

AI emerges as sugar's strategic moat

25 May 2026 | News

Speaking exclusively to AgroSpectrum, Guillermo Jose Medina LLarena, Chief Digital Architect (CDA) at Grupo Pantaleon outlines how data-driven mills are unlocking recovery gains, sustainability improvements and long-term operational advantages



Guillermo Jose Medina LLarena,
Chief Digital Architect (CDA),
Grupo Pantaleon

Speaking exclusively to AgroSpectrum, Guillermo Jose Medina LLarena, Chief Digital Architect (CDA) at Grupo Pantaleon outlines how data-driven mills are unlocking recovery gains, sustainability improvements and long-term operational advantages

In an exclusive interview with AgroSpectrum, Guillermo Jose Medina LLarena, Chief Digital Architect (CDA) at Grupo Pantaleon, argues that artificial intelligence has rapidly shifted from a digital advantage to a competitive necessity for the global sugar industry. He explains how AI-driven platforms such as the Stoma Suite are helping sugar producers unlock multimillion-dollar operational gains through predictive crop intelligence, mill optimisation and real-time decision-making across the value chain.

Medina notes that measurable improvements in sugar recovery, precision irrigation, logistics optimisation and sustainability metrics are fundamentally reshaping how mills manage profitability and climate risk. He further warns that within the next seven years, the industry will likely consolidate around producers that successfully embed AI into core operations, while laggards risk structural decline in an increasingly volatile agricultural economy.

Sugar has long operated on thin margins and high volatility. Is AI now a competitive necessity rather than a digital luxury for large producers?

It already crossed that line – most producers just haven't realized it yet.

For decades, sugar competed on scale, geography, and political relationships. Those advantages still matter, but they're no longer sufficient. GLP-1 drugs are structurally reducing sugar consumption. Corn-based ethanol is undercutting margins in

biofuels. Climate volatility is making historical yield models unreliable. These aren't cyclical pressures you can hedge through - they're permanent shifts in the competitive landscape.

What AI changes is the speed at which a producer can detect a problem, understand its root cause, and correct it. A mill running on intuition and weekly reports is operating blind compared to one receiving prescriptive recommendations every four hours based on 100+ process variables. The gap between those two operations will compound every harvest cycle. Within seven years, I believe the global sugar industry will consolidate decisively around producers who mastered this - and those who are still trying to catch up. It won't be a gradual transition. It will look like a cliff.

Stomata Labs claims measurable ROI, including double-digit efficiency gains. What specific operational levers are driving that value creation?

The honest answer is that the biggest lever isn't any single variable - it's the elimination of the lag between what's happening in the process and what the decision-maker knows about it.

In a traditional mill, by the time an anomaly in juice extraction or boiling crystallization is detected, reported, and acted upon, you've already lost hours of recovery. Our Global Recovery Optimizer analyzes the full process continuously and surfaces recommendations to operators in near real-time. In our deployments at Pantaleon's PSA and Monte Rosa mills, we saw sugar recovery move from 81 per cent to 83 per cent in a single zafra - with statistical significance that removes any doubt about causality (p-value of 1.09×10^{-10}). That half-percentage-point difference, sustained across a full season, is worth millions.

On the agronomic side, Stoma Sense combines satellite imagery with proprietary cloud-fill algorithms that give you continuous NDVI monitoring even in tropical cloud cover - which has historically made remote sensing unreliable in precisely the regions where sugar is grown. Knowing actual crop stress at the parcel level, before the cane arrives at the mill, changes the entire harvesting and milling schedule from reactive to planned.

Across a 2-million-ton facility, you estimate up to \$6.6 million in annual value unlock. How much of that comes from field optimization versus mill performance improvements?

We model approximately 60 per cent from mill optimization and 40 per cent from field intelligence - but that framing, while useful for budgeting conversations, obscures the more important point.

The \$6.6 million figure is conservative and is predicated on the two being connected. Field data that doesn't inform the mill schedule creates agronomic insights that never become operational decisions. Mill optimization without crop-stage predictive data is reactive by definition. The value unlocks because you close the loop: you know what's coming before it arrives, and you configure the mill accordingly.

The reason we use \$6.6 million for a 2-million-ton facility - rather than a larger number - is credibility. We only claim what we can demonstrate with auditable, statistically validated data. Our philosophy is to under-promise on the model and over-deliver on the harvest.

The sugar industry is deeply data-fragmented. What was the biggest challenge in integrating agronomic, operational, and commercial datasets into a unified AI-ready system?

The honest answer is that the industry spent years trying to solve this - and the challenge wasn't technical. The technology to integrate these datasets has existed for some time. The challenge was building the right data models: understanding which variables actually drive yield and recovery, which signals are noise, and how agronomic data and industrial process data need to be structured to speak to each other.

