

Practical AI in Tortilla Manufacturing

What's Working, What's Hype,
and How to Evaluate It

TIA Convention 2026 | Educational Session



Why AI Is on Every Plant Manager's Radar

The Production Floor Reality

Lines running 200+ units per minute outpace human consistency

Operator fatigue drives quality drift over long shifts

Defect root causes are hard to isolate without data

Customer returns erode margins before you know there's a problem

The AI Promise vs. Reality

Vendors claim AI will revolutionize your plant overnight.

The truth is more nuanced. Some AI applications are delivering measurable ROI right now. Others are years away from practical deployment.

This session separates what works from what doesn't, and gives you a framework for making smart decisions.

No vendor pitches. No hype. Just practical knowledge you can take back to your plant.

What We'll Cover Today



Where AI Delivers ROI Today

Inspection, process feedback, production analytics

1



Where AI Falls Short (For Now)

What's genuinely promising vs. pure vaporware

2



The Vendor Evaluation Checklist

10 questions to ask before signing anything

3



What's Realistic in 2-3 Years

An honest look at where this technology is heading

4

First, Let's Define Our Terms

Not everything labeled "AI" is actually AI. Understanding the difference matters for your purchasing decisions.

Classical Computer Vision

Best for: Geometric measurements (diameter, area, shape)

How it works: Fixed rules, threshold-based decisions

Strength: Deterministic, sub-pixel accuracy, predictable

Limitation: Struggles with ambiguous defects, color variation

AI / Machine Learning Vision

Best for: Pattern recognition, defect classification, trend analysis

How it works: Learns from production data, adapts over time

Strength: Handles variation, improves with volume, no recalibration

Limitation: Requires training data, less transparent decisions

The best systems combine both: classical CV for precision measurements, AI for pattern recognition and statistical insight.

WHAT'S WORKING

Automated Quality Inspection

30+

Quality parameters
measured per unit

200+

Units inspected
per minute

100%

Consistency
no fatigue, no drift

What Modern Inspection Actually Measures

- ✓ Size & diameter consistency
- ✓ Color & browning patterns
- ✓ Spotting & burn marks
- ✓ Shape & roundness
- ✓ Thickness uniformity
- ✓ Tears, folds, structural defects

Available today from multiple vendors. Key differentiators: how the system learns, what data you own, and total cost of ownership.

Closed-Loop Process Feedback

When inspection data feeds back into production equipment, you move from catching problems to preventing them.

1

Inspect After Oven

Secondary inspection point with production-specific thresholds catches variation before it compounds.



2

AI Interprets Trends

System ranks defect causes by frequency (Pareto) and translates data into operator-readable guidance.



3

Operators Adjust

Clear, actionable recommendations: which oven, which press, which parameter to change and by how much.

Greatest impact during recipe changeover — The initial setup period for a new recipe is where waste is highest and where real-time feedback cuts losses fastest.

WHAT'S WORKING

Production Analytics for Management

Turn inspection data into business intelligence

Shift-by-shift performance

Compare yield, waste rates, and defect distribution across shifts, lines, and operators

Defect trending over time

Spot emerging quality issues before they become customer complaints

Recipe performance tracking

Quantify which recipes run cleanest and which generate the most waste

Compliance documentation

Automated reporting for customer audits and food safety requirements

Metrics That Matter

- ✓ Reject rate by defect type
- ✓ Throughput (units per minute)
- ✓ Size distribution curves
- ✓ Waste percentage by line
- ✓ Yield per batch/recipe
- ✓ Uptime vs. downtime correlation

Ask any vendor: Can I export my data via API? If the answer is no, you're renting insights you can never own.

The Hidden Cost of Inconsistent QC

False Rejects = Wasted Product

Every good tortilla flagged as defective is raw material, energy, and labor thrown away. A 2% false reject rate on a line producing 15,000 units/hour means 300 perfectly good tortillas wasted every hour.

That's 2,400 units per shift.

Missed Defects = Customer Returns

Every bad tortilla that reaches a customer costs 10-50x more than catching it on the line. Returns, chargebacks, lost shelf space, and damaged relationships with retailers add up fast.

One major return can erase a quarter's margin.

The Real Question Isn't "Should We Automate QC?"

It's how to do it without overspending on technology that doesn't fit your operation. The approaches below range from \$0 (better training) to \$250K+ (full vision systems). The right answer depends on your volume, your products, and your current pain points.

