INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 8 Batch, Run-to-Run, and Hybrid Control

CHAPTER 16

Summary and Outlook

Pages 377–390

Edited by John D. Hedengren

How to cite this chapter. "Summary and Outlook," in *Batch, Run-to-Run, and Hybrid Control*, John D. Hedengren, Ed., vol. 8 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 16, pp. 377–390.

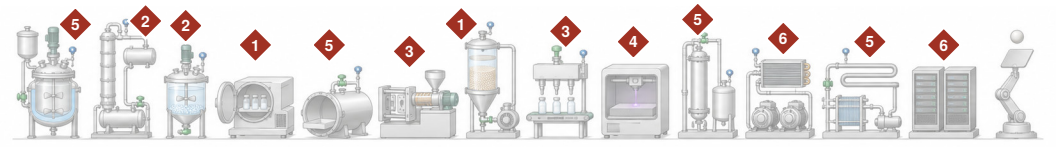
The page numbers in this file are those of the printed volume.

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v08-batch-hybrid/

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

CHAPTER 16

Summary and Outlook



Open gaps 1 end quality 2 trajectories 3 sparse results 4 learning 5 supervisors 6 coupling

Notice that the gap of verified supervisors joins four plants, the reactor, the autoclave, the capture train and the pasteurizer, more than any other gap. Illustration, not to scale.

IN PLAIN TERMS **What Fourteen Batches, Runs and Switches Taught**

The main idea. Each chapter of this book ran a plant that starts and stops: a batch judged once at its end, a run repeated many times, or a machine that jumps between modes. Nothing settled at a steady value that a controller could hold.

The challenge. What decided the result was the path the batch followed, the reading that ended each step and the simple guards around every switch. A controller that held its target well still made a poor batch when it was told to follow the wrong path.

Why it matters. Judging the end result on what was really made tells a good batch from one that only looked good on the screen.

Fourteen chapters put batches, repeated runs and switched plants in front of a controller. In each one the result was judged on something no loop held at a setpoint: a dose finished without a runaway, a drum at its purity, a vial dry and intact, a part of the right mass, a bounce that came back to its apex. Table 1.3 of Chapter 1 listed what each chapter judges at the end, the strategies it compares and what keeps its switching safe, and left the results to the chapters. This chapter adds the results, draws a short guide for choosing a strategy, looks at where practice is heading, reports how the volume itself was made with AI agents and what they got wrong, and lists the problems that remain open.

16.1 What the Chapters Showed

Tables 16.1 and 16.2 compare the chapters. The first holds the plants judged at the end of one batch, where the strategy is a trajectory, a guard that ends a phase or a correction between shots. The second holds the plants that repeat or switch, where the strategy is a correction from one run to the next, a learned profile or a supervisor over a switched plant. Each row names what is judged, a strategy that worked with the device that kept its switching safe, one that failed, and the measured result with the table or section that reports it. Figure 16.1 sets one like pair of numbers from each row side by side.

The columns of the two tables make five points.

The path decided the product, not the tracking. On the crystallizer the best-tracked curve, 0.142 °C at its worst, made 96.3 % fines, and the cubic curve with a larger deviation made 13.8 %. On the batch column the open-loop ramp collected 15.73 mol against the composition loop's 15.29 mol without any feedback. The semi-batch reactor held its temperature within 0.2 K while it stored enough reagent to rupture its disc. In each case the loop did what it was asked. The question that decided the batch was what it was asked to follow (Section 1.4.3).

