INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 12 Control in Petrochemicals and Polymers

CHAPTER 13

Summary and Outlook

Pages 305–318

Edited by John D. Hedengren

How to cite this chapter. "Summary and Outlook," in *Control in Petrochemicals and Polymers*, John D. Hedengren, Ed., vol. 12 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 13, pp. 305–318.

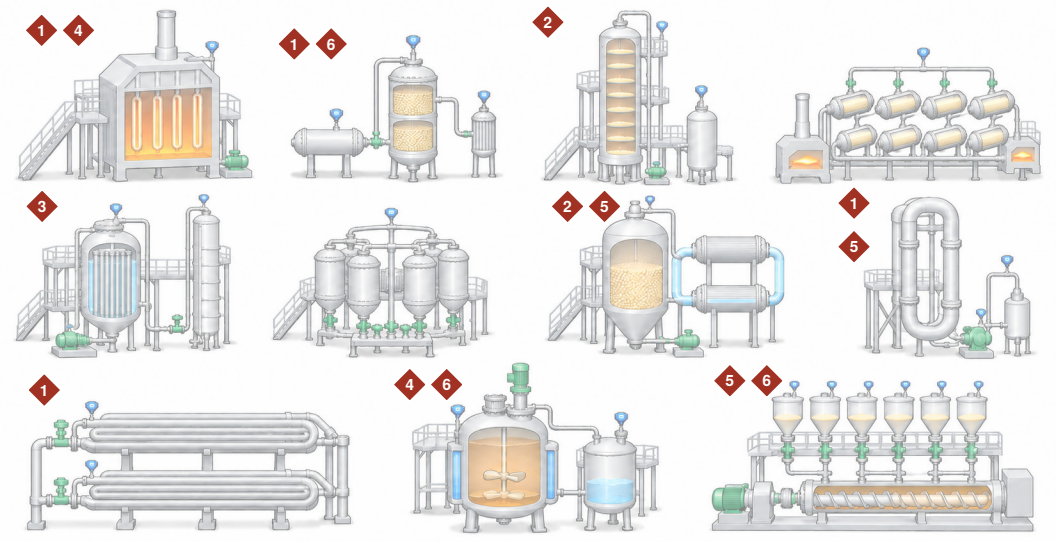
The page numbers in this file are those of the printed volume.

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v12-petrochemicals-polymers/

Copyright © 2026 Advanced Process Solutions, LLC. All rights reserved. Produced by Advanced Process Solutions, LLC.

CHAPTER 13

Summary and Outlook



Open gaps 1 limit where it is 2 inference validity 3 flat optimum 4 run length 5 whole grade changes 6 unweighed holdups

Notice that each number stands beside every unit where a chapter met that gap, so the first gap appears four times and only the dehydrogenation beds and the moving bed carry none. Illustration, not to scale.

IN PLAIN TERMS **What Eleven Petrochemical Units Taught**

The main idea. Eleven chapters followed gas on its way to plastic pellets, from the furnace that cracks it to the feeders that add the last ingredients. Each unit had to keep a product quality that is checked only now and then close to a limit it must never cross.

The challenge. The limit is often measured in the wrong place, so the true value can run ahead of the reading. The quality comes late, and a tester that has stopped can still look normal.

Why it matters. A unit that trusts a wrong reading makes plastic nobody can sell, or shuts down, while every display looks fine.

Eleven chapters followed the chain from the cracking furnace to the additive-dosed pellet. Table 1.2 listed, for each unit, the quality it holds and how often that quality is seen, the limit that binds and the structure that guards it. This chapter adds what worked, what failed and the measured result, and reads the comparison for the question of the volume: how a quality that is rarely measured is held against a safety limit that must never be crossed. It then gives a short guide for choosing a design, looks at where practice is heading, reports how the volume was made with AI agents and what they got wrong, and lists the problems that remain open.

13.1 What the Chapters Showed

Tables 13.1 and 13.2 compare the chapters, split between the units that make and purify monomers and intermediates and the units that polymerize and dose. The second column condenses Table 1.2

Table 13.1 How Each Unit Held a Rarely Measured Quality Against Its Limit, Chapters 2 to 7

Chapter	Quality, how often, limit	What worked	What failed	Measured result
2, cracking furnace	Conversion on a 21 min analyzer. Tube metal behind the coke, which makes a coked pass read cold.	Balancing the passes on tube metal. A decoke at 1090 °C with the tuned outlet loop.	Balancing on coil outlet temperature. Nobody decoking, as shipped.	Runs of 910, 888 and 841 h balanced on metal, not balanced and balanced on outlet (Table 2.4). Challenge 20.67 t/h with no time over the limit, against 858.2 h over (Section 2.6).
3, acetylene converter	Outlet under 1 ppm on a 5 min chromatograph. A window open on one side. Bed rise trip at 70 K.	Feedforwards from the inlet analyzer with a trim on log ppm at 0.1 ppm. A balance inference when the analyzer fails.	A fixed ratio, which breaks through at 1.5 and spends ethylene at 2.0.	Worst sample 0.53 ppm at a loss of 0.297 %, against 185 of 288 samples over at 1.5 and 0.566 % at 2.0 (Table 3.6).
4, propylene splitter	Purity on a 20 min chromatograph against a response of hours. A tray temperature that reads pressure first.	The analyzer on the reflux ratio while it is healthy. The cascade with its validity hold armed from the start.	A cascade that integrates a frozen analyzer until a late hold. An uncompensated tray.	Peak departure 0.063, 0.143 and 0.373 mol% for analyzer, cascade and inferential alone. Frozen analyzer: 98.90 with the late hold, 99.17 with it from the start (Table 4.4).
5, dehydrogenation beds	Conversion once a 3 min slot. Bed limit 680 °C. The handle acts on the next cycle's stored heat.	A cycle-to-cycle PI on the regeneration, with two beds on line after one is lost.	Three beds on line under any control. Fixed regeneration, as shipped.	Worst hour 1.73 points from target, against 3.19 shipped, 8.14 with three on line under the PI and 9.14 under the MPC (Table 5.5).
6, ethylene oxide	Selectivity on a 30 min analyzer. Hot spot, 4.2 °C of coolant from the end of the cool state. Flammable oxygen.	A retuned oxygen loop, all four composition loops closed, the override on the rate.	A chloride PID started above the optimum. The cautious oxygen loop, as shipped.	From 3.30 ppm the PID trips at 381.8 min, from 1.50 it creeps (Table 6.5). Challenge oxygen error 0.105 against 0.314 mol% (Section 6.5).
7, paraxylene moving bed	Both purities as 12 min cycle composites. Edges of the triangle, which move as the adsorbent ages.	Two PIs once a cycle on the logarithm of each impurity. The repetitive MPC.	Fixed ratios at the fresh point, as shipped. Every fixed point on a 0.01 grid.	PIs pass at 32.78 t/h with recovery no lower than 97.47 %. Shipped recovery falls to 82.62 %, best fixed point 27.00 t/h (Table 7.5).

into the quality, how often it is known, and the limit. The last column gives the measured result with the table or section of the chapter that reports it.

