SensoryOps

Sense the trend.
Perfect the batch.

sensoryops.com

PRE-SEED · £750,000
PILOT TARGET · DIAGEO

A physics-informed AI for the alcoholic beverage industry.

COMP0039 · TECH ENTREPRENEURSHIP · 2026 · PAUL REYNOLDS

Where the loop closes.

VISION

A world where beverage manufacturers answer every consumer trend without losing product consistency – where the distance between a social signal and a perfect shelf-ready can is measured in days, not development cycles.

MISSION

We build a physics-informed AI that closes the feedback loop between consumer demand signals and mash-tun process control, so every batch of alcohol hits its predicted flavour profile before the first litre runs.

Alcohol is made at industrial scale – and a meaningful share of every batch never makes it to bottle.

Why? Each batch drifts away from its target formulation while it is being made. Why can't the line stop it? By the time the sensors catch the drift, the chemistry has already gone past the point of recovery. A 200,000-litre mash cannot be rewound. SensoryOps closes that gap.

85%

CPG failure rate within two years of launch

NielsenIQ · multi-year

US$5-10m

Cost per failed beverage launch

PDMA industry estimates

12-18 mo

Typical reformulation cycle

Industry benchmark · CPG Timeline 2025

The category is shifting faster than the production lines can reformulate.

US$69.7bn

UK alcoholic beverages market 2025

Statista Market Insights · 2025

+175%

Non-alc beer volume growth 2019 – 2024

ProSight Financial · 2025

US$830bn

Listed-drinks market cap lost since June 2021

Bloomberg Alcohol Index · 30 October 2025

Not a dying market. A restructuring one.
The winners will be manufacturers who can adapt.

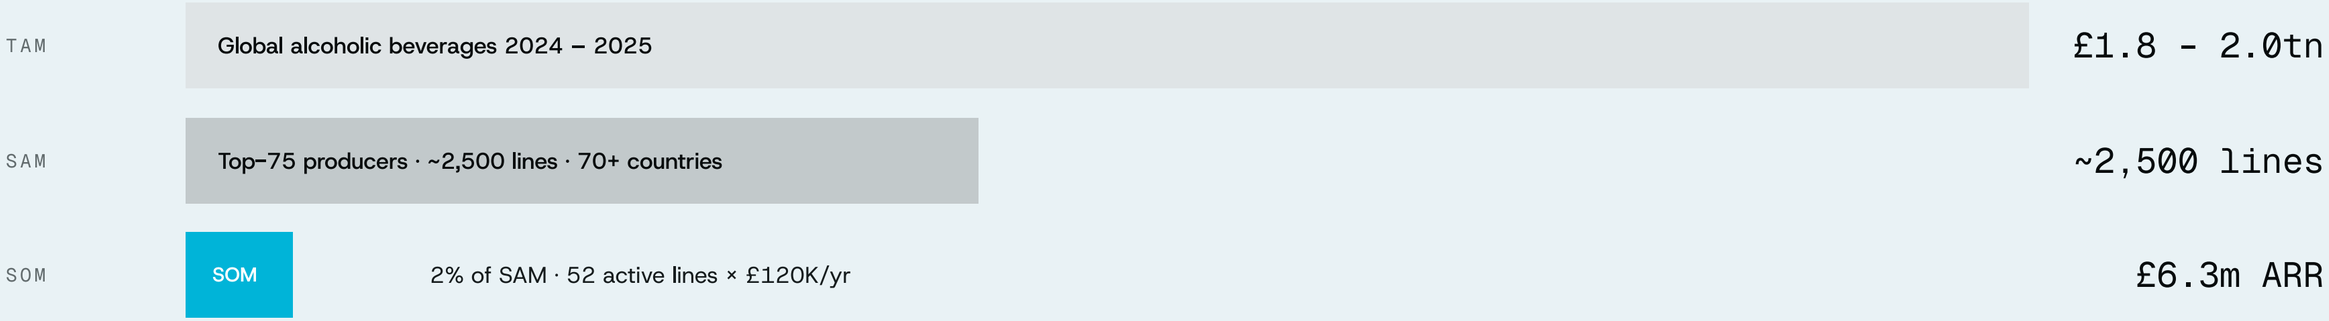

STATISTA / MORDOR INTELLIGENCE 2024 - 2025 · COMPANY FILINGS 2024 - 2025 · INTERNAL MODEL · BARS ARE ILLUSTRATIVE SCALE, NOT EXACT PROPORTION

THE TWO ENDS

Consumer
↔
Factory

Trend prediction and manufacturing consistency are not two problems. They are one feedback loop.

A flavour house can tell you what consumers will want. A factory-automation platform can tell you whether a line is running efficiently. Neither closes the loop. SensoryOps does – by treating the factory and the consumer as two ends of a single physics-informed control system.

01 **A PINN, not generic ML.** Every prediction is constrained by real thermodynamics and enzyme kinetics – physically viable by construction.

02 **A Digital Twin of the mash tun** that learns plant-specific corrections from 100 – 500 examples, not 10,000+. Physics embedded, so less data needed.

03 **A direct actuation layer to the PLC** – temperature, pressure, grain size – not another dashboard to read.

