make possible, grounded in the work of building it.