The columns make five points about the question the volume asked. Figure 13.1 draws one like pair of results from each chapter's row.

The limit is measured in the wrong place, and the guard is only as good as that reading. On the cracking furnace the coil outlets agreed and the tube metal did not. Balancing on the outlet took feed away from the coked pass that needed it, and the run fell to 841 h against 888 h with no balancing at all. On the acetylene converter the ethylene loss passed one percent 1.58 min after the hydrogen was raised by hand, and the bed rise reached its trip only at 3.95 min (Table 3.5). The selectivity was lost long before the limit was. On the LDPE tube a trip set 5 K below the decomposition line fired 1.5 s before a decomposition that one scan earlier it would have prevented (Section 10.9). The back-off of Equation (1.13) has a term for how far the truth runs ahead of its measurement, and each of these units needed that term.

Table 13.2 How Each Unit Held a Rarely Measured Quality Against Its Limit, Chapters 8 to 12

Chapter	Quality, how often, limit	What worked	What failed	Measured result
8, gas-phase polyethylene	Melt index and density from a laboratory 45 to 105 min old. Sticking temperature and cooling capacity.	The retuned quality cascade, with the bed drawn down to 55 % before the grade change.	The shipped melt-index tuning, which never used its hydrogen limit. A rate cut held through the change.	Off-grade 77.5 against 179.1 t as shipped, and 128.7 t at 20 t/h throughout (Table 8.9). The cascades read the bed quality without delay.
9, polypropylene loop	Melt flow rate on a 6 min analyzer behind a 52 min average. Pump cavitation as hydrogen spends the bubble point.	A bubble-point controller at a low selector. For the challenge, 41 bar first and a slower quality loop.	Quality cascades alone. Pressure alone or the slow tuning alone in the challenge.	Cavitation 42.3 against 0 min, polymer 45.6 against 110.6 t, of which 33.4 and 103.3 t off specification (Table 9.5).
10, tubular LDPE	Melt index every 6 min, density from the laboratory. Decomposition at 335 °C, gas ahead of every thermocouple.	Peak loops at 310 °C with the pulse feedforward. A predictive controller that carries the ceiling.	The feedforward off. Peaks at 312 °C. A supervisor reserve of 20 K blind to the lot change.	Conversion 30.79 % with 16.46 K to spare, against 27.72 as shipped. Without feedforward 13.55, with the supervisor 11.71 K (Table 10.7).
11, butyl rubber	Mooney from the laboratory every 2 h. A jacket walking to its refrigeration floor as the wall fouls.	The Mooney cascade with tuned loops and a wash on a cooling margin of 0.10. A setpoint near the warm optimum.	No wash, as shipped. A wash on a clock written for another fouling rate.	Challenge 3.83 t/h, worst error 0.83 °C, against a trip at 68.5 h (Section 11.6). On specification 4.37 at −93 °C, 1.25 at −101 °C (Table 11.3).
12, additive feeders	Additive content from a composite every 3 h, 45 min late. A sock below the scale that the weight cannot see.	Ratio control with the laboratory trim and the grade preview. Diversion on the sock switch.	Either setting alone. Any estimator against the holdup.	Shift 9574 kg on specification with 450 kg of transition, against none on specification as shipped (Table 12.4). Holdup 403 kg off specification in all four architectures (Table 12.5).

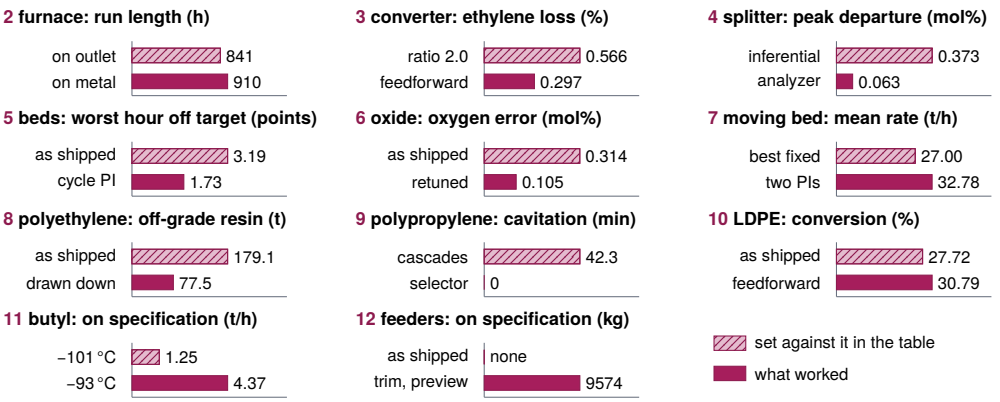


Figure 13.1 One pair of results from each chapter, with the numbers as the "Measured result" column of Tables 13.1 and 13.2 prints them. Each panel has its own scale from zero. A larger value is better in the panels of Chapters 2, 7, 10, 11 and 12. Notice that on the polypropylene pump and the feeders the structure that worked changed all of the result.

A guard buys the equipment, not the product. The polypropylene override kept the pump out of cavitation and the reactor running, and made 110.6 t against 45.6 t. It also left the melt flow rate at 6.55 against a setpoint of 8, so 103.3 t of its polymer was off specification, more than the run that destroyed the pump. Read by off-specification tonnage alone, the two runs rank the wrong way round. The butyl reactor shows the converse. With its conversion loop in manual, as an override that gave way to the cooling margin would leave it, the poisoned reactor lost tonnes faster than it saved hours, 4.00 t/h against 4.12 t/h in automatic (Section 11.6). Whether a guard pays depends on whether the end of a run is abrupt or creeping.

