IPPTA ZONAL SEMINAR • TECHNICAL PAPER

Industrial AI for Future-Ready Paper Mills

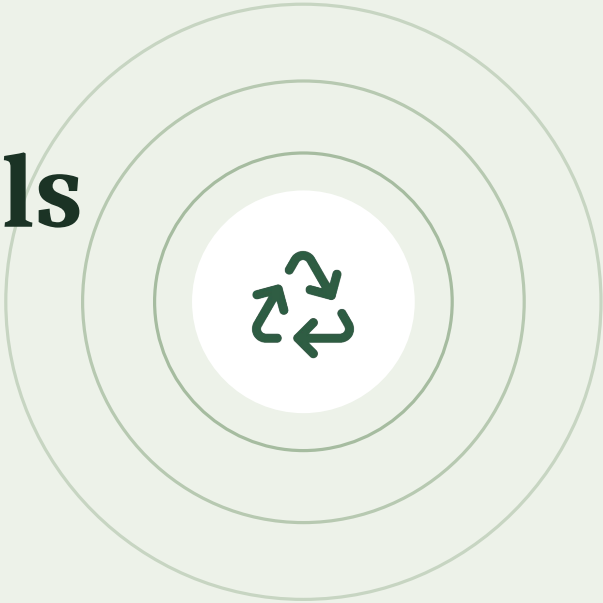

One operator-facing intelligence layer for mid-sized recycled-fibre mills.

Sneh Patel

COO, JMC Paper Tech USA LLC
Balaji JMC Paper Mill, Ciudad Juárez, Mexico

sneh@jmcengineers.com

THE PRESSURE · WHY NOW

Three pressures are arriving together

Policy

Compliance needs auditable data.

Sustainability

Customers buy verified proof.

Industrial AI

The tech is production-ready.

THE COMMON THREAD

Granular · time-stamped · machine-verified data

The backbone: mid-sized recycled-fibre mills

50–300
tonnes per day

1–2
machines per mill

100s
of mills nationwide

Same obligations as the giants — a fraction of the capital.

Tight margins, competing with larger producers and imports — with far less capital to digitalise.

THE PRESSURE · FORCE 1 — POLICY

Compliance now runs on machine-verified data

EPR
Auditable fibre & recycled-content traceability.

BRSR
ESG disclosure to SEBI.

PAT & net-zero
Emission-intensity targets.

Effluent norms
Tighter discharge limits.

One requirement: machine-verified data.

Customers now buy verified performance

- **Verified footprint reporting**
  written into supplier contracts
- **Specific water below 40 m^3/t**
  a hard purchasing threshold
- **Rising renewable-energy share**
  expected year on year

Measure it — or lose access

Prove it from process data — or lose premium customers.

Sustainability is now a purchasing decision.

Industrial AI has crossed into production

Falling data cost	Time-series ingestion is cheap at scale.
Open connectivity	OPC-UA and MQTT connect any equipment.
Production-ready models	Reliable for forecasting and soft sensing.
Standard components	Edge, databases and UIs — off the shelf.

So why aren't mid-sized mills using it?

Because the systems don't talk

DCS | QCS | ERP | LAB | MAINT

Five systems — disconnected. Nothing crosses within a shift.

Reactive decisions

Off-spec found only after the lab.

Energy: a black box

Per-tonne energy invisible live.

Knowledge leaves

Intuition gone at retirement.

The barrier isn't the algorithms — it's structural.

THE RESPONSE

One intelligence layer over the whole mill

4 **Decision & Insight**
dashboards · alerts · reports

3 **AI & Analytics**
forecasting · maintenance · energy

2 **Data Platform**
time-series store · context model

1 **Connectivity & Edge**
OPC-UA · MQTT · edge · gateways

MillMind — an approach, not a dashboard.

- **Real-time**
- **Predictive**
- **Multilingual**
- **Audit-ready**

Signals rise — insight returns to the operator.

Two choices carry the design

The time-series database is first-class

Mill data is high-frequency and trend-dominated.

