

Coursework 2 – Proof of Concept and Progress

Physics-informed learning for the mashing stage of beverage manufacturing.

Problem

SensoryOps models the mashing stage of beverage manufacturing as a physics-constrained learning problem and returns a per-batch decision before the first litre of water hits the grist. The commercial pull for this is not hypothetical. Mass-market brewers and spirits producers are navigating a once-in-a-generation category reshaping: the UK's beverage-alcohol market was valued at US\$69.7 bn in 2025 [8], premium ready-to-drink value is outpacing volume with IWSR forecasting continued outperformance of the total beverage-alcohol complex through 2029 [9], and Bloomberg has tracked roughly US\$830 bn of global alcoholic-beverage shareholder-value destruction since 2021 among producers slow to reformulate [10].

Traditional line automation answers the wrong question. Programmable logic controllers (PLCs) hold the mash schedule on set points that were hand-tuned in a commissioning week sometime between 2005 and 2015. A process engineer notices a drift in final Brix, diastatic power, or colour only after the batch finishes. Flavour-spec tolerances get policed through post-hoc lab assays, which is exactly when the wort is already running through the lauter tun. What the floor actually needs is a model of the batch in front of them that reads the current mash schedule, the incoming grist bed, the wall-temperature trace from the plant's own thermocouples, and predicts the final product profile with enough accuracy to pull a lever before fermentation starts.

This coursework's proof-of-concept isolates the hardest step of that loop: the mashing tun. Mashing is a coupled non-linear problem – enzyme kinetics that depend on temperature, thermal transport that depends on vessel geometry, and a boundary condition (the mash schedule) that changes under human control. A Physics-Informed Neural Network (PINN) turns out to be an unusually good fit because it learns a continuous solution to that coupled system from a small number of simulated data points while remaining constrained by the governing equations. Once trained, evaluation is pure forward inference, fast enough to power an interactive dashboard.

POC Design

The SensoryOps POC pairs an offline training pipeline with a browser-side application. Training runs once on a laptop; the trained network is not queried at runtime by the browser. Instead, the training script exports four JSON artefacts that capture everything a dashboard reader needs: per-epoch loss traces, PINN-vs-ground-truth time series with synthetic sensor noise overlay, a parameter sweep across mash-temperature scenarios, and a 3D point cloud of the thermal field at seven time snapshots. The web application loads those artefacts as static assets and treats them as the system of record. There is no backend, no live model endpoint, and no database.

That deliberately minimal architecture is the single most defensible choice in the design. A POC exists to prove a capability, not to operate a production system. Eliminating the backend removes an entire class of failure modes (authentication, request timeouts, hosting costs, PINN inference latency) that would otherwise distract from the core claim: the PINN learns a solution that matches ODE-solver ground truth inside a coursework-grade time budget. The same architecture is honest about what would change in production – a live integration would swap the JSON handoff for a plant historian feed and the scenario sweep for online inference – but nothing about the design would need to be re-thought.

