

“Global market will be won by whoever can unleash most reliable dairy decisions at the lowest integration cost” : Anand Mahurkar, Founder & CEO, Findability Sciences

14 July 2026 | Interviews

Why the next wave of industrial AI in dairy will be measured not by dashboards, but by decisions and outcomes; that is explored by Anand Mahurkar, Founder & CEO of Findability Sciences, in an exclusive interaction with AgroSpectrum



ANAND MAHURKAR
Founder & CEO, Findability Sciences

Discussing the launch of LactaAI, Mahurkar argues that dairy processors need AI systems capable of translating fragmented plant data into real-time operational actions that improve yield, reduce energy consumption, and accelerate decision-making. He explains how the platform's proprietary AI Factory architecture and industry-specific intelligence can unlock annual value ranging from Rs 2.35 crore to Rs 28.2 crore per plant, while integrating seamlessly with existing ERP, MES, SCADA, and plant systems. The conversation also explores AI adoption on the factory floor, the future of outcome-based industrial AI, and Findability Sciences' global ambitions as it takes LactaAI from India to major dairy markets worldwide.

Your claim that dairy plants “do not need more dashboards” challenges a crowded industrial analytics space. What fundamentally differentiates LactaAI from existing AI and Industry 4.0 solutions, and why hasn't this problem been solved effectively until now?

Dashboards describe. AI Factories decide. That is the difference, and it sits at the centre of everything LactaAI does.

After fifteen years and meaningful capital, the industrial analytics space has converged on a familiar pattern: collect telemetry, surface it on a dashboard, and leave the plant manager to figure out what to do with it. The result is what the industry now openly calls dashboard fatigue. Operators have more visibility than ever, and exactly the same decisions to make.

LactaAI starts from a different premise. We treat the dairy plant as an AI Factory, an environment that produces decisions the way a power plant produces electricity. Each decision is tied to a specific outcome: a yield gain, an energy unit saved, a

quality deviation prevented, a changeover shortened. The platform is closed-loop by design. It ingests operational data, computes a recommendation, routes it to the person who can act, and then measures what actually happened.

Two reasons the problem has not been solved before. The industrial analytics market has been dominated by horizontal Industry 4.0 platforms that are deep in operational technology plumbing but shallow in dairy physics and biochemistry. They digitise signals; they do not understand the difference between fat recovery in a separator and SNF loss in evaporation. The commercial models of legacy vendors reward dashboards over decisions, because a dashboard scales as a license while a decision requires ownership of an outcome. We accept that ownership.

Dairy plants generate massive volumes of data that often go unused. What are the core technical and organizational barriers that prevent this data from translating into decisions, and how does LactaAI overcome both?

Two layers of friction sit between data and decisions in dairy. Both are structural, not aspirational.

The technical layer first. Data lives across SCADA, MES, ERP, LIMS, IoT gateways, and laboratory notebooks. Each system has its own schema, its own time signature, and its own owner. Plant historians capture telemetry at one cadence, ERP captures transactions at another, LIMS captures test results at a third. Joining these into a unified picture of a single batch or a single shift is non-trivial. Add legacy PLCs on proprietary protocols, and a meaningful fraction of operational data never reaches a place where a model can reason about it.

The organizational layer is the harder one. Even when data is consolidated, recommendations stall. Operators do not trust outputs they cannot interrogate. Plant managers do not have a closed-loop feedback mechanism to verify whether an AI suggestion actually moved the metric it claimed. Ownership of the outcome remains diffuse between IT, OT, quality, and operations.

LactaAI addresses both. The platform sits on our I-CUPP architecture, which unifies the data layer before any model touches it, so the AI reasons over a coherent operational picture rather than fragments. On top of that, our Business Process Co-Pilot embeds each recommendation into the systems operators already use, with a transparent rationale, an accountable owner, and a measured outcome. The decision and its consequences are bonded.

Industrial environments are notoriously resistant to change. How do you ensure adoption among plant operators and decision-makers who may be skeptical of AI-driven recommendations, especially in high-risk production settings?

Operator adoption is earned, not assumed. Three principles guide how we earn it.

First, we deploy as a Co-Pilot, not an autopilot. Every recommendation comes with the reasoning behind it in plain plant language, the variables it considered, and a confidence band. The operator remains the decision-maker. This matters because in a high-risk production setting, the person closest to the asset must retain authority.