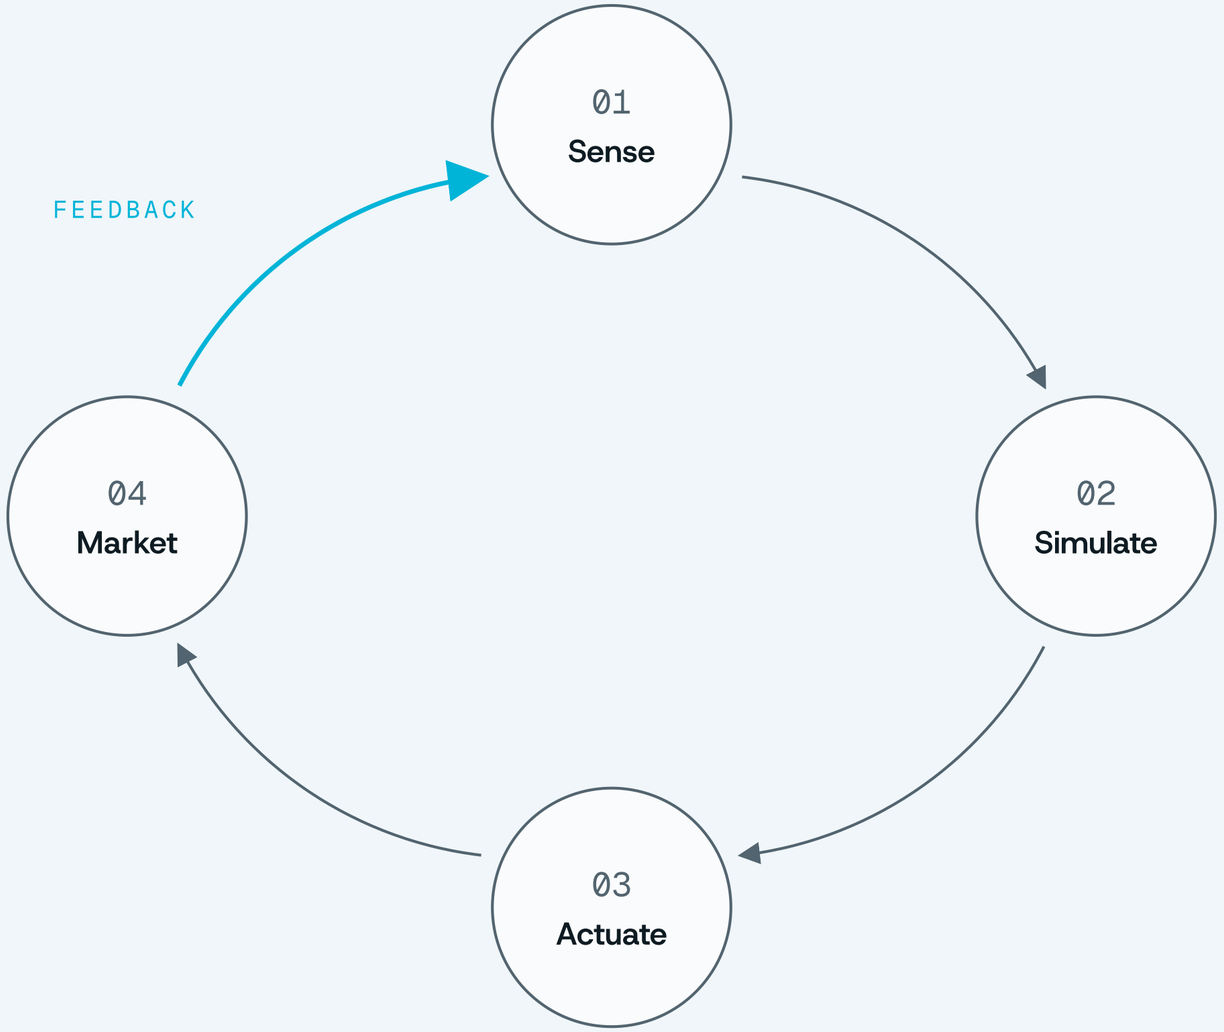

Sense. Simulate. Actuate. Feed back.

STEP #1 — SENSE THE MARKET

NLP reads social, POS and review data; translates preferences into chemical targets (target Brix, pH, ester profile).

STEP #2 — SIMULATE THE PHYSICS

A PINN models the mash with Michaelis–Menten enzyme kinetics and Fourier's Law thermal dynamics; every output is physically viable.

STEP #3 — ACTUATE THE FACTORY

Optimised parameters push to the PLC — temperature, pressure, grain size — so the batch hits the predicted flavour profile.

FEEDBACK

Every shipped batch's market performance feeds back into Step 1; the system learns continuously.

Three modules. One loop.

01 Trend Intelligence Engine

NLP over social, point-of-sale and review data. Output: chemical-target vectors (Brix, pH, ester profile) per SKU.

02 Process Optimiser

PINN + Digital Twin of the mash tun. Couples Michaelis–Menten enzyme kinetics with Fourier heat transport. Output: predicted batch outcome and recommended parameter adjustments before the first litre runs.

03 Actuation Layer

OPC-UA / MQTT bridge to the plant PLC. Four-week wire-in. No PLC replacement.

MVP is the mash-tun twin on a canned-cocktail base. Deep and specific beats broad and shallow.

Nobody else owns the full stack.

Consumer-side AI ends at the lab door. Factory-side AI is generic.
We are the only layer that closes the loop with physics.

