IPPTA ZONAL SEMINAR-I 2026 I JAIPUR I 10–11 JULY 2026

From Reactive Control to Autonomous Operations

Applying AI, Machine Learning and Digital Twins for Predictive Maintenance, Quality Stabilisation and Process Simulation in Paper Manufacturing

Ajay Jindal
EVP – Strategy & Partnership, Decimal Point Analytics

Laurent Bourrouilhou
Partner Manager India, Braincube SAS • 35 years in Industrial AI

Tracks: AI/ML for Predictive Maintenance & Quality Stabilisation | Digital Twins & Data-Driven Modelling

The Problem

A Complex System That Will Not Hold Still

A modern paper machine processes hundreds of interacting variables in real time - furnish, headbox pressure, nip loads, steam pressures and quality setpoints.

Any one variable can drift. When several drift together - during grade changes, furnish switches or seasonal humidity shifts – **the system becomes too complex to steer optimally.**

Why the conventional toolkit reaches its ceiling

Static in a dynamic world

DCS, QCS and SPC hold targets and raise alarms - but a setpoint optimal at the start of a run is no longer optimal 30 minutes later when the fibre blend shifts or the felt loads up.

Data exists but is siloed

Process, quality, lab and manual records sit in separate systems, poorly time-stamped and rarely integrated into one analytical layer. The data is generated - it is not yet use

Reaction, not anticipation

Correction happens minutes-to-hours after a deviation is detected. By then broke, off-spec reels or a break have already occurred.

The Cost of Operating Static Control

The financial consequence of static control in a dynamic process is substantial and measurable across four dimensions.

USD 3k–10k
per paper break

A mill averaging 5 breaks/day loses USD 5–15 million a year - in lost production, wasted fibre, re-threading and energy.

3–8%
of production value

Off-spec reels, basis-weight variation, moisture stripes and coating defects generate broke, rework and rejections.

10–20%
excess energy

Steam and electricity are 25–35% of operating cost; dryer inefficiency and over-refining push consumption above the achievable minimum.

Unplanned
downtime

Felt, roll and bearing failures that could have been predicted and planned drain maintenance budgets and wreck schedules.

Why This Matters Most for Indian Mills

Indian operating conditions amplify the variability challenge - the framework must begin from these realities, not from idealised virgin-fibre conditions.

01

Furnish heterogeneity

25–30% of production depends on imported wastepaper; agro-residue mills face shift-to-shift swings in fibre quality.

02

Operator skill variability

Mills vary widely in scale and workforce; data literacy differs too, favouring AI-assisted guidance over AI replacement.

03

Maintenance maturity

Condition-based maintenance is still rare, causing over- and under-maintenance across assets.

04

Data availability

Most mills have DCS and QCS, but data is siloed and rarely integrated with lab or MES records.

Where Variability Enters the Machine

A standard fourdrinier machine - and the primary variability entry points an AI model must track and compensate for at each section.

Stock Prep

- Pulper
- Blending
- Refiner
- Chest

Variability

- Furnish
- Freeness
- Consistency

Forming

- Headbox
- Fourdrinier
- Wire table
- Couch roll

Variability

- Jet / wire
- Drainage
- Retention

Press

- Press nips
- Felts
- Doctors

Variability

- Nip pressure
- Felt load
- Sheet tension

Dryer

- Steam cans
- Humidity
- Steam profile

Variability

- Steam pressure
- Temp profile
- Hood humidity

Finishing

- Calendering
- Reeling
- Winding

Variability

- Caliper
- Moisture
- Tension

Conceptual Framework

From Process Control to Process Steering

The difference between control and steering is a thermostat versus a skilled driver. One reacts after the fact; the other anticipates and adjusts before the car leaves the lane.

	Process Control (today)	Process Steering (target)
Setpoints	Fixed targets set by recipe	Dynamic targets recalculated continuously
Response time	Minutes-to-hours after a deviation	Real-time, before the deviation develops
Knowledge source	Operator experience & historical recipes	AI working on the mill's own contextualised data
Output to operator	Alarm: 'You are out of spec'	Guidance: 'Hold steam at 4.4–4.6 bar to stay stable'
Equipment context	Settings unchanged as assets age	Recommendations account for asset age & maintenance

Conceptual Framework

The Five-Level Maturity Model

The path from traditional control to autonomous operations is a progression - each level builds on the last and delivers measurable value independently.

LEVEL 1
Reactive Control
DCS/QCS holds

Baseline

LEVEL 2
Descriptive Analytics
Historical dashboards, OEE and SPC charts

5–10%

LEVEL 3
Predictive Models
AI detects early warning signatures of breaks, drift, wear

15–30%

LEVEL 4
Prescriptive Guidance
AI calculates optimal setpoint windows; real-time recommendations

25–45%

LEVEL 5
Autonomous Steering
Closed-loop: system adjusts setpoints within approved bounds

40–60%

Indicative payback

Target: Level 3–4 in 12–18 months, path to Level 5.