A late quality is held by what is measured early, and lost to a measurement that has stopped. The acetylene converter passed its day only with feedforward from the inlet analyzer, ahead of a 20 min drum. The LDPE pulse feedforward raised the worst margin over two pulses from 33.1 to 41.9 K (Table 10.5), and the grade preview cut the feeders' transition material from 541 to 410 kg (Table 12.4). Where the analyzer is fast against the plant, as a 20 min chromatograph is against a splitter that settles in about twelve hours, it can close the loop directly. The danger is the analyzer that stops. The splitter's cascade that kept integrating a frozen result settled at 98.90 mol%, worse than the tray temperature alone, and the feeders' scale read normal for the whole time that 403 kg of wrong product went to prime.

An optimum is not a setpoint. The chloride of the ethylene oxide reactor is flat in selectivity and steep in hot spot, 22.6 °C across the same two parts per million (Table 6.4). A PID with one tuning creeps toward its setpoint from below the optimum and trips the reactor from above it. The furnace's best severity, near 850 °C at 20.31 t/h per elapsed hour (Table 2.5), is visible to no continuous instrument, and the butyl reactor's warm optimum is bounded by the Mooney grade and not by the equipment. On the LDPE tube the grade set the floor as well: peaks at 305 °C kept ample margin and let the density leave its band. Each of these optima is found by running the unit to the end of a campaign or a grade, which no loop does.

Cycles and grade changes are prepared, not chased. Prediction paid where it carried a ceiling. With three dehydrogenation beds on line no controller met the target, because the regeneration could not reach it under the bed limit. Drawing the polyethylene bed down before the grade change cut the off-grade resin from 179.1 to 77.5 t, and the polypropylene challenge passed by raising the pressure first. The predictive controllers were mixed. On the LDPE tube the plan that carried the ceiling converted 31.54 % with 12.8 K to spare against 31.09 for the loops with the requests clipped by hand (Table 10.6). On the coking furnace it cut the time above the tube limit from 260.5 to 228.5 h (Table 2.6). It lost to the prepared polyethylene cascades, 85.3 against 80.9 t (Table 8.10), and tied the dehydrogenation PI. On the polypropylene loop most of its gain came from a colder reactor, and the margin written as a constraint was worth 18.6 t on the same temperature path (Section 9.6.6). Volume 9 treats that choice in full.

> ⧖ HISTORY **1991 to 2026: The Laboratory Still Corrects the Model**
>
> The timeline of Chapter 1 closed in 1991 with a scheme that predicted the melt index and density of an industrial polyethylene reactor from its temperatures and gas compositions, and corrected the model with infrequent laboratory results [1]. A 2013 review of soft sensors found model maintenance the most critical issue in their industrial use. It listed a recursive model with a bias update that predicted melt index through grade changes, and noted that recursive methods cannot cope with abrupt changes [2]. In 2026 a study of 425 industrial polypropylene runs found that models lost accuracy when moved between product families [3], and a second study kept the laboratory method as the certification reference, with the model estimating between results [4]. The design of 1991 is still the design. Gap 2 of Section 13.5 asks what it lacks.

13.2 Choosing a Method

The chapters suggest an order of questions for a unit that holds a rarely measured quality near a limit. Each names the chapter where its answer decided the result. Figure 13.2 sets the same questions out as a flowchart.

1. *Where is the limit measured, and how far can the truth run ahead of it?* Size the setpoint from Equation (1.13) with that distance in it, and balance or protect on the variable nearest the limit (Chapter 2, Chapter 10).
2. *Is the guard one-way, faster than the quality loop, and told what was applied?* A low selector with tracking keeps the losing loop from winding up, and its indicator should say whether the constraint holds the ceiling, not only whether it won the last scan (Chapter 3, Chapter 9).
3. *Does the process gain keep its sign over the range?* Near an optimum it does not. Hold a fixed policy, search slowly, or keep the handle on the side where the sign is known (Chapter 6).
4. *How late is the quality against the plant's own time scale?* An analyzer fast against the plant can close the loop. A slow one needs the load fed forward from what is measured early (Chapter 4, Chapter 3).
5. *What happens when the analyzer stops or lies?* Arm a validity hold, carry an inference through the outage, and never let a laboratory trim learn from a sample that the diagnostics already distrust (Chapter 4, Chapter 12).
6. *Does the unit run in cycles?* Correct once a cycle on the cycle's report, with the gain the delay of reports allows (Section 1.4.7), and first check that the cycle can reach the target at all (Chapter 5, Chapter 7).
7. *Is the binding constraint the grade or the equipment?* Find the optimum per elapsed hour, wash or decoke on the quantity that ends the run, and expect the grade to bound it (Chapter 2, Chapter 11).
8. *Can preparation do what control cannot?* Draw the bed down, move the pressure, preview the grade (Chapter 8, Chapter 9, Chapter 12). Then ask what a predictive controller adds over the prepared structure at the same target and ceiling.

13.3 Where Practice Is Heading

Inferred quality learns from data. Soft sensors are widely built in the process industries for quality prediction and monitoring, and a 2021 survey reviews deep learning for them, the demands it meets in practical applications and the problems still open [5]. In the editor's view the inference will grow more capable faster than the logic that decides when to trust it. The chapters of this volume failed at that logic, not at the inference: the splitter's cascade integrated a stale analyzer, and the feeders' trim had to be kept from learning a holdup as a loss.

Nonlinear prediction reaches polymer reactors. Nonlinear model predictive control has been reported on an industrial polymerization process, where it controlled polymer quality across production rates better than reactor temperature targets computed offline from the same model [6]. For tubular LDPE reactors, optimization-based strategies combining nonlinear predictive control with fast state estimation were proposed to regulate the process and its economics under fouling [7]. In the editor's view such controllers will carry more of the limits of this volume in their plans, as the LDPE controller of Chapter 10 carried its ceiling. The comparisons of this volume suggest testing each against a structure given the same preparation.