That took years of working inside real mills, with real operators, through real harvests. There's no shortcut. You can't simulate your way to those models in a lab. You learn them by being embedded in the operation - watching what a master boiler knows intuitively and figuring out how to encode that into a system that works even when he's not on shift.

What we built at Pantaleon - which has 175 years of operational history and data - gave us a foundation that others would take a decade to replicate from scratch. That institutional knowledge, now encoded in our models, is arguably as valuable as the software itself.

How does embedding predictive intelligence at the crop stage change risk management compared to traditional reactive mill-based optimization?

Traditional mill optimization is like reading the news. You find out what happened, you understand why, and you adjust for next time. That's still valuable but it's fundamentally reactive.

Predictive crop-stage intelligence is more like having a weather forecast with a confidence interval attached. You know that Block 14 is showing early stress signatures six weeks before harvest. You know that if you delay cutting by eight days, you capture a meaningful tonnage improvement. You know that three sections of cane are maturing simultaneously and will create a mill throughput bottleneck in week seven unless you resequence the harvest plan now.

That shift from reacting to problems to engineering around them before they materialize changes the risk profile of the entire operation. You're not just optimizing the mill. You're optimizing the system. In an industry where a single bad week during peak zafrá can wipe out a season's margin, that's not a marginal improvement. It's a structural advantage.

In emerging markets where digital maturity varies, how scalable is the Stoma Suite across different geographies and operational complexities?

This is where most AgriTech companies fail and where we've deliberately designed differently.

The typical enterprise AgriTech approach assumes a level of digital infrastructure, data discipline, and technical staff that simply doesn't exist in most mills outside of Western Europe and North America. The result is a sophisticated product that requires a sophisticated client to operate it, which limits the addressable market to the top tier of an already-small industry.

We built Stoma Suite to deliver value at the operator level, not just the analytics team. The interfaces are mobile-first. The recommendations are expressed in operational language, not data science language. A senior boiler operator doesn't need to understand the model he needs to trust the recommendation and act on it. We've achieved 70 per cent adoption rates on AI recommendations at our active deployments, which in industrial settings is exceptionally high. That number tells you more about usability than any demo would.

On the geographic side we are already deployed across Mexico, Guatemala, India, and have active pipeline in Brazil and Honduras. Each geography required calibration, but the core models transfer. The data fragmentation problem is universal. The physics of cane processing don't change by latitude.

Beyond efficiency gains, how does AI integration strengthen sustainability metrics particularly in water use, recovery rates, and carbon footprint?

Sustainability in sugar has historically been a compliance conversation. AI makes it an operational one which is where the real gains live.

Recovery rate is the clearest example. Every percentage point of additional sugar recovered from the same quantity of cane is a percentage point less cane you need to grow, irrigate, harvest, and transport to produce the same output. The environmental math is straightforward. Our recovery improvements at Pantaleon translate directly to lower land use intensity per ton of sugar produced.

Water is more complex but equally tractable. Stoma Sense's parcel-level monitoring enables precision irrigation scheduling based on actual crop stress rather than calendar-based rules. In water-scarce regions which increasingly describes every major sugar-producing geography that's both an economic and an existential capability.

On carbon, the biggest opportunity isn't in the mill it's in logistics. Harvest sequencing optimization, which our field intelligence enables, reduces the dead mileage and idle time that accounts for a surprisingly large share of cane operations' fuel consumption. These aren't soft sustainability claims. They're measurable, reportable, and increasingly what institutional investors and export buyers require.

Looking ahead five years, do you see AI-enabled intelligence becoming a defining moat in global sugar competitiveness?

Seven years, not five but yes, and more definitively than most in the industry are prepared to accept.

Here's the dynamic that people underestimate: AI models get better with data, and data accumulates with deployment. A producer who begins deploying today will have seven harvest cycles of model refinement by the time a competitor decides to start. The models we run today are meaningfully better than the ones we ran two seasons ago, because they've learned from two more seasons of real operational decisions and outcomes. That gap is not static it compounds.

What makes this different from previous technology cycles in agriculture is that the moat isn't hardware or capital it's learning. You can buy the same sensors, hire the same data scientists, license the same cloud infrastructure. What you cannot buy is the accumulated harvest-by-harvest calibration of models that have learned what a specific mill, in a specific climate, with a specific variety of cane, does under a thousand different conditions. That takes time. It takes patience. And it takes a willingness to embed deeply in operations rather than sell software from the outside.

The producers who treat AI as a long-term operational capability not a procurement decision will look back in seven years and understand that this was the moment the industry permanently separated into two groups. We're helping the right clients make sure they're in the right group.

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