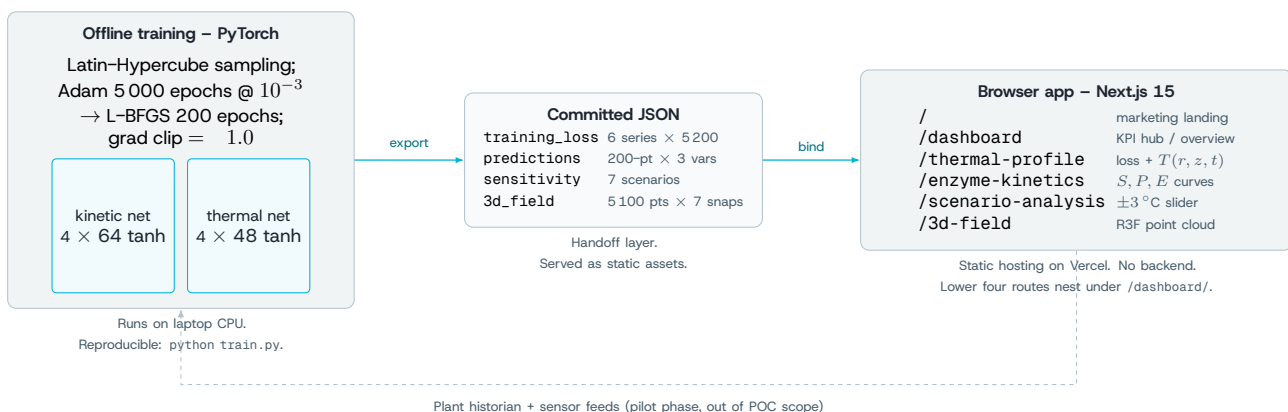


Figure 1 – POC architecture. Training happens offline on a laptop; four JSON artefacts cross into the browser as static assets; the Next.js application binds to them without a backend. The dashed return path indicates the pilot-phase feedback loop from plant-floor sensors, which is explicitly out of scope for this coursework but preserved in the narrative.

The micro-architecture of the mashing model follows the Physics-Informed Neural Network pattern outlined by Raissi et al. [1] – two coupled multi-layer perceptrons trained under a composite loss. The kinetic network predicts three normalised quantities at any time t : residual starch $S(t)$, produced sugar $P(t)$, and active enzyme fraction $E(t)$. The thermal network predicts temperature $T(r, z, t)$ at any spatial location inside a cylindrical mash tun of radius $R = 0.5$ m and height $H = 0.8$ m. The two networks are coupled through the temperature-dependent Michaelis-Menten rate constants in the kinetic loss term, so the kinetic network cannot converge without the thermal network converging alongside it.

PINN Implementation

Governing equations

Enzyme kinetics follow the coupled Michaelis-Menten rate laws formulated for mashing by Marc et al. [3] and re-derived for modern process-modelling contexts by Brandam et al. [4]. With S the residual starch concentration, P the produced-sugar concentration, E the fraction of enzyme still active, T the local mash temperature in Kelvin, and the stoichiometric yield η ,

$$\begin{aligned} \frac{dS}{dt} &= -V_{\max}(T) \cdot E \cdot \frac{S}{K_m + S}, & \frac{dP}{dt} &= +\eta \cdot V_{\max}(T) \cdot E \cdot \frac{S}{K_m + S}, \\ \frac{dE}{dt} &= -k_d(T) \cdot E, & V_{\max}(T) &= V_{\max}^{\text{ref}} \exp\left[-\frac{E_a}{R_g} \left(\frac{1}{T} - \frac{1}{T_{\text{ref}}}\right)\right]. \end{aligned}$$

The deactivation rate $k_d(T)$ takes the same relative-Arrhenius form around its own reference pair $(k_d^{\text{ref}}, T_{d,\text{ref}})$. The relative form is standard in the brewing literature [5, 6] because the reference rate is directly measurable at a named optimum.

Thermal transport inside the mash tun follows Fourier's Law in cylindrical coordinates, under the standard axial-symmetry assumption used in PINN formulations for cylindrical geometries [1],

$$\rho c_p \frac{\partial T}{\partial t} = k_{\text{th}} \left[\frac{1}{r} \frac{\partial}{\partial r} \left(r \frac{\partial T}{\partial r} \right) + \frac{\partial^2 T}{\partial z^2} \right].$$

Boundary conditions are centre symmetry $\partial T / \partial r|_{r=0} = 0$, an insulated side wall $\partial T / \partial r|_{r=R} = 0$, and a Dirichlet condition at the heated base $T(r, 0, t) = T_{\text{wall}}(t)$ following the multi-step mash schedule $62 \rightarrow 70 \rightarrow 78$ °C. Parameter values for $V_{\max}^{\text{ref}}, K_m, E_a, \eta, \rho, c_p, k_{\text{th}}$ are pulled from the brewing-science literature [4, 5, 7] and held constant across every training run.