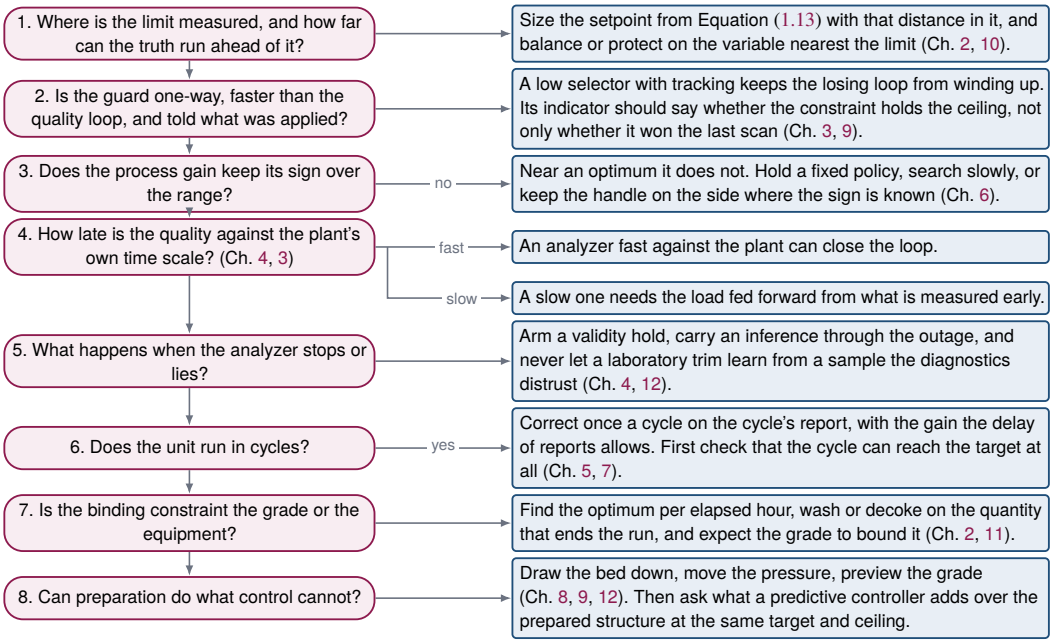


Figure 13.2 The eight questions of this section in order, with the list's answer and the chapters where it decided the result. Question 4 branches on the list's two cases. Notice that questions 1, 4 and 5 ask about a measurement before any structure is chosen.

Cyclic units are operated by optimization. Optimization-based control of a reactive simulated moving bed was reported in 2004 [8], and a 2007 article treats feedback control for optimal process operation [9]. The paraxylene ring of Chapter 7 found no fixed point that passed sixty cycles of aging, while two PIs on cycle reports did. In the editor's view the correction from one cycle to the next will stay the core of such control, with an optimizer setting what it corrects toward.

The cracker's fuel and feed change. Steam cracking emits large quantities of carbon dioxide, and a 2020 review names enhanced heat transfer in the radiation section, green energy and novel furnace designs as the keys to reducing it [10]. Pyrolysis oils from plastic waste are expected to become a steam cracker feed, but their contaminants drive corrosion, fouling and downstream catalyst poisoning, and the crackers rely on a stable and predictable feed [11]. In the editor's view both changes act on the limits of this volume. A new firing pattern moves the tube metal that Chapter 2 balances, a fouling feed shortens the run whose optimum it finds, and a poisoned converter moves the window of Chapter 3.

Electric heat reaches the cracker. Industrial steam crackers use gas-fired furnaces, which emit about 366 Mt of carbon dioxide a year worldwide. A 2022 review finds that improvements to the present designs are limited and cannot reach the reductions required. Electric heating avoids the flue gas losses, and resistance heating, promising for naphtha, is relatively easy to scale up and can retrofit existing crackers. It still has to be tested and demonstrated, and a cracker's heat is highly integrated: its steam drives the compressor turbines [12]. In the editor's view the tube metal limit of Chapter 2 would stay, held by power instead of fuel.

Spectroscopy moves into the melt. Near-infrared and Raman spectroscopy of molten polymer, measured in line, have been researched intensively and are still sparsely applied in industry, because

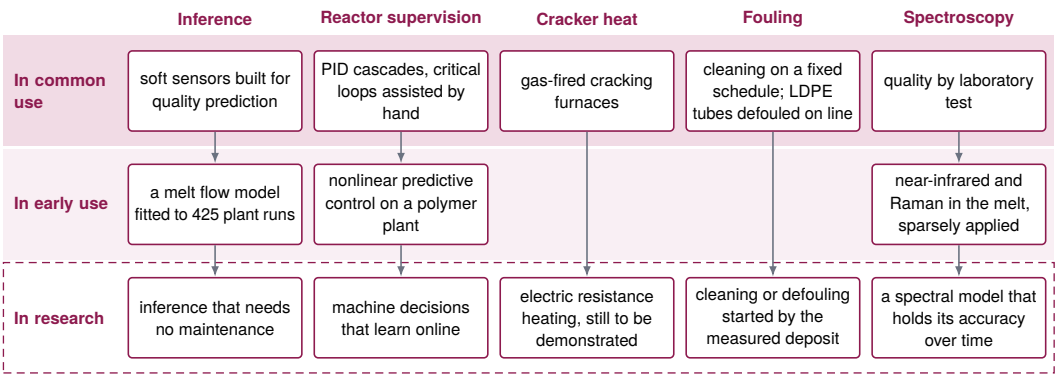


Figure 13.3 The developments of this section in five threads. The top lane holds what the cited source calls standard, established or current practice, the middle lane what it reports in service at some sites or on plant data, and the bottom lane what it reports from simulation or as a need. An empty cell has no cited development. The placing is the editor's reading of the sources, and no box carries a date. Sources by thread: inference [2, 3, 5], reactor supervision [6, 17], cracker heat [12], fouling [15, 16] and spectroscopy [4, 13, 14]. Notice that no thread has a box in common use that reads the product or the deposit in line.

their readings are unpredictable and unstable [13]. A 2021 review of extrusion monitoring reports a study of polymer additive content in which a model trained on near-infrared spectra from one time did not keep good accuracy on spectra taken at other times [14]. In the editor's view such a reading would shorten the three-hour wait of Chapter 12, once it passes the validity test of gap 2.

Cleaning follows the deposit. A 2025 review of fouling in thermal processing describes sensors that detect the deposit as it forms and can trigger cleaning on the actual buildup rather than on the conventional fixed schedule [15]. For tubular LDPE reactors, which are defouled on line by pressure cleaning or thermal shock, a 2022 study let a soft sensor of the layer's thickness start the defouling when the layer passed its limit, under a nonlinear predictive controller, on a process simulation [16]. The butyl wash of Chapter 11 applies the idea to a deposit inferred from the jacket. Figure 13.3 sets the developments of this section on three lanes, by how far each has come into use.