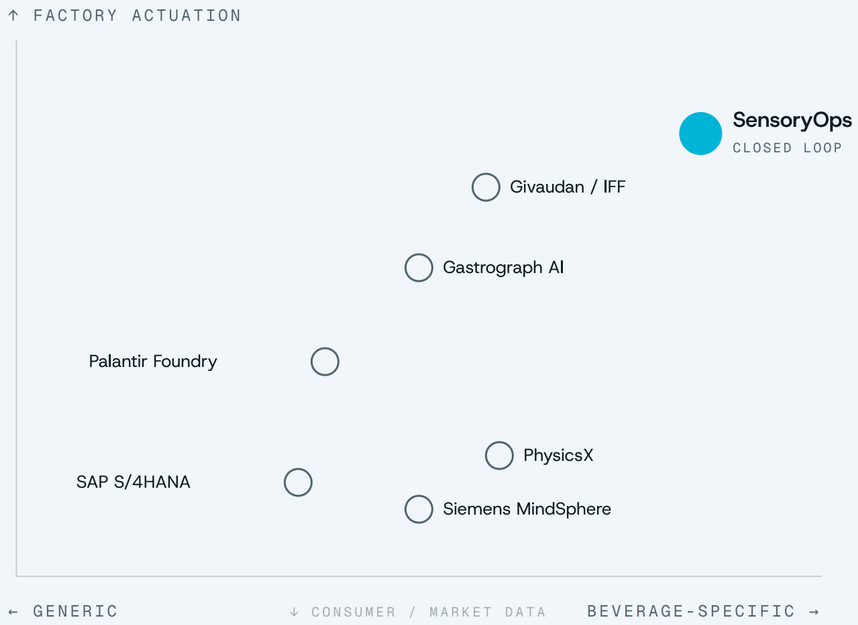

Gastrograph AI

Acquired by NielsenIQ · consumer-side validated

NIQ press release · 7 April 2025

PhysicsX · $135m

Series B led by Atomico · Siemens + NVIDIA · near-unicorn physics-AI

GlobeNewswire · 23 June 2025

For the production lines that already run at industrial scale.

The beachhead is the top-75 global alcohol producers — around 2,500 lines across 70+ countries. A pilot is one line. We land on a pilot and expand into the customer's portfolio from there.

01	**AB InBev**	175+ breweries · 40+ operations sites	US$58.6bn
02	**Heineken**	165+ breweries in 70+ countries	US$33.8bn
03	**Diageo** PILOT TARGET	140+ manufacturing sites in 180 countries	US$32.4bn
04	**Pernod Ricard**	~90 production sites	US$12.9bn
05	**Carlsberg**	~70 breweries	US$9.6bn

COMPANY ANNUAL REPORTS · 2024 - 2025 FILINGS

CEO has already validated the concept with a global alcohol manufacturer.

From PINN validation to Series A.

TRL	2026 Q2 - Q3	2026 Q4 - 2027 Q2	2027 H2 - 2028	2029+
	01 · VALIDATION	**02 · PILOT**	**03 · SCALE**	**04 · EXPAND**
	PINN operational on synthetic data · industry partner LOI.	4–week wire–in at one line · 3–month live validation.	6 customers · 15 active lines.	RTD premium → spirits mainstream.
	TRL 4 · CRL 2 → 3	TRL 6 · CRL 5	TRL 7 - 8 · CRL 6	TRL 8 - 9 · CRL 7

One–time integration	£50,000	per line	4–week plant wire–in · bespoke Digital Twin calibration · OPC–UA / MQTT bridge
Annual SaaS licence	£120,000	/yr per line	continuous trend feed · live predictions · parameter–adjustment recommendations · 24/7 support

COMPARABLE: PALANTIR FOUNDRY LEVEL 0 FROM £250K+/YR (UK G-CLOUD 14, MAY 2024) · SAP S/4HANA YEAR-1 £600K+ (NBS-US, 2025) · SENSORYOPS IS 5× LOWER ON YEAR-1 TCO

Domain access. Deep-tech build. Enterprise go-to-market.

01	CEO	**Tejas Rathod**	Direct industry access. Validated the concept with a global alcohol manufacturer.
02	CTO	**Paul Reynolds**	Maths and ML background. Owns the PINN engine and Digital Twin. Built sensoryops.com and the live dashboard.
03	COO	**Anushka Sutreja**	Operations and pilot delivery.
04	CMO	**Saranya Roy**	Market strategy and enterprise pipeline.

£750,000 pre-seed. 18 months runway. Near-breakeven on one raise.

ARR	£140K	£730K	£1.8m	£3.7m	£6.3m
LINES	Y1 · 1	Y2 · 6	Y3 · 15	Y4 · 31	Y5 · 52

GROSS MARGIN 77% STEADY · EBITDA-POSITIVE YEAR 4 (£498K) · YEAR 5 EBITDA £1.56M

£550K
LTV per line (5-yr, 93% retention)

£30 - 50K
CAC per customer

11 - 18×
LTV : CAC

USE OF FUNDS · £750,000

The alcohol industry is restructuring. US$830 billion has been wiped from listed producers since 2021 – not because people stopped drinking, but because legacy production cannot track how they drink now. SensoryOps closes that gap with physics, not guesswork.

SensoryOps

Sense the trend. Perfect the batch.

A physics-informed AI for the alcoholic beverage industry.

The batch that never makes it to bottle

Every batch of alcohol begins as a chemical target. A specific Brix. A pH window. An aromatic profile a taster has signed off. It ends as the bottle, can, or cask that reaches a customer. Most of the value lost between those two points is lost in the middle, and the plant rarely notices in time.

Traditional line automation holds a schedule that was hand-tuned at commissioning. The in-line sensors can confirm what has already happened, but not what is about to. By the time a process engineer catches a drift in the mash, the chemistry is already past the point of recovery. A 200,000-litre mash cannot be rewound. The batch either finishes off-spec or gets blended down, and the margin goes with it.

85%
of new consumer-packaged-goods launches fail within two years of shelf
NielsenIQ

US$5 – 10m
average cost of a single failed beverage launch
PDMA

US$830bn
listed-drinks market capitalisation erased since 2021
Bloomberg, 30 Oct 2025

Two problems, one feedback loop

Consumers are shifting categories faster than production lines can keep up. Flavour houses can tell a manufacturer what a shelf ought to taste like. Factory-automation platforms can tell a manufacturer whether a line is running efficiently. Neither closes the loop between the two.