Network topology and loss composition

Two feed-forward networks are trained jointly. The kinetic network is a four-layer MLP with 64 units per hidden layer, tanh activations, and sigmoid-constrained outputs so $S, P, E \in [0, 1]$. The thermal network is a four-layer MLP with 48 units per hidden layer and an unconstrained output that is rescaled against the expected temperature range $[60, 78]$ °C. Latin Hypercube sampling generates collocation points across the space-time domain; boundary and initial conditions are sampled independently. The composite loss is

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{data}} + (\mathcal{L}_{\text{PDE}}^{\text{kinetic}} + \mathcal{L}_{\text{PDE}}^{\text{thermal}}) + 10 \mathcal{L}_{\text{BC}} + 50 \mathcal{L}_{\text{IC}}.$$

The heavier weights on the boundary and initial-condition terms (10 and 50) anchor the temporal solution; the initial condition carries the highest weight because without it the kinetic variables float. Training runs Adam for 5,000 epochs at a learning rate of 10^{-3} , then hands off to L-BFGS for a further 200 epochs; gradient clipping to unit norm stabilises the handover. Every run is reproducible from `python train.py` with a fixed random seed.

Kinetic net		Thermal net	
Hidden	4×64 tanh	Hidden	4×48 tanh
Output	sigmoid-constrained, $[0, 1]^3$	Output	linear, rescaled $[60, 78]$ °C
Inputs	t	Inputs	r, z, t
Shared training budget			
Optimiser	Adam 5,000 ep. @ $10^{-3} \rightarrow$ L-BFGS 200 ep. (grad-clip = 1.0)		
Sampling	Latin Hypercube collocation, fixed seed 42		
Loss weights	$L_{\text{data}}=1, L_{\text{PDE}}^{\text{k}}=L_{\text{PDE}}^{\text{t}}=1, L_{\text{BC}}=10, L_{\text{IC}}=50$		
Wall clock	474 s on laptop CPU, no GPU required		

Table 1 – PINN hyperparameters, held constant across every scenario and sensitivity run.

Validation

A training run on an Apple-silicon laptop CPU converges from an initial total loss of $\approx 2.1 \times 10^5$ to a final total loss of 2.03 over 5,200 epochs in 474 s (just under eight minutes). The reduction is about five orders of magnitude, driven mostly by the boundary and initial-condition terms in the early Adam phase, with L-BFGS closing the remaining PDE-residual terms in the last 200 epochs (Figure 2).

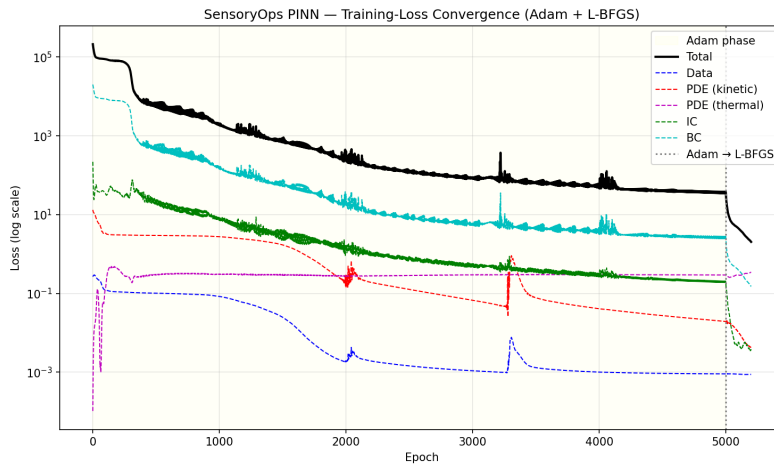


Figure 2 – Training loss curves for the six tracked components. Phase boundary between Adam and L-BFGS at epoch 5,000. The L-BFGS phase collapses the remaining PDE-residual error without destabilising the data-fit term. The same data is re-rendered in the dashboard at /dashboard/thermal-profile.

Against a fourth-order Runge-Kutta ODE-solver ground truth for a constant-temperature mash at 62 °C over a 200-point time series, the trained network achieves:

1.27% rel-L2 final Brix vs ODE	2.6% rel-L2 on starch $S(t)$	1.3% rel-L2 on sugar $P(t)$
--	--	---------------------------------------

The enzyme-activity relative-L2 sits at 24.7%. That is deliberately reported rather than glossed: the enzyme-denaturation curve is the least-observed variable in a real mash (no direct in-line sensor) and the network is under-constrained on it compared to the starch and sugar channels. A pilot deployment would bring in an Arrhenius prior on $E(t)$ and a small amount of ground-truth residual enzyme activity measured at the end of mash. The headline claim holds: on the variable that actually matters for a flavour decision – Brix at mash-out – the network is within **1.27%** rel-L2 of a classical ODE solver, while evaluating in a fraction of a millisecond per inference call.