13.4 Engineering With AI Agents

An AI agent is a language model that works toward a goal in a loop. It plans, calls tools such as a compiler, a simulator or a browser, reads the results and continues [18]. In the editor's view, agents can now carry out much of the work that used to fill the hours of a control project. They derive a model from balances or fit one to plant data. They write the simulator and its tests, code the controller, tune it against stated criteria, script commissioning checks, write the documentation and watch running loops for degradation. The apps and chapters of this series were produced that way under the editor's direction, and the account that follows this section reports what the agents got wrong in this volume and how it was caught.

In the editor's view, this changes where an engineer's value lies. When a model, a controller and a report can each be produced in minutes, producing them is no longer the scarce skill. The work moves to five things.

Quality control. An agent's output arrives as a claim: this model fits, this tuning settles in four minutes, this code implements the equation. The engineer checks claims against the plant or the

simulator. A reported number is measured again on a fresh run at the stated seed. The code is read for what it clamps, what it initializes and what it reads, because a display that shows the controller's model in place of the plant looks correct until the two differ. A tuning is tried on a disturbance the agent did not test. The most useful checks are independent of the author, made by a second agent or a person who has not seen the first answer, and they are checks that can fail.

Supervision. The engineer sets the scope of each task and decides what an agent may touch. Reading a historian and proposing a tuning is one level of access. Writing a setpoint, downloading a configuration or changing an interlock is another, and it passes through the plant's management of change with a named approver. Supervision also fixes the order of work: a baseline before a change, one change at a time and a comparison that shows what moved.

Creativity. An agent answers the question it is given. Choosing the question remains the engineer's task: which loop is worth improving, whether a disturbance calls for feedforward, a cascade or a change to the process, which step test the operation can tolerate, and which comparison would show that a predictive controller earns its place over a retuned PID loop. The experiment that separates two explanations is designed by someone who can imagine both.

Process insight. A model contains what its author put in. The equipment does more. A valve sticks, the dead time of an analyzer grows when its sample flow drops, a thermowell lags, a heat exchanger fouls and an operator intervenes at the same hour each night. The engineer who knows the unit recognizes a result that is too good, asks what the model left out and knows which measurement to distrust first.

Accountability. Responsibility does not transfer to a tool. Functional safety standards assign the activities of the safety lifecycle to persons and organizations of stated competence, and they require an assessment that includes someone independent of the design team [19]. An agent may draft the hazard analysis or the test procedure. A person owns the decision to accept it, to start up and to keep running.

These five are the professional form of the habit introduced in Section 1.7: check before trusting, and know enough to check.

13.4.1 How This Volume Was Made

Each of the eleven apps was built and tested by agents. An agent then wrote the paper on it, keeping verification notes that trace every number to a run, and a referee report was written on each paper. Four of the eleven reports are titled self-referee reports. The others state a reviewer's stance without saying who wrote them. Seven papers had a further referee pass when a predictive controller was added. Defects went into a log for each app with a severity, a cause, a suggested fix and a status. For this book one agent per chapter reran the numbers its exercises use. The editor set the scope of each round and the rules on what each agent could change, and statements that runs contradicted were left for the editor to decide.

The defect logs. Ten of the eleven logs hold 99 entries with a severity. Of these, 37 are wrong behavior or output of an app, 47 are documentation, 12 are cosmetic or display and 3 are other kinds, such as a fragile test. The splitter's log is written under headings without a severity field and is not counted. Two logs hold a single entry each, and four apps keep a second log in their own folders, which is not counted either. Five cases show the errors that matter for a volume about quality and limits.

An inferred margin that read the truth. The polypropylene override and its predictive controller acted on a bubble-point margin described in the code, the app and the paper as inferred from the pressure transmitter, the thermowell and the chromatograph. Only the pressure was a measurement. The bubble point came from the plant's true composition. An audit that compared two plants with identical readings, one with a quarter more light inerts, found the margins 0.29 bar apart after 1.5 min and the hydrogen valve 16.8 % lower within 12 min. The margin is now computed from instruments, and every number of a run in which the override or the predictive controller acts moved. An audit of the same kind found the acetylene converter's predictive model built from the run's own kinetic settings. Doubling a rate constant moved its ratio setpoint from 1.827 to 1.906 with nothing measured changing. Three more entries, on the dehydrogenation and moving-bed apps, record start-up values and an operator aid taken from the plant rather than its instruments, with no shipped run moved.

A laboratory trim that learned a holdup. On the feeders a powder holdup below the scale made the first assay read 626 ppm against a prediction of 1000 ppm, and the trim took the shortfall as a lasting loss of delivery. After the sock was cleaned the lane was asked for 1/0.813 of its recipe, the product averaged 1217 ppm against a band edge of 1100 ppm, and the estimate still read 1000 ppm while the product stayed on prime. By 3 h the off-specification prime reached 1741 kg against 598 kg with the trim off. The paper agent found it on a second reading, reported it with the fix, and the owner approved the change. A result taken under the holdup alarms now leaves the trim at zero.

A selector that looked idle, and a cycle that did not exist. The polypropylene referee report made three findings that bear on the selector. The indicator that said the constraint "has the pen" was true for 34.8 of the 198 minutes after selection began, while the margin held a mean of 1.99 against its 2.00 bar setpoint, because tracking pulls the losing controller onto the winning value. The app's source, its guide and the brief all claimed a forty-minute limit cycle when the idle constraint controller was parked at the top of its range. The referee emulated it and found none. The draft's central comparison changed the architecture and the tuning at once, and the comparison with one variable changed became Table 9.5.

Protections that did not count what they guarded. The furnace's tube metal could jump past its limit at the start of a decoke, to 1109.0 °C on one noise seed, while the over-limit tile and the alarm skipped decoking time and the run passed. The feed is now withdrawn no faster than the firebox cools, and decoking time counts. On the LDPE tube the decomposition of case 5 turned on one thermocouple scan, so a change in the order of random draws could have made the case uncompletable. The text now says so, and a test fails if the case stops decomposing.