A Five-Layer Intelligence Stack

Every steering capability rests on one foundation: a clean, contextualised representation of the process - built from data the mill already generates.

01

Operator Digital Assistants / UI
Real-time guidance | setpoint recommendations | golden-run reporting

02

AI Scoring Engine
Variable scoring | best-range identification | rule discovery

03

Process Digital Twin + Product-Centric Model
Contextualised data | grade-tagged | lag-corrected | single source of truth

04

Data Integration (edge + cloud)
DCS | QCS | MES | LIMS | manual inputs | utility meters

05

Shop Floor / Sensor Layer
Paper-machine sensors | drives | actuators | lab instruments

Architecture

The Process Digital Twin

A continuously updated, contextualised replica of the machine - the single source of truth for analysis, not a 3D model or SCADA replacement.

Three traits of real paper-machine data that defeat naïve analysis - and how the Twin fixes each

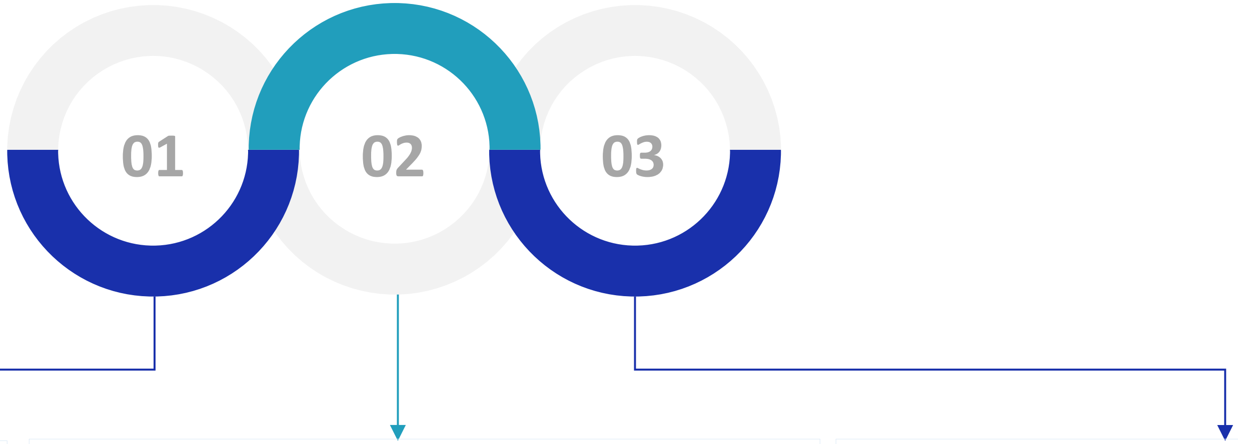

Time lags

A refiner change may not reach tensile strength for 20–40 minutes. The Twin maps these lags so it learns real cause-and-effect, not coincidence.

Grade transitions

The 30–60 minutes around a grade change is neither old grade nor new. The Twin tags these periods so analysis isn't skewed.

Sensor noise & outliers

Pearson-style correlation breaks with a single bad reading. Probabilistic scoring is inherently robust to outliers and non-normal data - reliable even on raw data.

Organising Data Around the Product, Not the Clock

Conventional - time-centric

A historian records each sensor value at each timestamp. It tells you how the machine was set at a given moment - necessary, but not enough to understand quality.

Product-centric model

The same data, reorganised around each unit of production - each reel or batch - and the full set of conditions that produced it. Good and bad outcomes trace directly back to the settings that caused them.

"What was the temperature at 14:32?"

"What makes a good reel?"

Time series show how machines are set; product data shows how products are made.

Every product instance links to the conditions that produced it, making quality traceable.

Architecture

A Scoring Method Built for Real Industrial Data

Real plant data is non-normal, full of outliers and mixes continuous with discrete variables - conditions that invalidate classical Pearson regression.

	PEARSON R^2	RANGE-SCORING ENGINE
Output	Regression equation (no range)	Best impacting range of the variable
Outlier sensitivity	One outlier destroys the score	Inherently robust to outliers
Non-normal data	Unreliable	Reliable across all distributions
Discrete variables	Cannot handle	Native support for binary / categorical
Threshold effects	Undetected	Identifies threshold & sweet-spot responses

A score of 84 means the variable correctly separates 84% of good output points from bad ones - **and the method simultaneously identifies the specific input range where that separation is maximised.** That gives the operator an actionable setpoint target, not merely an alarm.

