

Cut-point determination in the vacuum fractionation of odour-critical products

A note on the measurement problem, and what can be inferred without new analysers

Sunil Kamat

September 2026

1. The decisions that determine yield

In batch fractionation of aroma chemicals and terpene fractions under vacuum, most of the yield is decided at three moments: fore-run to main cut, main cut to intermediate, and intermediate to residue. The periods between those transitions are largely a matter of holding conditions steady.

The operator making those decisions typically has a head temperature, one or two tray temperatures, a column pressure reading, a distillate rate, and a gas chromatograph result that arrives every thirty to forty minutes. The specification he is working to is frequently odour and colour rather than assay, and a batch can pass on GC while failing on the nose.

So the decision that determines the yield is taken on a temperature and an intermittent laboratory result, against a specification that neither of them measures.

2. Why head temperature is a weak proxy under vacuum

Head temperature relates to distillate composition through the vapour–liquid equilibrium, but only at constant pressure. Under vacuum that condition is not reliably met.

From the Clausius–Clapeyron relation, $dT/dP = RT^2 / (P \cdot \Delta H)$, where ΔH is the enthalpy of vaporisation. For a typical organic at 20 mbar and 400 K with ΔH near 40 kJ/mol, this gives roughly 1.5 to 2 °C per mbar. A three-millibar excursion — a condenser fouling slightly, a gland leaking, a vacuum pump warming through the shift — therefore moves the boiling temperature by four to six degrees.

Set that against the separation being attempted. Close-boiling terpene fractions differ by only a few degrees at atmospheric pressure and have relative volatilities in the region of 1.05 to 1.15. The temperature signature of the transition the operator is trying to detect is comparable to, and often smaller than, the temperature noise created by ordinary pressure drift.

The operator sees the temperature move and cannot distinguish a change in composition from a change in pressure. In practice he assumes the less favourable explanation and acts on it.

There is a second effect. Head temperature reports the vapour arriving now, which reflects tray compositions established some minutes earlier. During a transition, when composition is changing fastest, the indication lags most.

3. An asymmetry that biases every plant in the same direction

Consider the two ways of being wrong.

Cutting late puts transition material into the prime fraction. The consequence is an odour or colour failure: visible immediately, discussed at length, sometimes a customer complaint, occasionally a requalification. Everyone remembers it.

Cutting early puts prime product into the intermediate fraction. The consequence is a slightly smaller drum. It appears in no alarm, no laboratory report and no shift log. It is visible only in a per-batch mass balance, which most plants do not compute.

One error is loud and the other is silent. After some years of that feedback, operators converge on the same behaviour: cut early, and keep cutting early. This is not carelessness. It is the rational response to asymmetric consequences, and it is why the bias runs in one direction in almost every plant rather than averaging out.

It is also straightforward to test. Assay the intermediate cut. **If it contains a significant fraction of prime product, the bias is real and its size is measurable.** If it does not, the column is already being run close to its limit and there is little here worth pursuing.

4. What can be estimated from instruments already present

The variable the operator actually needs is distillate composition, and continuous online analysis for it is expensive. But it does not have to be measured directly. It can be inferred, continuously and to useful accuracy, from instrumentation most columns already have or can have at modest cost:

- Four to eight temperature points down the column — four-wire RTDs, not dial gauges
- Absolute pressure at top and bottom, accurate at the operating range rather than at atmosphere
- Column differential pressure
- Reflux and distillate flows
- Distillate density

What makes this tractable is the gap between the underlying idealised vapour–liquid equilibrium for terpene and common aroma systems and any particular column — tray efficiency, holdup, heat loss, the gradual ageing of packing — has to be resolved against the plant itself.

The estimate does not need to be exact. It needs to be better than a temperature whose reading is confounded by pressure, and to update continuously rather than every half hour. **Nor is it difficult to check: the plant's own GC results provide the reference, batch after batch.**

5. Reflux is not a constant

A second point, largely independent of the first.

In batch distillation, holding the reflux ratio constant is not the optimal policy. For a fixed distillate purity, the recovery-maximising policy is a reflux ratio that rises through the batch as the still composition depletes. A constant setting over-refluxes at the start, wasting reboiler duty and lengthening the batch, and under-refluxes at the end, pushing prime product into the intermediate cut.

Most plants run a fixed reflux (or depending on the temperature), chosen once for the most difficult case expected, and never revisited.

Related, and cheaper still: many plants analyse every incoming feedstock lot and then make no use of the result. Natural feedstocks vary, and the isomer ratio in turpentine differs by origin and

by season. That assay is a feed-forward signal, already paid for, which can be used to pre-position the column before the batch begins rather than discovering the difference partway through it.

6. What this approach cannot do

It cannot loosen an odour specification. Odour-active minor components must enter the objective function as hard constraints, never as terms to be traded against yield. A scheme that improves recovery by relaxing an impurity limit has not solved the problem; it has moved it downstream to the customer.

It cannot defeat a phase equilibrium. Where a separation is capped by an azeotrope, no control policy changes that. The answer in that case is a different separation, not better control of this one.

It cannot compensate for an under-staged, fouled or flooding column. If the stages are not there, a good estimator will simply report accurately that they are not there.

It does not replace the gas chromatograph. It interpolates between GC results and holds the column steady in between. The GC remains the reference and the arbiter.

And one commercial caution that matters more than any of the above. In flavour and fragrance, a supplier's material is approved against a reference sample. Any change in cut policy that shifts the impurity profile, even in a direction that looks like improvement, can trigger requalification. For that reason the first objective in this kind of work should be reduced batch-to-batch variation, not increased yield. Consistency is what the customer actually pays for, it is the easier internal case to make, and the recovery improvement tends to follow anyway as a consequence of running closer to a well-characterised optimum.

Sunil Kamat is a B.Tech graduate of IIT Bombay and works on model-based supervisory control for thermally constrained and composition-limited separations. An earlier application of the same approach, a hybrid thermodynamic and kinetic model solving several hundred simultaneous equations in real time against live plant data, was deployed on the fourteen carbonation towers at Saurashtra Chemicals, a division of Nirma Ltd, Porbandar, and validated over six months of plant operation. A related article, "Optimisation of chemical processes using hybrid technology — artificial intelligence and process knowledge", appeared in Chemical Weekly on 14 October 2025.

snlkamat@gmail.com · +91 94230 07004