A steady state that had not converged. The ethylene oxide app computes "the steady state with every feed fixed" with a helper that stops after 220 iterations. The chapter agent found that off the design point the helper had not converged. The oxygen margin printed at a coolant of 241.0 °C in the fixed-feed column of Table 6.6, 0.39 mol%, is −0.50 mol% when the solve is run to convergence, inside the trip guard, and a long run of the app with every station in manual trips on oxygen at 1279.5 min. The slow part is the purge, with a time constant near 2000 min. The printed values stay as the paper has them, under the rule that leaves a contradicted statement to the editor, and the defect is logged for the editor and the app owner. The fixed-feed column is therefore not a measured result, and this chapter uses none of its values. The chapter's lesson holds and is stronger: cooling with the feeds fixed eventually trips the unit. The chapter agents that followed were told to check the convergence of any helper behind a number they used. Five of their records report the check, and in each the helper converged.

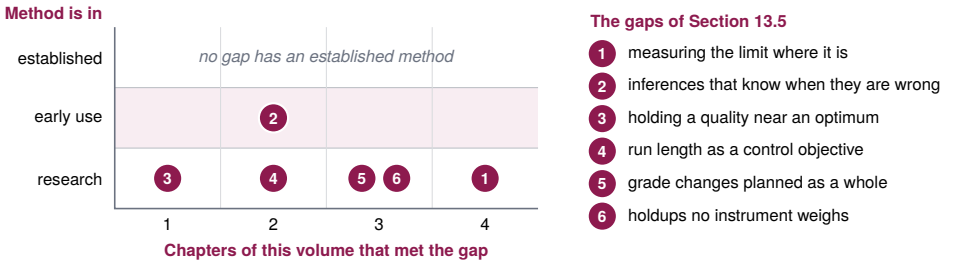


Figure 13.4 The six gaps of this section. Across is the number of chapters that each item names as evidence: four for gap 1 (Chapters 2, 3, 9 and 10), two for gap 2 (4 and 8), one for gap 3 (6), two for gap 4 (2 and 11), three for gap 5 (8, 9 and 12) and three for gap 6 (3, 11 and 12). Up is the maturity of the method available, which is the editor's reading of the sources cited in each item and not a measurement. Notice that only the gap of inferred qualities has a method in use anywhere.

The book round. Each chapter agent first confirmed that its paper had not changed since conversion, and none of the eleven had. The eleven chapter records list 170 rows of numeric verification, each a stated value or group of values set beside its reproduction. Ten records report that every rerun number reproduced, and the eleventh is the ethylene oxide record above. The book round's defect list holds ten entries for this volume. Beside the ethylene oxide helper they are mostly wording, display and stale numbers in the apps' own texts. One briefing calls a speed of 200 min/s "200 times", which is 12 000 times real time, and one exported column holds ten times the chloride its name gives. Three chapter statements were aligned with their own tables. The acetylene conclusion had the loss passing one percent 1.58 min before the trip, where the trip came 3.95 min after the change. The agent writing Chapter 1 first claimed that each chapter meets at least two of the industry's defining problems, and the splitter meets one.

Every error above was caught by a check that could fail: two plants with the same readings, a selector flag set beside the margin it should explain, a cycle emulated instead of quoted, a solver run to convergence and then against the dynamic app. Four working rules follow. Ask what every guard reads, and from which instrument. Count the time a protection is meant to prevent, including the time it is busy. Do not let a quality correction learn from a sample its own diagnostics distrust. Run every helper to convergence before printing what it returns.

13.5 Gaps in Petrochemical and Polymer Control

Six gaps belong to the control of petrochemical and polymer units, and the gaps common to agent-era practice follow them in Section 13.5.1. In the editor's view none is closed. A sentence on difficulty or progress that carries no citation is the editor's view. Figure 13.4 places the gaps by the chapters that met them and by the maturity of the method available, and the same numbers mark the units in the picture that opens this chapter.

1. *Measuring the limit where it is.* The furnace's tube metal sits behind coke, the LDPE gas runs ahead of every thermocouple and the pump's margin is inferred from a composition (Chapter 2, Chapter 10, Chapter 9). On the acetylene converter the selectivity was lost minutes before the bed rise reached its trip (Chapter 3). Methods that bound the distance between the measured and the true limit, and size the back-off from that bound, are not routine. The Open challenge box after this list takes the gap further.

2. *Inferred qualities that know when they are wrong.* The splitter's cascade trusted a frozen analyzer until a hold was armed (Chapter 4), and the polyethylene cascades of Chapter 8 read the bed quality without delay, a luxury no plant has. A validity layer that ranks the analyzer, the laboratory and the inference, and switches between them without a bump, is still designed by hand for each

unit. A 2026 study of industrial polypropylene notes that most published soft sensors are validated within a single grade family or a limited operating region [3]. It is hard because recursive updates cannot cope with abrupt changes [2], and a grade change is one. Progress would be an inference that says when it has left the region of its data.

3. *Holding a quality near an optimum.* The ethylene oxide chloride defeats a feedback loop at its optimum, and a search hunts across the top (Chapter 6). A method that holds a flat optimum next to a steep limit, with a guarantee on the limit, remains open. A 2022 review of optimization by feedback explains why: near the optimum the steady-state relation between input and cost is typically flat, so a gradient fitted to plant data is least accurate exactly there. Search without a model converges in hours to days on chemical processes, one of the main reasons it is not used widely there [20]. Progress would be a search that keeps a stated margin on the steep side.

4. *Run length as a control objective.* The furnace's best severity and the butyl reactor's wash are decided over a campaign that no loop sees (Chapter 2, Chapter 11). Joining the decoke and wash schedule to the loops that spend the margin, with the coking or fouling rate estimated as it changes, is an engineering problem with little routine practice behind it. Across reported coke properties that differ by orders of magnitude, a 2020 study found that the predicted run of an industrial ethane furnace ranged from 20 to 227 days [21]. A 2025 review finds models still far from scheduling whole exchangers over their cleaning cycles [15]. The gap is hard because the rate that ends the run is the least known number in the model. Progress would be a run planned on a fouling rate estimated from the running unit and revised at each decoke or wash.

5. *Grade changes planned as a whole.* The polyethylene bed drawdown and the polypropylene pressure were preparation decided by hand, and the preview of the feeders cut transition material by about a quarter (Chapter 8, Chapter 9, Chapter 12). A planner that chooses inventory, pressure and the quality trajectory together, within the limits of every unit downstream, is not yet standard. A 1992 study optimized the grade changes of a gas-phase polyethylene reactor on a model: lowering the bed level and the catalyst feed during a change cut the off-specification polymer significantly, and the policy needed feedback, because disturbances and model mismatch move the trajectory [22].