Second, we sequence the decisions. We begin with high-confidence, low-risk recommendations where the operator can verify the outcome quickly, such as CIP cycle optimization or boiler load scheduling. Trust accumulates. Only once a track record is established do we move to higher-leverage decisions like composition-aware standardization or fat-recovery setpoints.

Third, we localise. Our interfaces run in the languages operators actually speak. In India that means Marathi, Hindi, Gujarati, Kannada, and Tamil. The recommendation is useless if the operator has to translate it. Training materials and onboarding are built around the working vocabulary of shift supervisors and floor leads, not corporate slide decks.

The result is that operators move from skepticism to ownership within a few weeks of go-live, because they see the platform working with them rather than around them.

LactaAI integrates with legacy systems like ERP, MES, and SCADA without requiring replacement. What were the biggest engineering challenges in achieving this interoperability, and how do you handle data inconsistencies across systems?

Interoperability with legacy systems is the single hardest engineering problem in industrial AI. Dairy plants run a cocktail of vintages: SAP or Oracle ERP from the 2010s, Wonderware or Siemens SCADA, custom MES, LIMS from a half dozen vendors, and OEM-controlled PLCs that sometimes predate the people running them.

Our approach has three layers. At the bottom, a protocol adapter layer that speaks OPC-UA, Modbus, MQTT, and the vendor-specific protocols that matter in dairy. We read before we write. In the first phase of every deployment, we are a non-intrusive observer; we do not push setpoints back into the control system until governance and trust are in place.

In the middle, the unification layer, which is the U in I-CUPP. This is where data inconsistencies are resolved: entity resolution across systems that use different identifiers for the same batch or SKU, time synchronization across instruments with drift, unit normalization, and semantic harmonization so that what one system calls fat percentage and another calls FAT_PCT become the same field.

At the top, the decision layer, which is where the AI Factory runs. By the time models see the data, it is coherent and reliable.

The biggest engineering challenges were time-series at scale, particularly hot-cold tiering for plants generating millions of tag-seconds per day, and reconciling real-time streams with batch-loaded reference data without compounding errors. We solved them by treating the unification layer as a first-class product, not a connector marketplace.

Dairy processing varies widely—from fluid milk to complex whey derivatives. How scalable is LactaAI across different product categories, and what customization is required for each segment?

Dairy looks like one industry from the outside and a dozen sub-industries from the inside. Fluid milk, butter, ghee, paneer, cheese, milk powders, condensed milks, whey derivatives, infant nutrition, ice cream, fermented products. Each has its own process flow, its own quality regime, and its own economic drivers.

What scales across all of them is the I-CUPP architecture. The way we collect, unify, process, and present data is invariant to product category. What customises per category is what sits on top: the process model, the BOM derivation logic, and the decision packs that encode the levers that matter for that product.

We have built a library of pre-configured process models for the major dairy segments. For a fluid milk plant, the high-value decisions cluster around standardization, route optimization, and shelf-life prediction. For a milk powder plant, around evaporator and dryer thermal efficiency, particle size control, and lactose crystallization windows. For cheese and paneer, around culture management, syneresis control, and ageing. For whey, around membrane fouling and protein recovery.

The customization required for a new segment is the process model and the decision pack, not the platform. Typically a new segment takes us six to ten weeks from first engagement to first production decisions.

India is the launchpad, but dairy is a global industry. How do you see LactaAI competing internationally, particularly against established industrial AI players, and what markets are you targeting next?

Our near-term international roadmap targets the United States first, where our Chicago go-to-market is engaging the top-25 processors and where the consolidation of the industry into a small number of large producers makes the unit economics attractive. Australia and New Zealand follow, with channel work underway into the Fonterra ecosystem and the broader Oceania industry. Japan is open through our SoftBank relationship, with strong specialty dairy and functional nutrition demand. Latin America is the natural next step from our existing presence in adjacent agri-industrial categories.

The global market will not be won by whoever has the largest platform. It will be won by whoever can produce the most reliable dairy decisions at the lowest cost of integration.

--- Suchetana Choudhury (suchetana.choudhuri@agrospectrumindia.com)