

Manufacturing's Next AI Wave: Notes from the Frontier

I've spent the last few months in conversations with founders building AI for manufacturing, from pre-seed teams to growth-stage companies. I came at this as an operator at a contract electronics manufacturer, curious where AI moves the needle on the shop floor and on the P&L. I came away with a clear view: this is one of the most underrated spaces in AI right now, and the next wave is going to look very different from the current one.

Most of the AI-in-manufacturing conversation today is about robots. Robotics startups pulled in over **\$10 billion** in venture capital in 2025. **Mind Robotics**, a Rivian spinout, crossed \$1 billion in cumulative funding inside a year.

Robots are easy to film and easy to fund. That's why they get the headlines.

For most operators, though, robots are not the half of manufacturing AI that runs their day. That half lives in supplier inboxes, RFQ spreadsheets, BOMs with 300 line items, ERP screens, CAD files, and shop-floor dashboards. That's the half I want to talk about.

A Quick Map of the AI Manufacturing Space

The non-robotics AI manufacturing landscape sorts into a handful of problem areas.

Manufacturing intelligence and shop-floor visibility

The most mature subcategory. Augury (\$369M, Series F) leads on predictive maintenance. First Resonance (\$19M, Series A) is a manufacturing OS built for high-mix, low-volume electronics. Guidewheel (\$30M), Tulip (\$152M), and Sight Machine (\$87M) cover OEE, shop-floor apps, and industrial analytics. Zelos Cloud is the seed-stage entrant.

Visual inspection and quality

Instrumental (\$80M, Series C) is the most established play for PCB and electronics quality. Andrew Ng's Landing AI (\$57M, Series B) brings a data-centric approach for small datasets. At seed, Bucket Robotics is using CAD-trained models to deploy visual inspection without manual labeling.

RFQ and quoting

Paperless Parts (\$51M, Series B) has been working this problem for years. The seed wave (Korso, Soff, Arzana) is automating the full quote-to-order workflow.

Procurement and supply chain

The largest cluster of seed-stage activity. Mandel, Lumari, Comena, Vendra, and Cavela are each building AI agents on top of supplier emails. At Series A, Didero (\$37M, NYC) is bringing the same approach to enterprise scale.

Design, compliance, security, ESG

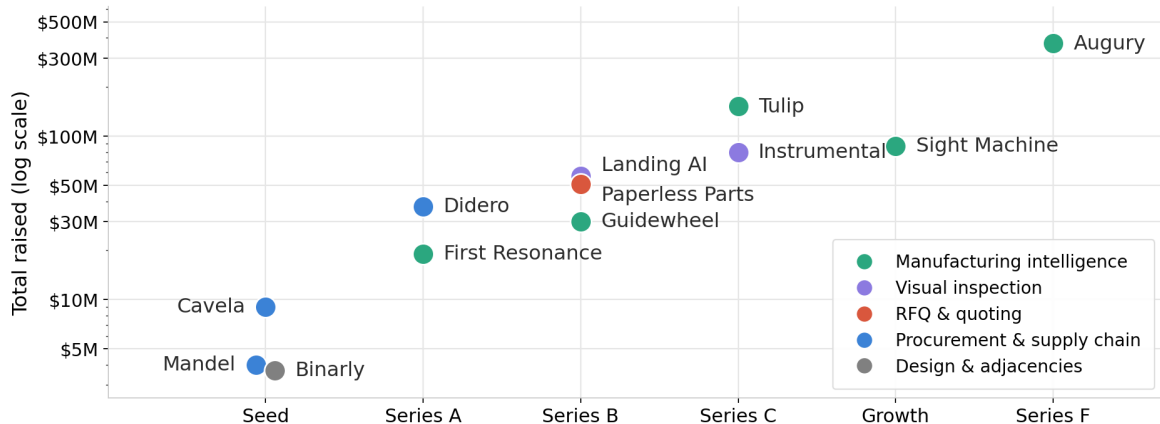
Smaller but real activity. Diode Computers and Prototyping.io in design automation. Noetic and Saphira in hardware compliance. Binarly in firmware security. Coolset in ESG reporting.

Whitespace

One area stood out as under-served: high-mix, low-volume production scheduling. Most MES tools are built for high-volume, repetitive runs. For an EMS shop juggling 50+ different jobs a week across SMT, wave-solder, and AOI lines, the scheduling layer is still mostly spreadsheets. No seed-stage AI startup is purpose-built for it today.