AI Application 1

Break Reduction: Early-Warning Signals Across the Machine

Breaks are costly and usually preceded by recognisable change. The method finds stable-running ranges from mill history and guides operators back as conditions drift.

Wet end

Drainage variability, headbox consistency fluctuation, white-water turbidity changes

Press section

Nip-load asymmetry, felt-permeability decline, sheet-moisture variance widening

Dryer section

Steam-pressure hunting, hood-humidity spikes, sheet-tension oscillation

QCS

Basis-weight variance increasing, moisture stripes forming, caliper excursions

GUIDANCE THE OPERATOR SEES - A PRIORITISED INSTRUCTION, NOT A BLACK-BOX ALARM

“Break risk elevated.

Primary driver: **headbox consistency 3.8%** (recommended 3.2–3.5%).

Secondary driver: **press steam pressure 4.1 bar** (recommended 4.4–4.6 bar).”

Break Reduction at Scale

Corrugating-medium / kraft-liner mill - recurrent breaks averaging 5–7 per day, root causes dispersed across wet end, press and dryer; manual analysis unable to isolate primary drivers.

68%	+5%	~1,850	< 4 mo
reduction in paper breaks	throughput increase	break-free hours / year	time to value

Method

A Process Digital Twin was built from historical DCS and QCS data; the scoring engine identified operating ranges for stable running, delivered to operators as updated production standards.

Key finding

The most impactful break drivers were cross-sectional combinations of variables across different machine sections -relationships extremely difficult to isolate manually, but surfaced directly from the data.

Quality Stabilisation

Holding finished-paper properties within spec through the run, across grade changes - by identifying conditions that separate on-spec from off-spec.

Basis weight (GSM)	Headbox & stock conditions associated with on-target GSM
Moisture profile	Dryer-steam conditions giving uniform target moisture across the reel
Moisture profile	Press & calender conditions for target caliper at given furnish
Tensile / tear	Refining, pressing & drying conditions behind lab strength results
Ash content	Filler-addition conditions for target ash at current freeness & retention

The essence of steering: not more alarms, but a small number of clear, prioritised, in-range targets the operator can hold.

Documented Result

32% → 50%+

Good-batch incidence

Isolating the two process temperatures that most strongly drive quality - and holding both in range - even at high line speed.

DOCUMENTED EVIDENCE

Runnability Recovery - A Phased Value Model

Newsprint mill - averaging 7 breaks/day and running below design speed. Calendar-based maintenance couldn't find root causes; each phase generated ROI that funded the next.

01 PHASE Optimised

New setpoints identified for the highest-impact variables. 2 breaks/day eliminated immediately.

USD 1.2M saved (Yr 1)

02 PHASE Amplified

Model coverage expanded across more sections - a further 4.5 breaks/day removed.

USD 4M total savings

03 PHASE Transformed

Machine speed ramped 12% to meet demand - no additional capital expenditure.

USD 40M added sales

7 → 0.5 breaks/day equivalent to roughly USD 40 million in recovered annual sales - achieved by starting with the most impactful variables and demonstrating value quickly.

Pre-Qualify the Business Case Before Any Capital

A structured audit of 12–24 months of DCS, QCS and production data (4–6 weeks) returns a ranked list of value pools - shown here for a 200 TPD mill.

	Break reduction	50–68% from baseline	₹3.5–6.0 Cr
	Throughput gain	3–8% from runnability	₹4.0–10.0 Cr
	Energy saving	8–15% steam & power	₹2.5–5.0 Cr
	Chemical / fibre	5–10% broke & off-spec	₹1.5–3.0 Cr
	Maintenance	20–30% unplanned downtime	₹1.0–2.5 Cr
Total potential (illustrative)		—	**₹12.5–26.5 Cr / year**

ROI PROFILE

Break-even
3–6 months

Year 1 ROI
2.5x – 4x

Year 3 ROI
5x – 8x

Payback < 4 mo
96% of cases

Do More With What You Have - No CAPEX Required

Operators build “safety margins” into every step to protect KPIs. Together, they hide productivity pockets the AI can surface with no capital spend.

Chemical over-dosing	A little extra bleaching agent, a little extra retention aid - “just to be safe” on whiteness and formation. The AI identifies the true minimum-effective dose for current furnish and conditions, cutting chemical bills 5–10% with no loss of spec.
Over-drying the sheet	Operators push steam higher than needed so moisture never breaches the upper spec. Every percentage point below target is wasted gas and wasted basis weight. Holding moisture inside the real window saves 8–15% steam.
Over-refining the pulp	Extra refining energy is added to guarantee strength - at the cost of drainage, runnability and electricity. The scoring engine locates the refining range that hits strength targets at the lowest specific energy, often 5–12% kWh/t lower.
More recycled fibre, safely	Conservative recycled-furnish ratios protect strength and runnability -but leave money on the table. The Twin shows exactly when the process tolerates a higher recycled share, lifting it grade-by-grade without breaks or rejects.