Comparing QC Approaches: Honest Trade-Offs

Approach	Upfront Cost	Speed	Consistency	Data Output	Best For
Manual Human Inspection	Low	Limited	Variable	None	Low-volume, specialty products
Basic Optical / Laser	\$150-250K	High	Good (needs recalibration)	Limited	High-volume, simple criteria
AI Vision (Purchase)	\$170-250K	High	High (self-adapting)	Rich API/Export	High-volume, complex products
AI Vision (Subscription)	\$0 (monthly fee)	High	High (self-adapting)	Rich API/Export	Any volume, flexible commitment

There is no single "right" answer. The best approach depends on your line speed, product complexity, labor availability, and budget.

Regardless of approach, the Pareto principle applies: rank defect causes by their share of total defects and fix the biggest contributors first.

Promising Ideas That Aren't Ready Yet

These concepts are real, but the gap between demo and deployment is wider than most vendors will tell you.



Fully Automated Recipe Generation

Reality: Equipment varies wildly between plants. Ovens, mixers, and presses have different designs and automation levels. AI cannot yet account for this heterogeneity reliably.



Complete Digital Twin of Production

Reality: Modeling dough formation and intermediate product states with enough precision to simulate final output is not yet achievable. Raw material variability alone makes this extremely difficult.



Zero-Human Production Floor

Reality: Ingredient properties change batch to batch. Human expertise in reading dough, adjusting feel, and making nuanced decisions remains essential. AI augments this; it doesn't replace it.

If a vendor can't show you this working in a production environment (not a lab), it's not ready.

How to Spot AI Hype in Vendor Pitches

Red Flags

- ❌ "Our AI does everything" with no specifics
- ❌ No reference customers in food production
- ❌ Demo only works on perfect samples in a lab
- ❌ Can't explain what happens when the system is wrong
- ❌ Data stays in their cloud with no export option
- ❌ ROI projections with no methodology shown

Green Flags

- ✅ Transparent about what AI does vs. classical CV
- ✅ Named production references you can call
- ✅ Offers paid pilot before full commitment
- ✅ You own your data and can export via API
- ✅ Explains limitations honestly
- ✅ Pricing model that aligns with your risk tolerance

A vendor who admits their technology has limitations is more trustworthy than one who claims it has none.

TAKE THIS HOME

10 Questions to Ask Any AI Vendor

- 1 What specifically does your AI do vs. classical algorithms?
- 2 Can I visit a production facility where this runs 24/7?
- 3 What happens when the system makes a wrong call?
- 4 Do I own my data? Can I export it via API?
- 5 What does your system need from my existing infrastructure?
- 6 How does the system handle new products or recipe changes?
- 7 What's the total cost of ownership over 3 years?
- 8 What does your maintenance and support model look like?
- 9 Can I run a paid pilot before committing?
- 10 What are the known limitations of your system?

Any vendor worth working with will welcome these questions. The ones who dodge them are telling you something.

What's Realistic in the Next 2-3 Years

NOW

- ✓ Visual inspection at line speed
- ✓ Defect classification and trending
- ✓ Operator guidance from inspection data
- ✓ Production reporting and analytics

1-2 YEARS

- ⚠ Tighter closed-loop oven/press feedback
- ⚠ Predictive maintenance from sensor fusion
- ⚠ Cross-plant benchmarking via cloud analytics
- ⚠ AI-assisted recipe optimization support

3-5 YEARS

- ❓ Partial digital twins for specific processes
- ❓ Semi-automated recipe parameter suggestion
- ❓ Integrated multi-vendor data platforms
- ❓ AI-assisted supply chain quality tracking

Focus your investment on what works today. Keep an eye on what's emerging. Be skeptical of anything that requires "perfect" conditions to function.

The Practical Evaluation Framework

Does it solve a real problem?

If you can't name the specific production issue it addresses and the cost of that issue today, you don't need it yet.

Does it work with your equipment?

If the solution requires perfect standardization or equipment you don't have, it's not practical for your plant.

Can you measure the impact?

Demand specific, verifiable metrics: waste reduction, throughput improvement, labor savings. If they can't quantify it, neither can you.

Does it respect your expertise?

The best AI amplifies what your operators already know. If a system claims to eliminate the need for process knowledge, walk away.

AI should not replace process knowledge. It should amplify it.

Questions & Discussion

[Speaker Name & Title]

[Contact Information]

TIA Convention 2026 | Las Vegas