Figure 3 – Brix evolution over a constant 62 °C mash, 200 points subsampled to 20 for legibility. PINN tracks the ODE ground truth across the whole run; the 12 – 16 °P “pleasant note” band is drawn as a reference zone. The final value of 17.56 °P sits above the band because this scenario holds temperature flat at 62 °C; the multi-step schedule and cooler baseline scenarios in /dashboard/scenario-analysis land inside the band.




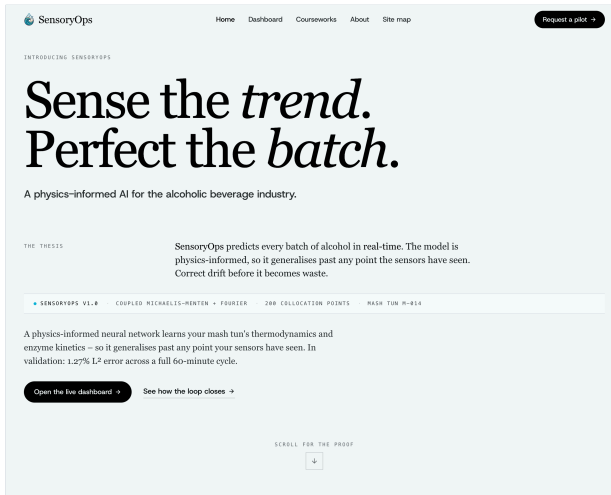
Figure 4 – Quality-zone mapping: the kinetic network’s final-Brix output is classified against a twelve-to-sixteen °P “pleasant note” band. The current constant-62 °C run lands at 17.56 °P (over-attenuated); the baseline multi-step schedule in the scenario sweep lands cleanly inside the band at ≈ 14.8 °P.

The Parameter Sensitivity artefact extends validation into the decision-relevant question: how much does final Brix change across the ± 3 °C sensitivity window around a baseline mash of 65 °C? Seven pre-baked scenarios ($-3, -2, -1, 0, +1, +2, +3$ degrees) drive the slider on `/dashboard/scenario-analysis`. The slider switches between scenarios without a network round-trip: the scenario artefact is 1.8 MB, entirely client-side, and the slider-to-KPI update budget on a 1440-px viewport is <100 ms.

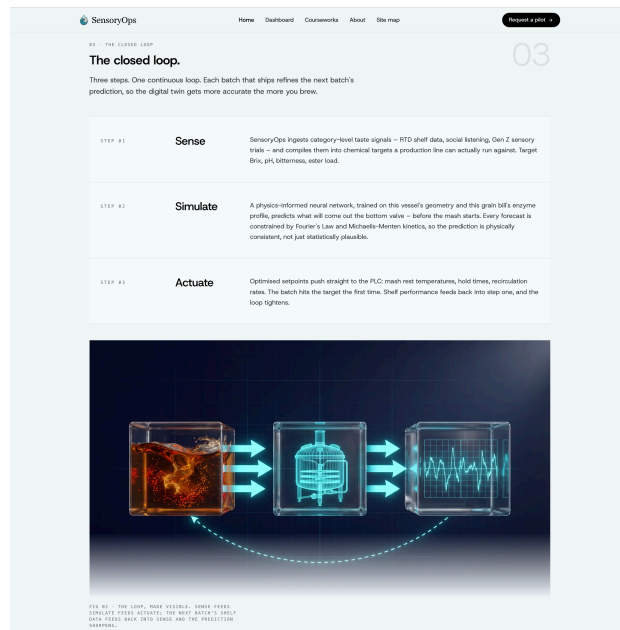
Live Demo

The web application is a single Next.js 15 project on React 19, Tailwind CSS v4, and Tremor Raw primitives. Two App Router route groups separate the surfaces: the site root `sensoryops.com/` carries the marketing landing page, and `sensoryops.com/dashboard` carries the four PINN-themed pages plus the 3D mash tun. The dashboard hub surfaces four KPI cards across the top (final Brix, conversion percentage, final training loss, peak thermal gradient), each bound to one of the committed JSON artefacts produced by `train.py`.