6. *Holdups no instrument weighs.* The feeders' sock, a bed's green oil and a jacket's fouling are inventories inferred from their effects (Chapter 12, Chapter 3, Chapter 11). Estimating them from torque, temperature or a laboratory result, and saying when the estimate cannot see them, is needed before a correction is allowed to learn. A 2025 review finds a notable gap in the in-line measurement and prediction of fouling [15], and on LDPE reactors fouling of the inner wall is difficult to assess and its prediction is not systematically formulated in normal practice [16]. It is hard because a holdup shows only through effects that a change of feed also produces.

> ⚠ **OPEN CHALLENGE** **A Limit Read Behind the Truth**
>
> Gap 1 is the one most chapters met. A 2021 review of reactor runaway finds that knowledge of the phenomenon has improved a lot but is not fully integrated into practice, and it names as a research direction the handling of parameter uncertainty in detecting a runaway during operation [23]. A margin kept from a constraint to avoid violating it in transients costs a loss that grows linearly with the margin [20]. Progress would be a stated bound on how far the truth can run ahead of each reading, carried into the setpoint and the trip.

13.5.1 Gaps Common to Agent-Era Practice

The final chapter of each volume lists gaps specific to its topic. Seven more are common to every topic in the series. In the editor's view none of them is closed.

1. *Verifying agent-written control code and models before they reach a plant.* Tests show that a controller behaves on the cases tested. For controllers learned from data, guarantees of stability and constraint satisfaction rest on assumptions about the model and its uncertainty, and open problems remain [24]. In the editor's view, code written by an agent has no comparable body of method yet. Practice needs a routine path from an agent's draft to code a plant can trust: independent derivation, tests of what the code reads and clamps, and trials with hardware in the loop.

2. *Trustworthy models from sparse or poor plant data.* Plant records are short, taken in closed loop, noisy and marked by unmeasured disturbances. A fitting routine returns a model from any record. What is missing is a dependable statement, delivered with the model, of where it can be trusted, and a method that says so when the data cannot identify a parameter.

3. *Safety cases and standards for agents in commissioning and operation.* Functional safety standards define a lifecycle of specified requirements, verification and assigned human roles [19]. In the editor's view, how an agent that writes configuration, runs commissioning checks or advises an operator is qualified, limited and audited under them is not yet settled.

4. *Security of agent access to control systems.* An agent with credentials to a control system is a new path into it. Industrial security standards state the requirements a control system must meet at each security level [25]. An agent that reads documents, messages or web pages can also be steered by text placed there. In the editor's view, least privilege, read-only access by default and human approval of every write are a starting point and not a solution.

5. *Benchmarks and reproducibility.* An agent asked twice may deliver two different controllers, and a result reported without its plant, seed and scoring cannot be checked. Control needs shared benchmark plants with fixed disturbances and published measures, on which agent-built and conventional designs are compared like for like. The apps of this series, with a default seed and an exported record of every run, are a start.

6. *Transfer from simulation to hardware.* A controller developed on a simulator inherits what the simulator omits. The robotics literature documents the gap between simulated and real performance and the methods used to narrow it, such as randomizing the simulated plant and adapting on the real one [26]. Process plants add slow dynamics, rare upsets and little freedom to experiment, so each transfer still needs a staged plan from the simulator to a bench rig to the unit.

7. *Keeping engineers' judgment sharp when agents do the routine work.* Automation removes the routine work through which skill is built and leaves the person the hard cases, an irony described for process operators in 1983 [27]. The same now applies to engineers. An engineer who has never tuned a loop by hand or closed a balance on paper is poorly placed to review an agent that has. Deliberate practice, of the kind the exercises in this book provide, has to be designed into education and into work.

KEY POINTS

- A limit is often read in the wrong place, and a guard is only as good as that reading.
- A guard buys the equipment, not the product.
- A late quality is held by what is measured early, and lost to a measurement that has stopped.
- Cycles and grade changes are prepared, not chased, and an optimum is not a setpoint. Prediction paid where it carried a ceiling.

References

[1] K. B. McAuley and J. F. MacGregor, "On-line inference of polymer properties in an industrial polyethylene reactor," *AIChE Journal*, vol. 37, no. 6, pp. 825–835, 1991.

[2] M. Kano and K. Fujiwara, "Virtual sensing technology in process industries: Trends and challenges revealed by recent industrial applications," *Journal of Chemical Engineering of Japan*, vol. 46, no. 1, pp. 1–17, 2013.

[3] J. Hernández-Fernández and J. Lopez-Martinez, "Reactor-grade non-additive polypropylene architecture defines the physicochemical limits of melt flow index model transferability: An industrial chemometric and machine learning study," *Polymers*, vol. 18, no. 15, p. 1880, 2026.

[4] J. Hernández-Fernández, "Temporally validated and laboratory-reconciled soft sensing of xylene-soluble fraction in industrial polypropylene manufacturing," *Polymers*, vol. 18, no. 19, p. 2401, 2026.

[5] Q. Sun and Z. Ge, "A survey on deep learning for data-driven soft sensors," *IEEE Transactions on Industrial Informatics*, vol. 17, no. 9, pp. 5853–5866, 2021.

[6] R. Bindlish, "Nonlinear model predictive control of an industrial polymerization process," *Computers & Chemical Engineering*, vol. 73, pp. 43–48, 2015.

[7] V. M. Zavala and L. T. Biegler, "Optimization-based strategies for the operation of low-density polyethylene tubular reactors: Nonlinear model predictive control," *Computers & Chemical Engineering*, vol. 33, no. 10, pp. 1735–1746, 2009.

[8] A. Toumi and S. Engell, "Optimization-based control of a reactive simulated moving bed process for glucose isomerization," *Chemical Engineering Science*, vol. 59, no. 18, pp. 3777–3792, 2004.

[9] S. Engell, "Feedback control for optimal process operation," *Journal of Process Control*, vol. 17, no. 3, pp. 203–219, 2007.

[10] I. Amghizar, J. N. Dedeyne, D. J. Brown, G. B. Marin, and K. M. Van Geem, "Sustainable innovations in steam cracking: CO_2 neutral olefin production," *Reaction Chemistry & Engineering*, vol. 5, no. 2, pp. 239–257, 2020.

