

CDU Energy and Carbon Optimisation

Key Performance Indicators

—5.2%

energy consumed per barrel

—4.6%

carbon dioxide per barrel

—6.8%

fired heater fuel

—9.4%

stripping steam

—3.7%

product quality giveaway



14 to 19%

of total refinery carbon dioxide is attributable to the crude distillation unit, one of its leading scope one sources



71%

of the column heat is recovered by the preheat train before the fired heater fires, so the rest sets the fuel and carbon bill




under 3%

model deviation against the plant's own measured temperatures, so its account of the energy and carbon can be trusted

The crude distillation unit is the largest single energy user on a refinery and one of its leading scope one carbon sources, because most of the heat a barrel needs is raised by burning fuel in the fired heater. Firing rate and excess air, pump around heat recovery, preflash and feed preheat, stripping steam, column pressure and cut points all interact, so a move that protects one product specification can quietly waste fuel elsewhere. To hold every specification safely across changing crude slates, operators run to conservative fixed set points with margin in hand. That margin costs energy and carbon on every barrel, and routine historian trends do not show where it is being lost or how much could be recovered without breaching a specification.

EntroMetrix builds a physics informed machine learning model of the unit's energy and carbon from the plant's own historian data, combining first principles heat and mass balances with learned corrections, then takes an exergy view to locate where useful energy is destroyed rather than where it merely flows. A reinforcement learning agent works against that model and continuously trims the operating levers, furnace excess air and firing, pump around duties, preflash and stripping steam, pressure and cut points, to lower energy per barrel and carbon while every product specification is held and the unit stays inside its safe operating window. Rather than a single answer it returns a Pareto set of policies that trade energy and carbon against throughput and product value, and the plant chooses the point it wants.

The loss map and the policy



An exergy view of the unit, on the left, ranks where useful energy is actually destroyed; the main column and the coolers dominate, and that is where the levers reach. On the right each point is a candidate policy; fixed set point practice is dominated, and the chosen policy sits on the frontier of carbon against throughput the plant selects from.

What the Model Does

- **Physics informed ML:** a physics informed machine learning model of the unit's energy and carbon, built from the plant's own historian data with no new sensors needed to start.
- **Exergy view:** a second law analysis locates where useful energy is destroyed, so effort goes where the loss is.
- **Levers:** furnace firing and excess air, pump around duties, preflash and preheat, stripping steam, pressure and cut points.
- **Multi objective reward:** energy and carbon per barrel are the reward, weighed against throughput and product value.
- **Specifications held:** product boiling points and quality are constraints, pushed toward their efficient limits, never past them.
- **Pareto set:** the agent returns a set of policies on the frontier and the plant selects the point it prefers.




Figure 1. Energy and carbon per barrel, fired heater fuel and stripping steam, baseline against the optimised policy. Every objective falls under one policy, held within product specifications.

The model is built from data the plant already holds. Historian records of temperatures, flows, duties and fuel across the crude unit are enough to build an energy and carbon picture of the process, and a second law, or exergy, view then shows where that energy is destroyed rather than merely where it flows. The operating limits and product specifications are set with the plant's own engineers before any recommendation is made.

On a crude distillation unit modelled and validated to within about three percent of measured temperatures, the optimised policy lowered energy per barrel by about five percent and carbon per barrel by a similar margin, chiefly by trimming fired heater firing and excess air and by tightening stripping steam, while holding every product specification. Most of the saving came from a small number of levers, so the value is reached with instruments the plant already has.

The work is retrospective and offline first. The agent is trained against a fast physics informed surrogate of the unit and tested against operating history before any set point is suggested, and its policy is constrained to product specifications and the safe operating window, so proposals stay operable. Nothing connects to the control system in the proof phase, and the agent's own critic flags when a crude slate change has moved the unit and a refit is due.

What the Model Does (cont.)

- **Retrospective:** learned and tested against operating history first, before any set point reaches the board.
- **Offline and isolated:** the model runs against a copy of the process, with no connection to control systems in the proof phase.
- **Fast surrogate:** a physics informed surrogate of the rigorous model gives a safe, quick training environment.
- **Safe operating window:** the policy is constrained to product specifications and the safe window, so proposals stay operable.
- **Drift monitor:** the agent's own critic flags when a crude slate change has moved the unit and a refit is due.
- **Field readiness:** the proof stage is built to graduate from offline recommendations to a supervised, online set point loop on the crude unit.

What the Model Delivers




Figure 2. Specific energy per barrel across the run, current practice against the optimised policy. Current practice drifts up as heat recovery declines and crude varies, while the optimised policy holds below the run mean, so the gap widens through the run.




Figure 3. Carbon avoided against current practice, accumulating across the run. The monthly avoided firing rises and falls with crude slate and rate and grows as the gap widens, so the total curves upward to a plant scale figure.




Figure 4. Where the saving comes from. Furnace firing and excess air is the largest lever, ahead of pump around recovery and stripping steam, so effort and instruments follow the value.

Lower Energy and Carbon per Barrel

Demonstrated on a crude distillation unit