100k+
tag updates/sec

30 days
full resolution

Context links tags to meaning

A vibration rise reads as:

> *“a felt-roll bearing event during a grade change”*

Data becomes a decision when it carries meaning.

A soft sensor, not an advisory

PROCESS TAGS

refining · freeness · vacuum · steam · dryer · ash

→

Gradient-boosted model
per grade & furnish

→

FORECAST, MINUTES AHEAD

burst · RCT · CMT · weight · moisture

Spec crossed → a specific corrective recommendation.

THE VALIDATION LOOP

- **Every reel scored**
- **Drift → re-calibrate**
- **An instrument, not advice**

The forecast anticipates the lab

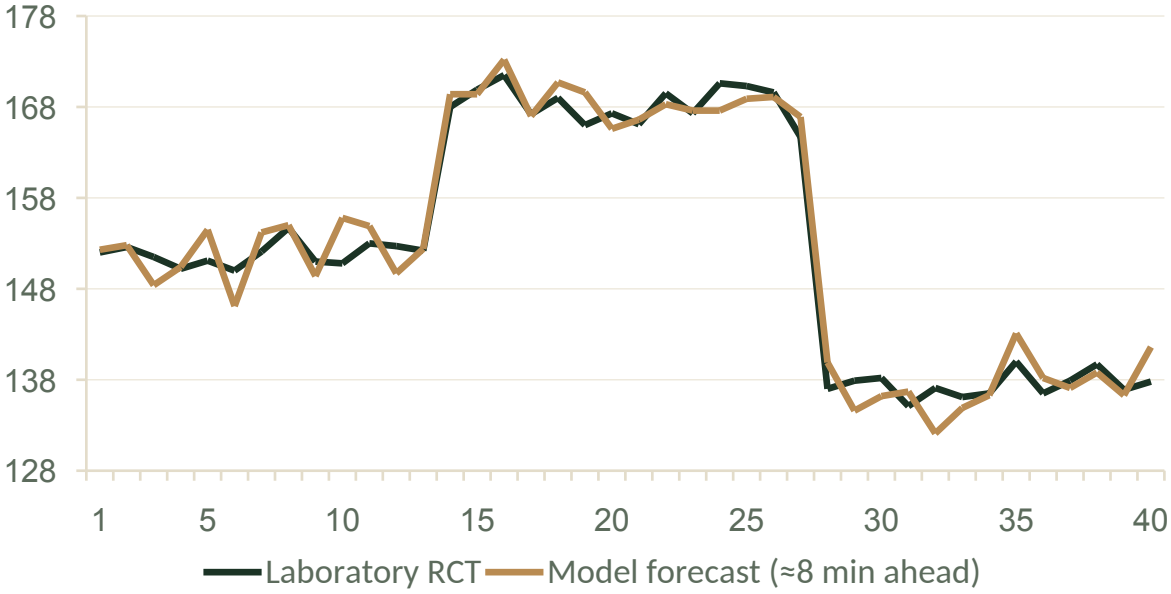

1.8 N
mean absolute error

0.97
R^2 against the lab

≈8 min
forecast lead time

40 reels
across three grades

Steps are grade changes — and it holds across all five properties, inside spec bands.

Failures caught days to weeks early

SIGNALS WATCHED

- Vibration
- Bearing temp
- Motor current
- Vacuum
- Dryer load

ASSET CLASSES COVERED

- Section drives & gearboxes
- Vacuum systems
- Refining lines
- Press bearings & rolls
- Dryer bearings & condensate
- Pulpers & screens

Failures move into planned windows.

The biggest cost is the biggest emitter

35–45%

of variable cost: steam, power, water, chemicals

…and the largest Scope 1 & 2 source.

One lever: margin and footprint.