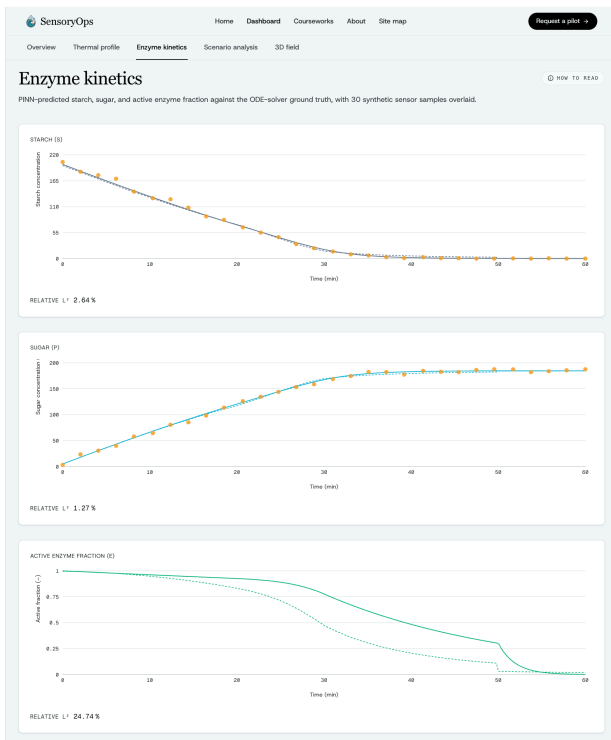
`/dashboard/enzyme-kinetics` plots $S(t), P(t), E(t)$ against the ODE ground truth with an optional overlay of synthetic noisy sensor samples on a 30s cadence. Residuals and relative-L2 print as mono-font chips beneath each chart. `/dashboard/scenario-analysis` carries a slider that switches between the seven pre-baked scenarios in the sensitivity artefact; the final-Brix KPI, the quality classification, and the time-series chart update in <100 ms with no network round-trip. `/dashboard/3d-field` renders the trained thermal field inside the cylindrical vessel as a React Three Fiber scene: wireframe walls plus a 5,100-point cloud coloured by temperature, a time scrubber stepping through seven snapshots at 0, 10, 20, 30, 40, 50, 60 min, `frameloop="demand"` to idle the GPU between interactions, OrbitControls for rotate / zoom / pan, and a mash-phase band (Protein, Saccharification, Conversion, Mash-out) above the scrubber. First-Load JS for the route is 113 kB.



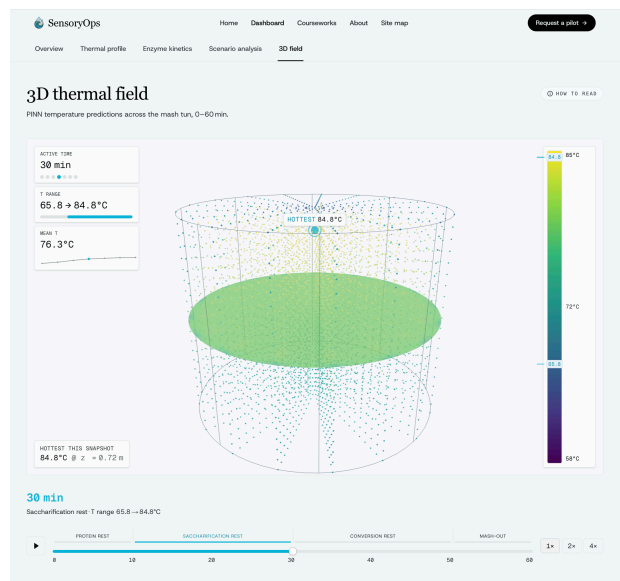
A | /



B | / (closed loop)