[11] M. Kusenberg, A. Eschenbacher, M. R. Djokic, A. Zayoud, K. Ragaert, S. De Meester, and K. M. Van Geem, "Opportunities and challenges for the application of post-consumer plastic waste pyrolysis oils as steam cracker feedstocks: To decontaminate or not to decontaminate?" *Waste Management*, vol. 138, pp. 83–115, 2022.

[12] M. E. H. Tijani, H. Zondag, and Y. Van Delft, "Review of electric cracking of hydrocarbons," *ACS Sustainable Chemistry & Engineering*, vol. 10, no. 49, pp. 16 070–16 089, 2022.

[13] S. Zhu, Z. Song, S. Shi, M. Wang, and G. Jin, "Fusion of near-infrared and raman spectroscopy for in-line measurement of component content of molten polymer blends," *Sensors*, vol. 19, no. 16, p. 3463, 2019.

[14] N. Munir, M. Nugent, D. Whitaker, and M. McAfee, "Machine learning for process monitoring and control of hot-melt extrusion: Current state of the art and future directions," *Pharmaceutics*, vol. 13, no. 9, p. 1432, 2021.

[15] J. Tarapata, T. R. Murphy, E. W. Finnegan, T. F. O'Callaghan, and J. A. O'Mahony, "Approaches for measuring and predicting fouling during thermal processing of dairy solutions," *Comprehensive Reviews in Food Science and Food Safety*, vol. 24, no. 4, p. e70209, 2025.

[16] F. S. Rohman, M. R. Othman, D. Muhammad, A. Azmi, I. Idris, R. A. Ilyas, S. E. Elkhatif, and M. N. Murat, "Nonlinear control of fouling in polyethylene reactors," *ACS Omega*, vol. 7, no. 44, pp. 39 648–39 661, 2022.

[17] Y. Liu, "Natural intelligence principle of supervisory-level control for condensed-mode polyethylene fluidized bed reactors: A proportional-integral-derivative perspective," *ACS Omega*, vol. 11, no. 39, pp. 60 064–60 072, 2026.

[18] L. Wang, C. Ma, X. Feng, Z. Zhang, H. Yang, J. Zhang, Z. Chen, J. Tang, X. Chen, Y. Lin, W. X. Zhao, Z. Wei, and J. Wen, "A survey on large language model based autonomous agents," *Frontiers of Computer Science*, vol. 18, no. 6, p. 186345, 2024.

[19] International Electrotechnical Commission, "Functional safety – Safety instrumented systems for the process industry sector – Part 1: Framework, definitions, system, hardware and application programming requirements," IEC 61511-1:2016, Edition 2.0, 2016.

[20] D. Krishnamoorthy and S. Skogestad, "Real-time optimization as a feedback control problem – a review," *Computers & Chemical Engineering*, vol. 161, p. 107723, 2022.

[21] M. Geerts, S. H. Symoens, P. A. Reyniers, G. B. Marin, M.-F. Reyniers, and K. M. Van Geem, "Steam cracking coke properties and their influence on furnace run length predictions: Experimental and modeling study," *Industrial & Engineering Chemistry Research*, vol. 59, no. 52, pp. 22 460–22 472, 2020.

[22] K. B. McAuley and J. F. MacGregor, "Optimal grade transitions in a gas phase polyethylene reactor," *AIChE Journal*, vol. 38, no. 10, pp. 1564–1576, 1992.

[23] A. Kummer and T. Varga, "What do we know already about reactor runaway? a review," *Process Safety and Environmental Protection*, vol. 147, pp. 460–476, 2021.

[24] L. Brunke, M. Greeff, A. W. Hall, Z. Yuan, S. Zhou, J. Panerati, and A. P. Schoellig, "Safe learning in robotics: From learning-based control to safe reinforcement learning," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 5, pp. 411–444, 2022.

[25] International Electrotechnical Commission, "Industrial communication networks – Network and system security – Part 3-3: System security requirements and security levels," IEC 62443-3-3:2013, Edition 1.0, 2013.

[26] W. Zhao, J. P. Queralta, and T. Westerlund, "Sim-to-real transfer in deep reinforcement learning for robotics: A survey," in *Proc. 2020 IEEE Symposium Series on Computational Intelligence (SSCI)*, 2020, pp. 737–744.

[27] L. Bainbridge, "Ironies of automation," *Automatica*, vol. 19, no. 6, pp. 775–779, 1983.

13.6 Exercises

13.1 Discussion · 8 min

An agent proposes a pass-balancing scheme for a new cracking furnace that equalizes the coil outlet temperatures, and reports that its simulation holds the four outlets within half a degree. Using Chapter 2 and Chapter 10, explain why that report does not show the scheme is safe, and list what you would ask for before accepting it.

13.2 Discussion · 7 min

A plant manager ranks two runs of the polypropylene loop of Chapter 9 by off-specification tonnage and concludes that the override made things worse. Use Table 9.5 to show why the ranking is wrong, and propose three measures a scorecard for a quality loop under a safety limit should report beside off-specification tonnage.

13.3 Discussion · 8 min

Three chapters show a quality measurement failing quietly: the frozen chromatograph of Chapter 4, the assay taken during a holdup in Chapter 12, and the cascades of Chapter 8 that read the bed quality without delay. For each, say what the loop believed, what was true, and what logic would have protected it. Then state a general rule for when a quality correction may learn from a measurement.

13.4 Discussion · 7 min

The chloride of Chapter 6 has a flat selectivity optimum near 2.5 ppm and a steep hot spot. Explain why a PID controller on selectivity cannot hold that optimum, why a step-and-search optimizer hunts across it, and argue for an operating policy that respects the hot spot limit.

13.5 Discussion · 10 min

Section 13.4.1 reports that an app's steady-state helper stopped after 220 iterations, so a printed oxygen margin of 0.39 mol% was −0.50 mol% when run to convergence. Explain why such an error passes ordinary review, describe the checks that caught it, and write a procedure of five to eight steps for an agent that computes steady states or limits that a person will rely on.

13.6 Discussion · 8 min

A vendor claims its predictive controller cut off-specification polymer on a polypropylene grade change from 83.2 to 43.7 t. Using Table 9.6, Chapter 8 and Chapter 10, describe the comparison you would require before crediting the controller, and say what each chapter's fair comparison showed.