That is the gap SensoryOps is built to close. We do not treat trend prediction and manufacturing consistency as two separate problems. They are two ends of the same feedback loop, and the only honest way to close it is with physics rather than guesswork.

How the platform works

The platform runs a three-step control loop.

Sense. An NLP engine reads social signal, point-of-sale and review data, and converts it into the chemical targets a plant already understands: a target Brix, a pH range, an ester profile, an ABV band. One target set per SKU.

Simulate. A Physics-Informed Neural Network models the mash tun under Michaelis-Menten enzyme kinetics and Fourier's Law thermal transport. Every prediction is constrained by the actual laws of chemistry and heat. The system cannot propose a batch that defies its own physics, which is the property that makes the output trustworthy on a plant floor.

Actuate. The optimised parameters push directly to the plant PLC over OPC-UA or MQTT. Temperature, pressure, grain size, set-point schedule. No dashboard to babysit. Every shipped batch's market performance then feeds back into the first step, and the loop improves.

What a pilot looks like

A pilot is a single production line. Four weeks to wire in. No PLC replacement required. SensoryOps provides the physics-informed model, the Digital Twin calibration, and the trend-intelligence feed. The manufacturer provides existing sensor telemetry and one process engineer to support commissioning. We instrument one recipe end-to-end, then let the system expand across the customer's portfolio one line at a time.

Pricing

£50,000	one-off integration, per production line
£120,000 / yr	SaaS licence, per production line

Comparable enterprise platforms start at multiples of this. Palantir Foundry Level 0 from £250,000 per year on UK G-Cloud 14 (May 2024); SAP S/4HANA Year-1 total cost of ownership above £600,000 (NBS-US, 2025). A single off-spec batch in a mid-sized production line pays for a year of SensoryOps.

INTRODUCING SENSORYOPS

Sense the *trend*. Perfect the *batch*.

A physics-informed AI for the alcoholic beverage industry.

THE THESIS

SensoryOps predicts every batch of alcohol in real-time. The model is physics-informed, so it generalises past any point the sensors have seen. Correct drift before it becomes waste.

SENSORYOPS V1.0 · COUPLED MICHAELIS–MENTEN + FOURIER · 200 COLLOCATION POINTS · MASH TUN M-014

A physics-informed neural network learns your mash tun's thermodynamics and enzyme kinetics – so it generalises past any point your sensors have seen. In validation: 1.27% L^2 error across a full 60-minute cycle.