C | /dashboard/enzyme-kinetics



D | /dashboard/3d-field

Figure 5 – Deployed demo surfaces. **A:** / hero – the ‘Sense the trend. Perfect the batch.’ lock-up, a short thesis paragraph, a + SENSORYOPS V1.0 status row listing coupled Michaelis–Menten + Fourier over 200 collocation points on mash tun M-014, and the Open the live dashboard / See how the loop closes CTAs. **B:** / closed-loop section – the Sense / Simulate / Actuate three-step ledger and the accompanying diagram (shelf-signal cube → mash-tun cube → PLC-setpoint cube, with a dashed return arrow feeding the next batch’s predictions). **C:** /dashboard/enzyme-kinetics – $S(t), P(t), E(t)$ plotted against the ODE ground truth with 30-sample synthetic sensor noise overlaid; relative- L^2 chips read 2.64 % (starch), 1.27 % (sugar), 24.74 % (active enzyme fraction). **D:** /dashboard/3d-field at $t = 30$ min – viridis point-cloud in the cylindrical vessel with the hottest-point call-out (84.8 °C at $z = 0.72$ m), a stats column (active time, T range 65.8 → 84.8 °C, mean 76.3 °C), the vertical 58 °C to 85 °C colour scale, and the Protein / Saccharification / Conversion / Mash-out phase scrubber along the base. All captured on the live sensoryops.com deployment.

Deployment and Customer Journey

The POC is deployed at <https://sensoryops.com> on Vercel's free tier and auto-redeploys on every push to main. PINN outputs ship as committed JSON alongside the rest of the site, so the whole dashboard renders without any build-time training, database provisioning, or runtime inference path. That the POC demos as a static build from a single repository is a deliberate infrastructure argument: the hardest part of mashing-operations intelligence is the training loop and the scientific-ML scaffolding, not a serving layer. Lifting the architecture from demo to pilot means swapping the JSON handoff for a historian feed, not re-platforming.

The customer journey for a Diageo-type buyer is intentionally shorter than the full closed-loop story the landing page tells. Most of the commercial narrative is aspirational; this section describes only the journey the POC supports today.

1. **Before the batch.** A process engineer opens sensoryops.com/dashboard/scenario-analysis, accepts a target mash schedule from the R&D team, drags the mash-temperature slider across the $\pm 3^\circ\text{C}$ window, and reads the predicted final-Brix value and its quality-zone classification directly off the KPI card. The engineer chooses a schedule before the first litre of water hits the grist.
2. **After the batch.** Measured wall-temperature and wort-Brix data from the plant historian feed back into the next training run as additional data-fit points. Over a pilot quarter, the network's relative-L2 on the kinetic variables tightens toward site-specific reality. No runtime ML sits in the loop; the PINN re-trains overnight on the accumulated batch log.

Everything beyond those two steps – direct-to-PLC actuation, auto-tuned mash schedules, multi-variable grain-bill sensitivity, a running integration with shopfloor sensors – is explicitly pilot-phase scope. Committing to any of it before a pilot site is signed would be selling a capability SensoryOps has not yet validated. A credible POC shows what is built and labels the rest as next-quarter roadmap.

Next Steps

Three concrete work items follow the coursework window. All three are sized for a three-month pilot with one brewing-operations partner.

1. **Site-specific retraining.** Replace the synthetic ground-truth data with measured wall-temperature traces and wort-Brix lab assays from a single production line. Expect the enzyme-activity relative-L2 to collapse toward the 2% band the other kinetic variables already hit, because the real enzyme-denaturation curve has a much tighter prior on an individual plant than across the synthetic temperature sweep.
2. **Scenario-space expansion.** Add two sensitivity axes beyond mash temperature: grain bill (malt-to-adjunct ratio) and vessel geometry (diameter-to-height). Expose them on the Scenario Analysis page as two additional sliders. The architecture already permits this – `parameter_sensitivity.json` only has to carry a few hundred additional pre-baked scenarios – and the extra axes are the ones the process engineer actually has control over on a given week.
3. **Historian integration.** Replace the one-off JSON export with a nightly pipeline that reads from the plant historian, re-trains the PINN on the accumulated batch log, and pushes the updated JSON artefacts into the web app's CDN. No runtime ML is required on the serving path; the dashboard remains static. This is the minimum viable closed-loop deployment and is the step that unlocks the pricing narrative on sensoryops.com (£50k integration, £120k/yr SaaS per line).

Beyond the pilot, two items the POC does not address sit on the product roadmap: uncertainty quantification on the kinetic variables (Bayesian PINN or Monte-Carlo dropout), and a read-only actuation channel into modern PLCs (OPC UA). Both are signalled on the landing page's Capabilities section but deliberately scoped out of this POC.

