

Fouling Predictions with Quantified Uncertainty

Key Performance Indicators

Model quality

—41%	—58%	—34%	—64%	+38%
prediction error versus a model built on data alone	training data to reach the target accuracy	narrower confidence ranges after calibration	fewer twin evaluations to the best plan	longer warning before the inefficient regime


What the plant gains

—6.9%	—6.6%	—8.7%	+2.2%	—11.9%
furnace fuel fired across the run	carbon dioxide at the fired heater	energy cost attributable to fouling	crude throughput at the furnace limit	avoidable cleaning and lost duty
				
170 days		0.33 K		93%
of routine plant history calibrates the physical parameters, with no trials on the asset		mean coil inlet temperature error against EntroMetrix's high fidelity reference model		of measured operating points fall inside the calibrated confidence range on the plant history

Fouling models for a preheat train come in two kinds, and each lets the operator down in its own way. Models built on physics are faithful to first principles, but their fouling parameters drift as crude and duty change, and they give one number with no sense of how far to trust it. Models built on data alone need long, clean histories, cannot see beyond the conditions they were trained on, and again give a point with no error bar. Yet the decisions that rest on them, when to take a shell out to clean and how hard to fire the furnace, are costly and slow to reverse, and the fouling record is thin and broken by cleans.

EntroMetrix builds a physics informed twin of the preheat train: a mechanistic law of deposit build up and removal at the core, with a small learned correction for what the physics misses, so the model stays accurate on little data and behaves sensibly outside the record. Its physical parameters are calibrated to the plant's own history by Bayesian inference, so every forecast comes with a confidence range, and a fast surrogate over the twin lets the optimiser test cleaning plans in a handful of runs. The work is retrospective and offline, checked against the plant's cleaning history before it informs any decision.

Accurate where data is thin



On the left, the physics informed twin stays on the measured coil inlet temperature across the run while a data only model drifts once it leaves its training data. On the right, the optimiser reaches the best cleaning plan in a handful of twin evaluations where a blind search needs many more.

What the Model Does

- **Physics at the core:** a mechanistic deposit and removal law drives the decline in coil inlet temperature between cleans.
- **Learned correction:** a small machine learning term carries only what the physics misses, so little data is needed.
- **Bayesian calibration:** the physical parameters are inferred from the plant record as ranges, not single values.
- **Confidence on every forecast:** each prediction arrives with a calibrated range and a stated confidence.
- **Fast surrogate:** a Gaussian process over the twin lets the optimiser find good plans in few evaluations.
- **Checked coverage:** the stated confidence is tested against measured operation, so it means what it says.




Figure 1. Furnace fuel, carbon, cleaning and total cost indexed against the plant's current practice, turnaround cleaning plus reactive cleans when the furnace nears its limit. Every term falls under the confidence weighted schedule.

The twin is built from data the plant already holds: daily stream temperatures, flows and duties over a single run of the hot end shells. About 170 days of ordinary history is enough to calibrate it, so no trials or interventions on the asset are needed. Against EntroMetrix's high fidelity reference model it tracks coil inlet temperature to about 0.33 K over a full run, and on a thin, clean punctuated record it roughly halves the error of a data only model while holding its accuracy with less than a fifth of the history.

Confidence is what turns a forecast into a decision. For every shell the twin projects the temperature a clean would recover, the fuel and carbon that saves and the duty lost while the shell is out, each as a range. A shell is cleaned when the low end of that range clears the cost of the clean and it can be isolated while the unit runs; the rest wait for the turnaround. The plan is made for the whole train, because cleaning one shell changes the duty on the others, and it weighs fuel, carbon, throughput and cleaning cost together, returning a frontier of plans the plant chooses from.

The operating gains are measured against the plant's current practice over a full run between turnarounds: about 7 percent less furnace fuel and carbon, about 9 percent off the energy cost of fouling, about 2 percent more crude through the furnace limit and about 12 percent less avoidable cleaning and lost duty. The study is retrospective and offline, and every clean it would have called is checked against what the plant actually did.

What the Model Does (cont.)

- **Blind validation:** calibrated on part of the history and tested on runs it did not see.
- **Interpretable parameters:** deposit, removal and ageing constants an engineer can read and check.
- **Network coupled:** every candidate schedule is scored on the whole train, not shell by shell.
- **Multi objective:** fuel, carbon, throughput and cleaning cost weighed together; the plant picks its point on the frontier.
- **Offline and isolated:** runs on historical data, never in the control loop.
- **Keeps learning:** each new clean and run updates the ranges, so the twin stays current as crude and duty move.

What the Model Delivers




Figure 2. A forward forecast of coil inlet temperature with its confidence band: where the shell is heading and how sure the twin is, which a cleaning or firing decision is weighed against.




Figure 3. The value of cleaning each shell at its best moment, with its 90 percent range. A shell is a priority clean when the low end of its range clears the cost of a clean.




Figure 4. Which shell to clean and when, over a three year run to the next turnaround: six cleans on the priority shells, with windows that widen further out; the other shells wait for the turnaround.

Predictions that Carry their own Confidence

Demonstrated on a crude distillation preheat train

Mohammed Ali, CEO · mo@entrometrix.ai · +44 (0) 773 800 0590