Open the live dashboard →

See how the loop closes →

SCROLL FOR THE PROOF

↓

FIG 01 · MASH-TUN M-014 · 62°C ISOTHERMAL REST

PINN prediction versus ground truth

200 collocation points · P-sugar relative L^2 error 1.27%

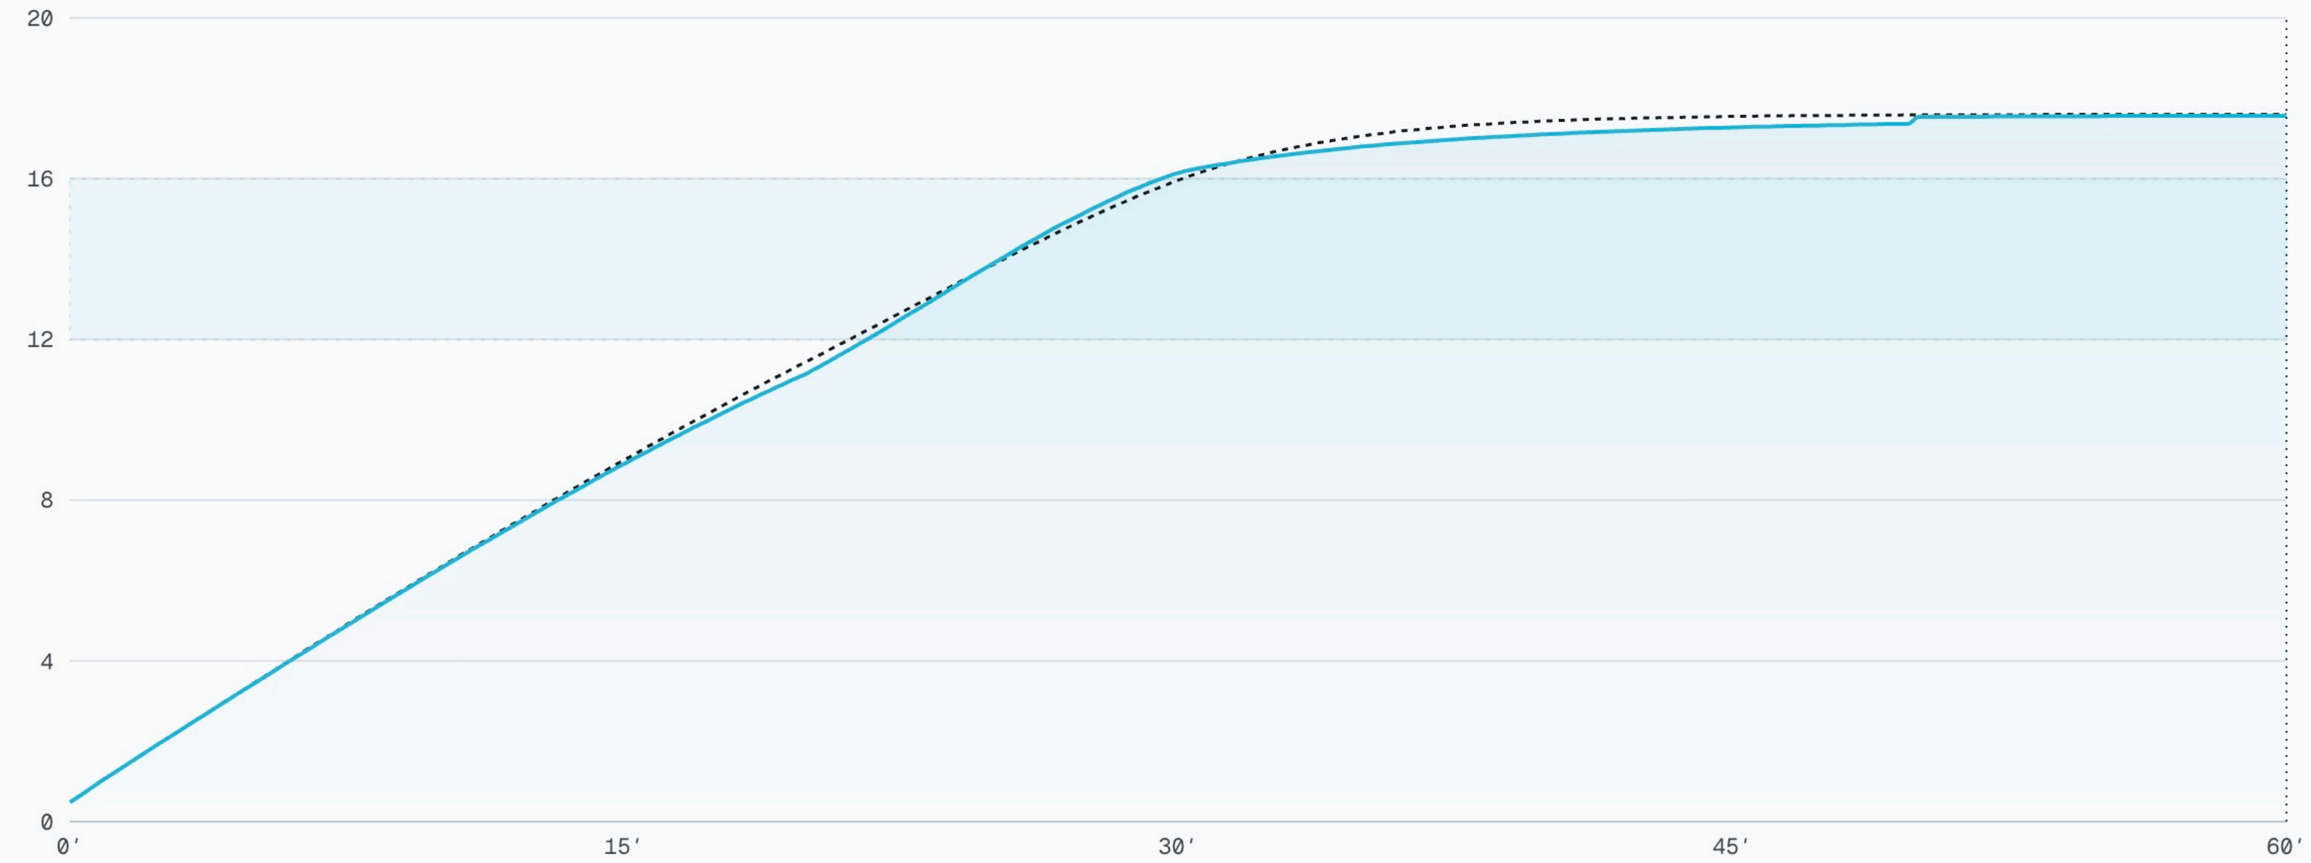

PINN prediction Ground truth Target window · 12–16 °Bx

T=60′ FINAL 17.56 °BX +1.56 ABOVE TARGET

01 TARGET VESSEL

Industrial mash tun

10,000 L · continuous production line

02 MODEL

2-head PINN

18 enzyme params · coupled ODE + 1-D heat

03 INTEGRATION

4-week wire-in

OPC-UA + MQTT · no PLC replacement

WHERE THE PHYSICS PLUGS IN

Four moments in every batch.

SensoryOps runs a physics-informed digital twin of each vessel on your line and surfaces a prediction at the four moments where flavour is set and drift is cheapest to correct.

01 MICHAELIS–MENTEN

Mash

Enzyme kinetics set 80 per cent of the finished flavour. The PINN forecasts Brix, pH, and α/β-amylase conversion for every minute of the rest curve.

02 DARCY FLOW

Lauter

Wort clarity and extract efficiency depend on grain-bed permeability. Off-spec runs are flagged before the first litre leaves the vessel.

03 ARRHENIUS

Boil

Hop isomerisation and DMS evolution. The model forecasts IBU and volatile load from the boil profile so adjustments land before a cool-down is wasted.

04 MONOD KINETICS

Fermentation

Yeast kinetics under varying pitch rate and temperature. The digital twin predicts attenuation curves and ester load for the cellar team.

01 · THE PROBLEM

01

Alcohol is made at industrial scale, and a meaningful share of every batch *never* makes it to bottle.

Why? Each batch drifts away from its target formulation while it's being made.

Why can't the production line stop it? By the time the on-line sensors catch the drift, the chemistry has already gone past the point of recovery – and a 200,000-litre mash cannot be rewound.

SensoryOps closes that gap.

01

85%

of new consumer-packaged goods fail within two years of launch.

NIELSENIQ

02

$5M

average cost of a failed product launch in R&D, marketing, and inventory.