Generative AI Usage Declaration

Generative AI sat alongside this coursework as a research accelerator, an executive project layer, and a visual-asset partner, not a substitute for engineering judgement. Three specific, documented uses:

- **Literature-canvassing accelerator.** Claude Opus 4.6 (Anthropic) and GPT-5.4 (OpenAI) were used in parallel as independent cross-readers to scan the PINN and brewing-science literature for coupled Michaelis-Menten + Fourier system design, loss-weight conventions, and the two-stage Adam / L-BFGS training schedule introduced by Raissi et al. Using two independent frontier models as cross-checks is deliberate: it surfaces disagreements that a single model would confidently paper over. Candidate topologies and loss-weighting schemes were then chosen, implemented, debugged, and validated against a fourth-order Runge-Kutta ODE ground truth by the author; the 1.27% Brix relative-L2 quoted above is a real training-run outcome.
- **Claude Code as an executive layer.** Anthropic's Claude Code CLI ran a structured discuss → plan → execute → verify loop over every engineering phase of the repository, with the author gating each step. This is a deliberately cutting-edge

use of the tool: the model operates less as a chatbot and more as an AI project manager under author direction, with every phase's decisions, atomic commits, and verification artefacts preserved on disk as an audit trail. The same loop drafted and compiled this LaTeX document.

- **Nano Banana Pro for landing-page illustration.** The two largest illustrative figures on sensoryops.com/ – the closed-loop process schematic in How-It-Works and the factory floor image – were generated with Google's **Nano Banana Pro** (the Gemini 3 image model) via Google AI Studio. Using a current-generation image model for brand-grade editorial illustration, rather than stock photography or free-form diagrams, is a creative choice: it lets a one-person technical team ship marketing-quality visuals without an illustrator.

The PINN code, every training run, the architectural decisions, the plagiarism distance from teammates' prose, the pricing narrative, and the pilot-phase roadmap are the author's. AI compressed the search for options and accelerated the labour-intensive craft work; judgement, training, and every decision that mattered stayed with the author.

Repository: github.com/pz-421/sensoryops-pinn. Live site: sensoryops.com.

References

- [1] Raissi, M., Perdikaris, P., Karniadakis, G.E. (2019). "Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations." *Journal of Computational Physics*, 378, 686 – 707.
- [2] Wang, S., Teng, Y., Perdikaris, P. (2021). "Understanding and Mitigating Gradient Flow Pathologies in Physics-Informed Neural Networks." *SIAM Journal on Scientific Computing*, 43(5), A3055 – A3081.
- [3] Marc, A., Engasser, J.M., Moll, M., Flayeux, R. (1983). "A kinetic model of starch hydrolysis by α - and β -amylase during mashing." *Biotechnology & Bioengineering*, 25(2), 481 – 496.
- [4] Brandam, C., Meyer, X.M., Proth, J., Strehaiano, P., Pingaud, H. (2003). "An original kinetic model for the enzymatic hydrolysis of starch during mashing." *Biochemical Engineering Journal*, 13(1), 43 – 52.
- [5] De Schepper, C.F., Michiels, P., Buvé, C., Van Loey, A.M., Courtin, C.M. (2021). "Starch hydrolysis during mashing: A study of the activity and thermal inactivation kinetics of barley malt α -amylase and β -amylase." *Carbohydrate Polymers*, 255, 117494.
- [6] Gunawardena, S. et al. (2019). "Modelling and optimisation of the mashing process in beer production with rice adjunct." University of Sri Jayawardenepura.
- [7] Hobani, A.I., Tolba, M.H. (1995). "Bulk thermal conductivity and diffusivity of barley." *Research Bulletin of the Agricultural Research Center, King Saud University*, 73, 5 – 18.
- [8] Statista Market Insights (2025). Alcoholic Drinks – United Kingdom. Statista. Available at: [statista.com/outlook/cmo/alcoholic-drinks/united-kingdom](https://www.statista.com/outlook/cmo/alcoholic-drinks/united-kingdom)
- [9] IWSR Drinks Market Analysis (2025). Premium RTDs Strategic Study. IWSR, London.
- [10] Bloomberg News (2025). "Shift in drinking habits wipes \$830 billion off alcohol stocks." Bloomberg, Global Consumer Staples coverage.