Steam loop
vs grade baseline

Power loop
vs design curve

Water loop
flags creep

Chemical loop
tied to furnish

THE EVIDENCE · ENERGY — RESULT

≈6% less steam in sixteen weeks

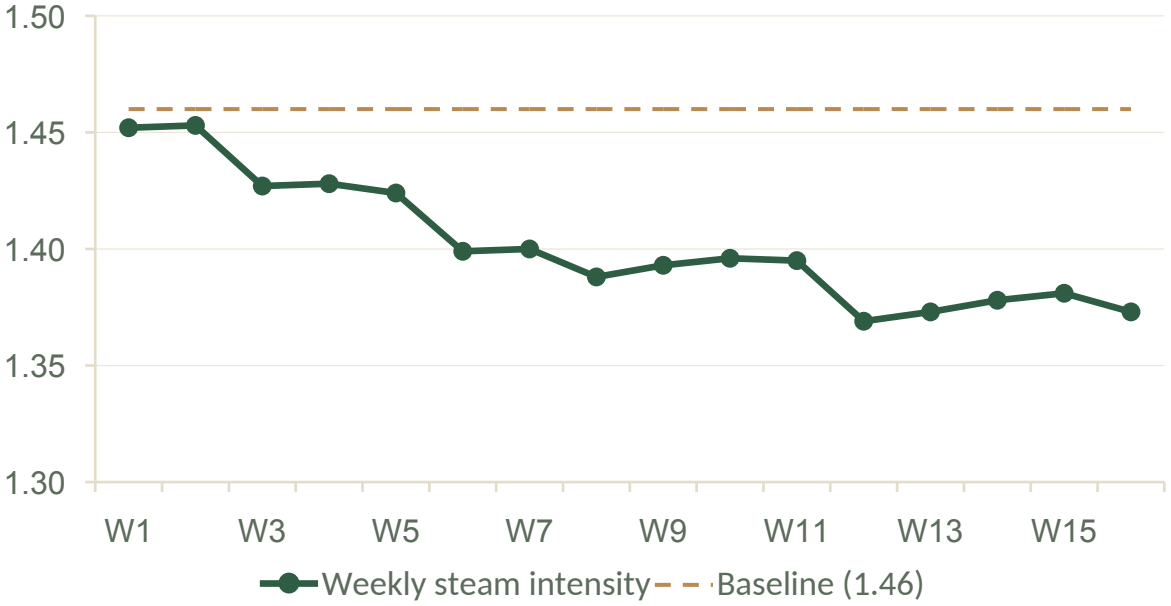

−6%
vs 1.46 t/t baseline

16 weeks
still trending down

Zero capex
tuning, not capital spend

Weekly dryer-steam intensity, t steam / t paper.

THE EVIDENCE · EARLY DEPLOYMENT

Measurable impact — during ramp-up

10–20%
lower unplanned downtime

3–7%
less steam per tonne

30–50%
faster grade changes

Always on
knowledge capture

Payback ≈ 12–18 months
typical 200 TPD mill

Feb 2026: 2,431 t at ≈93.5% saleable, still ramping

Indicative ranges from one early deployment — daily use reaches the top of the range.

Compliance falls out as a by-product

EPR	→	Lot-level fibre traceability, furnish to reel
BRSR	→	ESG KPIs with a full audit trail
Scope 1 & 2	→	Fuel & energy accounting, source-attributed
Effluent	→	Continuous pH, BOD, COD, TSS with alerts
Water	→	Per-tonne fresh-water draw & recycle
PAT / carbon	→	Energy baselines with provenance

Machine-derived · time-stamped · traceable — audit-ready.

THE PAYOFF · HONEST APPRAISAL

What we can — and can't yet — claim

THE LIMITS

- One greenfield site — ranges need wider proof.
- Value needs follow-through — it's socio-technical.

WHERE IT GOES NEXT

- **A plain-language assistant over mill history.**
- **Closed-loop control in validated envelopes.**

Both extend the same data fabric — neither replaces it.

Three pressures. One machine-verifiable data fabric.

Less downtime & steam

Faster grade changes

Knowledge that stays

Audit-ready by design

Payback ≈ 12–18 months. Build now — hold the advantage in cost, compliance and trust.

Thank you

Sneh Patel · sneh@jmcengineers.com