Plot these companies by stage and funding and the shape of the category jumps out: a dense seed-stage cluster on the left, a handful of growth-stage players up and to the right. The category is being defined right now.

Non-robotics AI in manufacturing — 25 companies by stage and funding



Seed stage · funding undisclosed (13 teams)

Zelos Cloud	Bucket Robotics	Korso	Soff	Arzana
Lumari	Comena	Vendra	Diode Computers	Prototyping.io
Noetic	Saphira	Coolset		

What I Heard Across the 'AI for Manufacturing' Conversations

Three patterns repeated in nearly every meeting. None of this is tied to any specific team. It's the shape of the category right now.

1. AI is being built on top of unstructured manufacturing data. Suppliers won't adopt portals, engineers won't redo their CAD workflow, and shop-floor operators have no interest in a new screen. So the dominant architectural choice is to read the data where it already lives. Inboxes and PDFs for procurement. CAD files and camera feeds for inspection. Machine-level signals for the shop floor. Same pattern across the category: read what's already there.
2. The workflows have converged. What separates teams now is configurability. Within each subcategory, the core loop is now standard. Procurement: ingest → distribute → extract → compare → push to ERP. Inspection: CAD → train → deploy → flag. Shop floor: sense → capture → surface → recommend. What separates teams is how easily a manufacturer can encode its own playbook: follow-up cadences, escalation rules, approval hierarchies, exception thresholds. Every shop runs its own playbook, and the better products are letting them keep it.
3. The customer counts are small, which is part of the opportunity. The teams I met range from zero customers to about a dozen. Even the well-funded ones are running pilots. I don't see that as a knock. It's the most exciting part of the category. The founders shipping today are writing the playbook for everyone who comes next.

“

*Automating tasks is the floor.
Augmenting human judgment is the ceiling.*

From "Manufacturing's Next AI Wave: Notes from the Frontier"

Four Takeaways for Founders Building Here

1. Automating tasks is the floor. Augmenting human judgment is the ceiling.

Today's products make existing manual workloads faster and cheaper. For a manufacturer where manual cognitive work is a meaningful share of opex, that's real money on the P&L. But faster execution of yesterday's workflow isn't the same as a fundamentally better business.

The bigger prize is expanding what an operator can *decide*. Surfacing patterns and possibilities that weren't visible before, so they can run categories and programs they couldn't run before. That's not "AI that does my job 10x faster." It's "AI that lets me do a job nobody could do at all."

2. The differentiator is data fusion, not the model.

The teams that define this category long-term won't be the ones with the cleverest agent architectures. They'll be the ones that fuse the customer's internal data (historical pricing, supplier scorecards, machine logs, quality history, BOMs) with meaningful external signals: public benchmarks, government data, geopolitical inputs, commodity trends, archival disruption data.

That fusion produces decisions the operator could not have reached alone. Anyone can wrap an LLM around a workflow. Very few can build the data graph that makes the model's output useful enough to act on.

3. Strong pitches start with the business outcome.

The pitches I came away most impressed by framed the product around a P&L line, not an architecture diagram. "We save \$X per shift on a \$100M EMS shop" is a different kind of sentence than "we use a 12-step pipeline with custom embeddings."

This matters beyond pitching. It shapes the product. Tech-first products optimize for elegance. Business-first products optimize for the customer's outcome. Over time, the second kind tends to win.

4. The best AI agents run in the background.

The use case I keep coming back to is the agent that works silently and only interrupts when it has something worth saying. In quality, that's an agent watching every line and flagging only the drifts worth attention. In scheduling, it's one that quietly re-sequences jobs around a delayed shipment. In sourcing, it scans new suppliers continuously and only surfaces when one beats your current vendor.

It's the opposite of chatbot-on-top-of-a-workflow. The better model is AI as a colleague who walks into your office only when something matters.

Closing Thoughts about Manufacturing's Next AI Wave

This category is being built right now. The current wave is laying the foundation: moving manual workloads off the page, building the data substrate, codifying the playbooks every manufacturer runs by hand today.

The next wave is the one I'm most excited about. Products that fuse internal and external data into decisions a single operator couldn't have reached alone. The agent layer underneath stays quiet until it has something worth saying. And the whole thing is built around the business outcome, not the architecture.

If you're a founder building in this space, or an operator at a manufacturing business thinking about where AI moves your P&L, I'd love to compare notes.