PDMA, 2024

03

12–18 mo

the R&D cycle before a production line is locked in and committed.

CPG TIMELINE GUIDE, 2025

PHYSICS OF DRIFT

Most products don't fail because the concept was wrong – there genuinely is demand for a mango hard seltzer or a low-calorie RTD gin cocktail. They fail because the formulation that tastes right in a 5-litre lab batch doesn't taste right in a 10,000-litre production run. The flavour drifts. The mouthfeel changes. By the time the manufacturer notices, the R&D cycle has already committed them to production, marketing, and distribution.

The lab-to-line gap is a physics problem dressed as a process problem. At 5 litres, a master distiller can hold a grain bed within ±0.2°C for every minute of the mash rest. At 10,000 litres, the grain bed is a 45-centimetre-deep thermal insulator – a temperature change applied at the recirculation loop takes 10 to 15 minutes to reach the bottom.[1] The enzymes do not wait: α-amylase inactivates on a sharp Arrhenius curve above 72°C, β-amylase above 65°C.[2] Those ten minutes of thermal drift are enough to shift the fermentability of the wort by several percentage points – which shows up in the finished batch as body, mouthfeel, and residual sweetness that the lab recipe never predicted.

Today's counter-measure is a tasting panel five hours after the mash is already done. That is too late to adjust this batch, and too slow to close the loop on the next one. What the industry needs is a way to predict the final flavour-critical numbers – Brix, pH, ester load – before the first litre runs, using the physics of this specific vessel, with this specific grist, at this specific ambient. That is the SensoryOps thesis.

▸ SOURCES

02 · WHY NOW

The category is shifting faster than the production lines can reformulate.

02

01

US$69.7bn

UK alcoholic beverages market – SensoryOps's beach-head.

STATISTA, 2025

02

2×

forecast global RTD growth 2019–2029 – premium value outpacing the rest of beverage alcohol.

IWSR RTDS STRATEGIC STUDY, 2025

03

US$830bn

in shareholder value destroyed since 2021 – reclaimable by producers who can reformulate batch-by-batch.

BLOOMBERG, 2025

WHERE THE VALUE GOES

Ready-to-drink cocktails are the fastest-growing category in alcohol, with independent market-research estimates putting the compound annual growth rate somewhere in the 11.8 to 15.7 per cent range through the end of the decade (Straits Research, 2025; Polaris Market Research, 2025). IWSR, in its own RTDs Strategic Study, forecasts the category doubling globally between 2019 and 2029 with premium value outpacing volume. Gen Z drinkers aged 21 to 27 now match the general drinking population at 74 per cent – they drink differently, not less (IWSR, 2025). Investors are already pricing in the manufacturers' inability to adapt: the companies that can reformulate line-by-line, batch-by-batch, will win that value back.

The bottleneck is not creativity – RTD R&D teams have never been more fluent in what Gen Z wants. The bottleneck is the 12-to-18-month journey from a prototype that tastes right in a 5-litre vessel to a production line that reliably makes the same thing at 10,000 litres, batch after batch. SensoryOps collapses that journey by closing the loop between the flavour target a manufacturer needs and the physics of the vessel it has to run in.

03 · THE CLOSED LOOP

03

The closed loop.

Three steps. One continuous loop. Each batch that ships refines the next batch's prediction, so the digital twin gets more accurate the more you brew.

STEP 01

Sense

SensoryOps ingests category-level taste signals – RTD shelf data, social listening, Gen Z sensory trials – and compiles them into chemical targets a production line can actually run against. Target Brix, pH, bitterness, ester load.

STEP 02

Simulate

A physics-informed neural network, trained on this vessel's geometry and this grain bill's enzyme profile, predicts what will come out the bottom valve – before the mash starts. Every forecast is constrained by Fourier's Law and Michaelis–Menten kinetics, so the prediction is physically consistent, not just statistically plausible.

STEP 03

Actuate

Optimised setpoints push straight to the PLC: mash rest temperatures, hold times, recirculation rates. The batch hits the target the first time. Shelf performance feeds back into step one, and the loop tightens.