Every Operational Gain Is a Sustainability Gain

Fewer breaks, tighter setpoints and better maintenance immediately cut resource use - a co-equal outcome, not a secondary benefit.

Lever	Metric	Reduction
Break reduction	Specific fibre consumption	3–8% ↓
Break reduction	Specific energy (GJ/t)	4–10% ↓
Dryer optimisation	Steam consumption (GJ/t)	8–15% ↓
Dryer optimisation	Direct CO_2 emissions	6–12% ↓
Chemical optimisation	COD load to effluent	5–10% ↓
Refiner optimisation	Electrical energy (kWh/t)	5–12% ↓

Aligned with India's Regulatory Agenda

For Indian mills, these gains are increasingly mandated - AI-driven optimisation is among the most cost-effective pathways to compliance.

01 BEE PAT Scheme

Gate-to-gate Specific Energy Consumption targets for designated paper mills - penalties for non-compliance, tradeable energy-saving certificates for over-achievement. Continuous optimisation is a direct route to meeting SEC targets.

02 India's NDC

A commitment to cut the emissions intensity of GDP by 45% by 2030. AI-driven mills can quantify and report per-tonne emissions reduction with precision - a growing advantage as carbon disclosure becomes standard.

03 ISO 50001

The continuous monitoring, data-driven optimisation and management reporting built into the Digital Twin and operator dashboards provide exactly the infrastructure the energy-management standard requires.

Challenges and Lessons from Real Deployments

Honest acknowledgement of what arises in practice - each challenge can be anticipated and mitigated.

01

Data quality

15–30% of sensor data shows calibration drift, stuck sensors or gaps. A mandatory audit during the Twin build fixes most issues in 4–8 weeks - value that extends beyond the AI project.

02

Operator adoption

Experienced operators may see AI as a challenge to their expertise. The fix is design: 'here is what the data says - you decide.' Including them in validation turns them into champions.

03

Integration complexity

OT-to-platform connectivity involves proprietary protocols, IT/OT segmentation and cybersecurity. Connectors for major DCS vendors and both edge and cloud architectures de-risk this.

04

Equipment ageing

Models drift as assets age and furnish shifts. The Twin tracks equipment age and time since maintenance, so guidance stays relevant as the machine evolves.

05

Pilot-to-scale gap

Many projects shine in pilot but fail to scale. Mitigation: design for scale from Day 1 with templated dashboards that clone across machines, plus governance and internal capability.

Performing Better With the Team You Have

The industry has already cut its workforce sharply. AI isn't a lever to cut deeper - it makes the people who remain measurably more effective.

Three workforce realities the AI addresses - turning fewer, less-experienced people into a stronger operating team

Hiring & retention

Process engineers and skilled operators are harder to hire and harder to keep. Embedding the optimal strategy into the system - recomputed continuously - means the mill is no longer one resignation away from a step-change in performance.

Less experience on shift

Newer operators do not yet have the intuition that came from a decade of grade changes and seasonal swings. Prioritised, in-range guidance gives them the equivalent of a veteran sitting beside them on every shift, on every machine, all the time.

Continuous strategy

Computing the optimal setpoint set across hundreds of variables, every few minutes, all day, every day - that is the capability of greatest worth when operating with fewer and less-experienced people. The team is freed to manage exceptions, not chase steady state.

Closing Thought

Operating in a VUCA World - The Power to Recompute

The best strategy today may be maximum throughput; tomorrow it may be minimum energy. The key word is recompute - now in minutes, not weeks.

V

Volatility

Energy prices, raw-material costs and demand swing faster than recipe books can be revised. Setpoints optimal last quarter are wrong this quarter.

U

Uncertainty

Imported wastepaper quality, monsoon humidity, supply disruptions - none can be forecast. The mill must be ready to re-optimise on whatever it actually receives.

C

Complexity

Hundreds of interacting variables, multiple grades, ageing assets, shifting regulations. No human team can hold all the trade-offs in mind simultaneously - the model can.

A

Ambiguity

The same symptom can have several causes; the same lever can have several effects. Probabilistic scoring on real plant data resolves the ambiguity that human intuition cannot.

We Don’t Just Connect Workflows

We Create Measurable Performance Gains

India: Mumbai	India: Nashik	India: Gandhinagar	United Kingdom	U.S.A.
Head Office	Corporate Office	Corporate Office	London	New York
+91 99670 66333	+91 88056 99353	+ 91 83290 68743	+44 20 3239 7604	+1 917 310 1906

info@decimalpointanalytics.com | Privacy Policy | Contact Us