Table 16.1 Batches Judged at the End, Chapters 2 to 8

Chapter	Judged at the end	What worked	What failed	Measured result
2, semi-batch reactor	The whole dose fed, with the MTSR below its limit if cooling fails.	Dosing limited on the calorimeter's MTSR. A feed cut with hysteresis and a latched interlock.	A fixed dose at 30 °C, held within 0.2 K while 4398 mol accumulate.	Circulation lost at 2:30, the fixed dose is lost at 3:55 and the limited dose peaks at 74.6 °C. Challenge 22.76 against 1.61 K (Table 2.5).
3, batch column	Product at a cumulative purity of 0.990 in 180 min.	An open-loop reflux ramp searched on the model. Tracking that opens the composition loop without a bump.	The same ramp on a charge of 0.42. Constant composition behind an 8 min analyzer.	15.10, 15.29 and 15.73 mol for constant reflux, composition and ramp (Table 3.5). Wrong charge 1.26, trimmed 4.74, re-optimized 10.69 mol (Table 3.7).
4, crystallizer	Mean size and fines of the filtered crystals.	A cubic cooling curve, slow first. A cascade that states the wall limit.	Curves that tracked better. A cubic batch shortened to 150 min.	Fines 96.3, 99.3 and 13.8 % at peak deviations 0.142, 0.637 and 0.237 °C (Table 4.4). At 150 min, 68.2 % (Table 4.7).
5, freeze-dryer	Every vial group dry, none collapsed.	An endpoint-aware shelf with a learned resistance and the Pirani endpoint.	The shipped recipe and a probe-limited shelf. The learned resistance with a gauge factor of 1.7, not 1.6.	Passes 0, 0 and 6 of 6 seeds. Stoppered at 24.74 against 26.12 h for a tuned recipe and 24.47 h reading the truth (Table 5.6). Every group collapsed at 1.7 (Table 5.8).
6, autoclave	Cure of the slowest part, exotherm of the hottest core.	Lead/lag selection, the max loop through a low selector, and the observer's exotherm feedforward.	The air's clock alone. A thick spar ramped at 3 °C/min without the feedforward.	Spar peak 243.9 against 184.8 °C with the feedforward, at the same cycle (Table 6.6). Challenge passed in 289.8 min (Table 6.7).
7, molding machine	Mass of every part.	Switchover by tracking. A correction from the scale with 40 MPa of authority.	The fixed recipe. Switchover at 50.5 % fill (short) or 103.6 % (flash).	Parts in tolerance 9, 26 and 30 of 30 (Table 7.6). Handover step at most 0.44 % in 20 shots (Table 7.5).
8, particle coater	Four layer thicknesses and the carbide spread.	An inferential endpoint corrected by a load cell. Ramped gas changes and a timeout.	A timer on a cold carbide step. The estimate alone at a coating efficiency of 0.72.	Carbide 26.2 by timer, 35.0 µm by estimate (Table 8.4). Gap 28.35 closed to −0.79 µm for 191 to 219 min (Table 8.6).

An endpoint is only as good as the estimate behind it. A clock cannot see a cold carbide step, and the coater's estimate cannot see a coating efficiency outside its model, which left the batch 28.35 µm thin with every setpoint held. A load cell closed that gap and was itself blind to a wrong charge ticket. The freeze-dryer's learned resistance held its margin over a wide range and failed outright on a gauge factor 6 % high. The guard that worked in each case read a second, independent measurement, as Table 1.1 anticipated.

A correction between runs pays only for what persists. The molding correction took the campaign from 9 to 30 parts in tolerance, and at its default gain it widened the shot-to-shot spread by 15 %, close to the 17 % its loop gain predicts (Section 7.5). On the vial filler, with 5.9 weighed vials of a head between a fill and its result, the plain per-head law spread the heads to 47 to 48 mg, where

Table 16.2 Runs That Repeat and Plants That Switch, Chapters 9 to 15

Chapter	Judged at the end	What worked	What failed	Measured result
9, vial filler	Fill of every vial, with results in flight.	A record-based estimate per head, or the plain law detuned. A validity window and a step limit.	Pooled and per-head corrections at a gain of 0.50.	Accepted 566 and 519 against 2325 (record-based) and 2324 (detuned), with 2342 for ideal information (Table 9.6).
10, powder bed fusion	Melt pool width along every layer.	Learning from the last layer with a Q-filter, feedback underneath, the profile kept at a new outline.	More feedback gain. Learning without the filter.	Defects 4.7 fixed, about 2.5 with feedback, 0 from layer 3 with learning. No filter, 24.8 by layer 25. Kept 13.2 against 23.2 reset, in percent (Table 10.4).
11, antibody capture	Antibody per liter of resin and hour, antibody lost.	A UV threshold or a learned capacity. Switch only when the next column is ready.	A 25-minute clock. Loading on through a long cleaning hold.	38.32 and 37.98 against 37.23 g/L/h, recovery 98.57 against 90.42 % (Table 11.6). Loss 9.05 against 0.05 % paused (Table 11.5).
12, refrigeration pack	Goods temperature, energy and compressor starts.	A scheduler that desynchronizes the cases, floating pressures, a 6 K superheat. Band, minimum run and rest.	Free-running thermostats, in step after 65.2 min. A 2 K superheat.	Summer day 9 trips and 383 starts against 0 and 108 (Table 12.7). Three hours 45 starts free against 8 scheduled (Table 12.8).
13, pasteurizer	Hold time and temperature of every parcel.	A cascade with feedforward under a diversion supervisor, return after a flush, conditional integration.	No flush before the return. Two thermometers drifting together.	Challenge 4 diversions and 1554 L against none (Table 13.6). No flush, 8 diversions (Table 13.5). 128 parcels pass unseen (Table 13.4).
14, autoscaled service	Latency percentile, dropped requests and cost.	Admission control, a retry budget, a forecast floor, a cooldown.	The platform's rule in a fourfold crowd. Retries with the fleet capped.	Crowd 14 breach minutes and 23.7 % dropped, 0 and 4.8 % admitted on requests (Table 14.4). Storm 129 to 0 minutes (Table 14.5).
15, juggling paddle	A period-one bounce at its apex.	An open-loop stroke whose paddle slows at impact, 6.3 cm at 1.655 Hz.	The shipped stroke of 4.1 cm. A mirror law behind a camera 120 ms late.	Challenge passed on 0 and 10 of 10 seeds, 1.6 cm rms (Table 15.6). Every gain dies at 120 ms (Table 15.5).