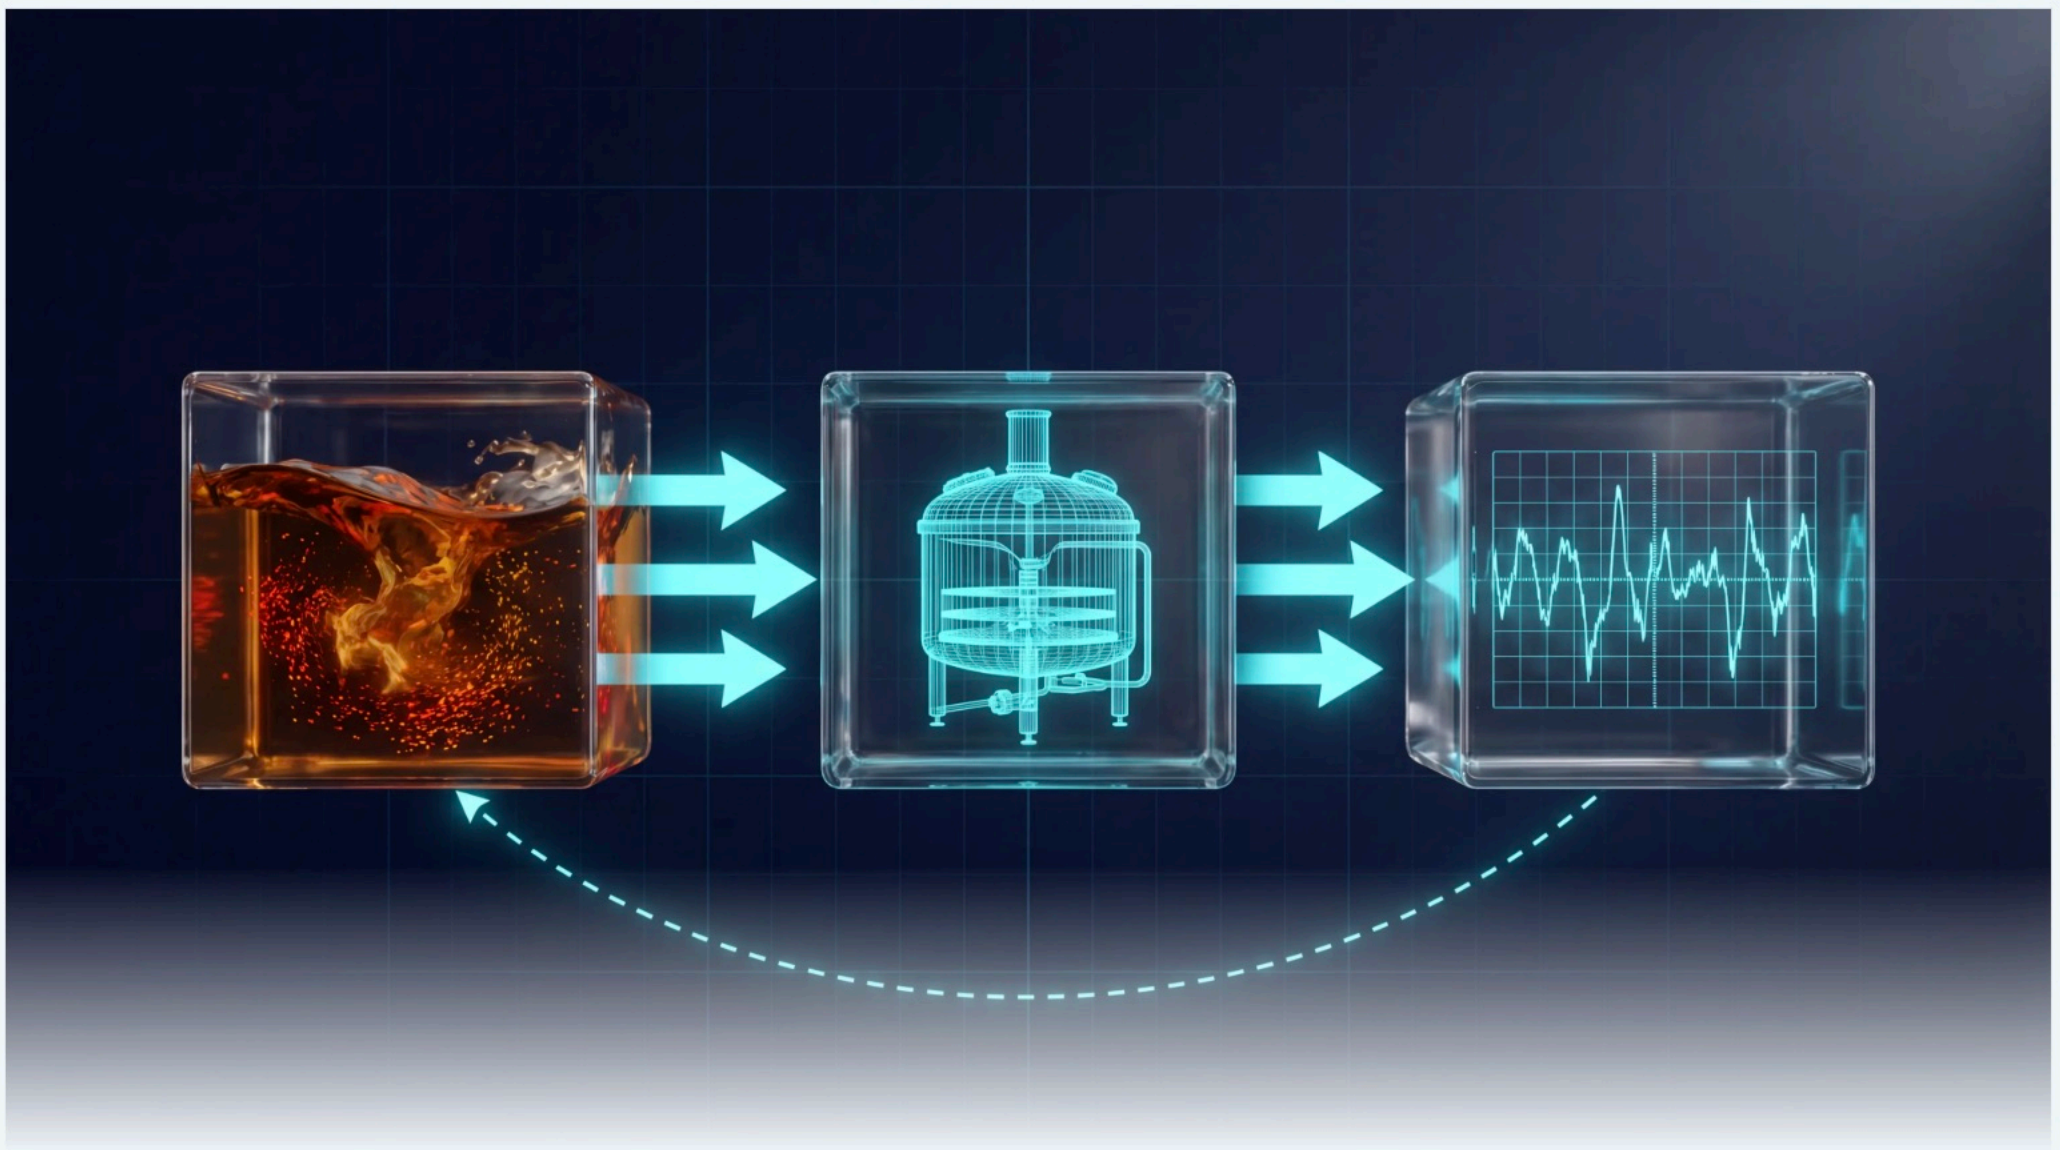

FIG 02 · THE LOOP, MADE VISIBLE. SENSE FEEDS SIMULATE FEEDS ACTUATE; THE NEXT BATCH'S SHELF DATA FEEDS BACK INTO SENSE AND THE PREDICTION SHARPENS.

04 · THE CAPABILITIES

04

Three things the PINN actually does.

Each pillar carries a proof. A sourced number from the current v1.0 run or the integration spec, not a hit rate invented for the pitch.

01 PROOF

Predict

Forecasts every batch's final Brix, pH, conversion, and thermal gradient before the first litre runs. One inference per scenario, sub-second on commodity hardware.

1.27%

L² ERROR · SUGAR TERM, 60-MINUTE REST · PREDICTIONS_VS_GROUND_TRUTH.JSON

02 PROOF

Explain

Every prediction ties back to a physics equation. The operator sees the Fourier residual, the Michaelis–Menten constants, and the L² error before accepting a setpoint change. Audit-ready on day one.

6 terms

PHYSICS-CONSTRAINED LOSS DECOMPOSITION · TRAINING_LOSS.JSON

03 PROOF

Deploy

Four-week wire-in using the standards the plant already speaks. OPC-UA + MQTT, no PLC replacement, no rip-and-replace. The digital twin runs alongside the line and hands setpoints forward, not backwards.

£50k + £120k/yr

INTEGRATION + PER-LINE SAAS · SEE /PILOT

05 · PILOT

05

For the production lines that already run at industrial scale.

SensoryOps is designed for the global brewers – Diageo, AB InBev, Pernod Ricard – whose production lines run continuously and whose R&D teams are under pressure to reformulate line-by-line. We are not building a tasting app or a shelf-trend report. We are building the physics layer that turns those reports into batches that hit the target the first time.

FIG 03 · THE KIND OF PLANT SENSORYOPS WIRES INTO – CONTINUOUS PRODUCTION, STAINLESS-STEEL VESSELS, OPC-UA CONTROL.

CATEGORY VALIDATION

Gastrograph AI, the closest 'AI + taste' comparable, was acquired by NielsenIQ in April 2025 (NIQ press release, 7 April 2025). PhysicsX, the closest 'physics-informed AI' comparable, raised $135M Series B in June 2025 at a near-unicorn valuation (GlobeNewswire, 23 June 2025). SensoryOps sits at the intersection – consumer-insight AI and physics-informed AI, applied to the one vertical where both matter at once.

06 · PRICING

06

Pays for itself against one off-spec batch.

