Industrial Automation That Learns While It Works

A solution for High Mix, Low Volume Production

There are two ways to deliver intelligent automation.

One requires training on millions of examples. The other learns while it works. High-mix, low-volume industrial processes where the work changes frequently require something that learns the way people learn – on the job, from experience, with human guidance when needed.

Why HMLV Is Different

High-Mix, Low-Volume manufacturing sits at the center of a structural contradiction. Customers increasingly demand flexibility, variation, and rapid turnaround. Factories, robots, and processes are optimized for stability and repeatability.

This mismatch creates problems that cut across every role. Executives see margins erode on work they can't refuse. Middle managers firefight changeovers and schedule disruptions. Operators face cognitive overload from constantly shifting instructions. Integrators avoid HMLV automation projects because they're unprofitable and high-risk. Quality teams struggle to maintain standards when production runs are short.

Nobody defends the status quo here, and that's unusual. In most manufacturing contexts, there's resistance somewhere: job protection concerns, cultural friction, skepticism about ROI. In HMLV, operators don't see automation as a threat. They see it as relief from the most unrewarding work on the floor.

The Strategic Reality

Large industrial accounts increasingly expect their suppliers to handle HMLV needs: quick-turn prototypes, variant handling, engineering changes, small-lot runs. Declining this work, or handling it poorly, jeopardizes the overall relationship.

The math is straightforward. A small percentage of customers often generate most of the revenue. Those customers nearly always require some degree of HMLV capability. Being unable to say yes leads to defection or re-sourcing.

This means HMLV work isn't about direct profitability. It's about account control and long-term customer retention. Even a nominally low-margin job protects much larger revenue streams into the future. The supplier who can say yes reliably wins.

Why Current Approaches Struggle Here

Vision-language-action models and similar approaches require collecting demonstration data, training a model, validating it, and deploying it to production. For a new production run, this process can take weeks to months. If the task changes mid-run, you retrain. For a new edge case, you add it to the training set and retrain again.

Shop floors have little patience for this, especially in HMLV where the job deadline might pass before training completes.

There are also systemic issues that won't resolve with more compute or better data.

Interpretability is a real concern. End-to-end learning systems encode knowledge in neural network weights. When something goes wrong, there's no way to ask "why did it do that?" and get a human-readable answer. For manufacturing environments where quality audits, regulatory compliance, and operator oversight matter, this opacity is a serious issue.

Human oversight requires a place for humans to intervene. When an operator sees something wrong, they need to correct it. In training-based systems, the only intervention is "provide more training data for it to learn from", which is not a substitute for in-situ oversight.

A further concern is knowledge ownership. Not only are the network weights inscrutable, they belong to the vendor. Your hard-won knowledge about fixture drift, supplier variability, and equipment-specific behaviors disappears into a model you can't inspect and don't own. Switch vendors and you start from zero. Worse, your operational knowledge may end up improving a model that also serves your competitors.

Our system inverts this. The generalizations it learns are human-readable artifacts stored at your site. You can inspect them, edit them, export them. Change cognitive infrastructure providers (eg, ChatGPT to Claude) and that knowledge comes with you. The learning belongs to your business, not the platform.

What We've Built

We have developed a system that learns from experience the way people do – by doing the work, observing what happens, and refining its understanding over time, with human guidance when needed. The underlying approach draws on research into how humans learn from narrative experience, using commercial foundation models as cognitive infrastructure. This system has broad application – our current focus is industrial process automation.

In our system, humans specify tasks in natural language and provide guidance when needed. The system discovers how the operational environment behaves, both through deliberate exploration and through adaptation during production. These are different kinds of knowledge and the system builds both.

Knowledge about the world accumulates over time. Base production rules and command vocabulary come first, provided by human operators. Learned facts about environment physics are discovered through exploration or production experience. Operational specifics emerge during execution, for example, a position on a robot workbench is locked, or a test fixture drifts. Knowledge is curated into the world model in real-time.

This is the practical difference: Training-based systems cannot operate until trained, while our system can start producing value immediately because proactive and reactive learning plus human escalation handling is a viable operating model from day one. Our system allows ad-hoc re-tasking in minutes, not days or weeks.

How Discovery Works

In validation testing, our system was given a sorting task in a work envelope but was not told that objects could be stacked. Through exploration, it discovered this on its own. When a pickup command didn't seem to move the object, the system noted its surprise, formed hypotheses, tested them, developed false theories, and eventually arrived at correct generalizations about stacking behavior.

Later, in production mode on a more complex task, the system encountered positions that unexpectedly rejected placements, analogous to fixture problems or locked positions in a real production floor work envelope. Without crashing or requiring reprogramming, it recorded the observation ("Some positions can be locked and will not accept placements"), adapted its approach, and completed the task using the remaining valid positions.

Every action includes a reasoning trace that operators can follow. When something goes wrong, they can see why. The system maintains hypotheses that are rationalized into generalizations that operators can review, edit, or delete. If the system learns something that is incorrect, operators can fix it directly rather than waiting for the retraining cycle to address it.

What This Changes

For manufacturers, the value is straightforward: the ability to say yes to work that matters without breaking processes, people, or margins. Revenue protection. Account expansion. Reduced chaos on the floor.

For integrators, this opens markets they've been forced to ignore. HMLV projects become viable because the system handles variability through learning rather than pre-programming. Reduced engineering risk, less custom development, and a recurring revenue opportunity from helping customers refine their knowledge over time.

For operators, this reduces cognitive load rather than replacing jobs. Fewer fire drills. Preservation of tribal knowledge. More predictable days.

Where We Are

The validation demonstrates the mechanism works: a system that discovers operational constraints it wasn't told about, reasons transparently about its actions, and maintains editable knowledge that operators control. Scaling to real production environments is the current development focus.

The open questions are about scale and robustness in messier real-world conditions. That's the work ahead.

But the core insight is validated: automation doesn't have to be trained before it's useful. It can learn while it works, building its world model through operation with human oversight. That difference means a radical change for HMLV economics – making it much easier for manufacturers to say “yes”.

This article describes research from Mossrake Group LLC on learning via episodic compression. Technical details are available on our website.

Version 1.1