PDMA estimates the average cost of a failed product launch at approximately $5M (PDMA, 2024).

INTEGRATION (ONE-OFF)

£50,000

- Vessel geometry capture and PLC integration survey
- PINN training on historical batch logs from the pilot line
- Onboarding: 4 sessions with the R&D team

SAAS, PER PRODUCTION LINE

£120,000 / yr

- Unlimited batch forecasts and scenario simulations
- Real-time digital-twin dashboard (live demo: /dashboard)
- Quarterly model refreshes as the line's data grows

07 · SEE THE PINN RUN

Open the live *dashboard*.

No sign-up, no login. The full proof-of-concept runs on pre-computed PINN output: thermal field, enzyme kinetics, scenario sweeps, 3D mash tun render.

Open the dashboard →

PINN V1.0 · LIVE
200 COLLOCATION POINTS
1.27% L^2 RELATIVE ERROR

01 · SITE MAP

01

Every page and section on sensoryops.com.

Flat top-nav architecture. The persistent bar at the top of every page routes between six sibling surfaces; the `/dashboard` branch adds five sub-views of the live PINN via its own tab strip. The wordmark returns to `/`.

01 Open ↗

/

Home · the landing page for commercial visitors

- `#process-map` Four moments in every batch · taxonomy of where the physics plugs in
- `#problem` 01 · The problem · Socratic framing + sourced statistics
- `#opportunity` 02 · Why now · market figures and bottleneck argument
- `#how-it-works` 03 · The closed loop · three-step process ledger
- `#capabilities` 04 · The capabilities · three pillars with sourced proof numbers
- `#pilot` 05 · Pilot · target customers + category validation
- `#pricing` 06 · Pricing · two-tier financial spec sheet
- `#cta` 07 · See it run · final dashboard CTA

02 Open ↗

/dashboard

Live PINN-driven proof of concept

- `(overview)` KPI grid against batch targets
- `/thermal-profile` 1-D heat-transfer field across mash bed depth
- `/enzyme-kinetics` Coupled Michaelis-Menten + Arrhenius curves
- `/scenario-analysis` What-if controls · live re-prediction
- `/3d-field` Volumetric thermal render (React-Three-Fiber)

03 Open ↗

/courseworks

COMP0039 submission index · CW1 / CW2 / CW3 / CW4

04 Open ↗

/about

Thesis, team, and where the company is right now

05 Open ↗

/pilot

Pilot enquiry form · routes to pilot@sensoryops.com

06 Open ↗

/sitemap

This page