the record-based estimate held 4.6 to 6.8 mg (Table 9.5). Learning on the powder bed removed defects that repeated at every corner and diverged without its filter. Section 1.4.5 gives each result its formula: the noise widening $2/(2-\lambda)$, the delay polynomial and the monotone convergence factor.

Switching was kept safe by devices outside the controller. Every chapter that switched relied on a device that had no model in it: a feed cut with hysteresis and a latch, tracking at a switchover, a flush before the return to forward flow, a minimum run and rest, a cooldown, a retry budget, a column that is ready before the switch. Removing one was measurable. The pasteurizer without its flush diverted 8 times instead of once. The autoclave's predictive controller met its limits only after the selector's protections were kept above it (Section 6.7.7).

A stable cycle can be the goal or the fault. The juggler's stroke makes the period-one bounce an attracting cycle with no sensor at all, and the measured eigenvalues, 0.79 to 0.83, match the closed form's 0.80 (Section 15.4). The refrigeration thermostats find a stable cycle too, all four in step

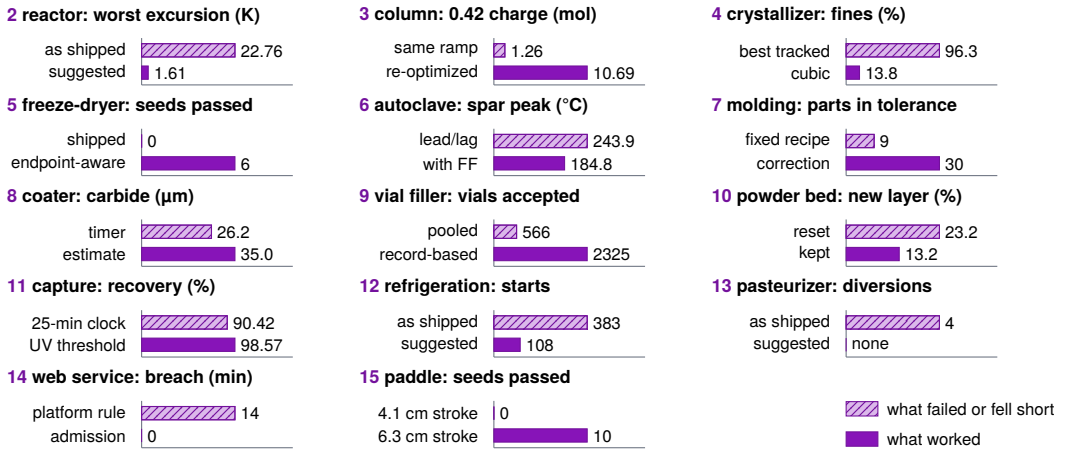


Figure 16.1 One like pair of numbers from each row of Tables 16.1 and 16.2, digit for digit, each panel on its own scale. The hatched bar is the policy or setting that failed or fell short, the solid bar the one that worked. Higher is better for the column, the freeze-dryer, the molding machine, the coater, the vial filler, the capture train and the paddle, lower for the other seven. The pasteurizer's "none" is drawn at zero. Notice that every better number came with a change of path, endpoint, correction or guard, and none with a tighter loop alone.

through one suction pressure, and that cycle is the fault. Section 1.4.9 reads both the same way, as the fixed point of a cycle-to-cycle map.

16.2 Choosing a Method

The chapters suggest an order of questions for a new batch, cycle or switched plant. Each names the chapter where its answer decided the result.

1. *What is judged at the end, and what during the run measures it?* Write the end result as a function of the path, as the dose of Section 1.4.1 is, before choosing a loop. If nothing in the run measures it, plan an estimate and a check on it (Chapter 8).
2. *Which variable carries the hazard?* Limit that variable, not the one the trend shows. The reactor's hazard was the unreacted reagent, and the temperature showed none of it (Chapter 2).
3. *Is the trajectory known well enough to run open loop?* A ramp searched on a correct model beat feedback, and on a wrong charge it collapsed to 1.26 mol. Re-optimize when the charge changes and trim only small errors (Chapter 3).
4. *What ends each phase?* A clock is blind to a slow batch, an event to a missing trigger and an estimate to its own model. Put a time limit behind each trigger (Chapter 7, Chapter 8).
5. *Does the disturbance repeat within a run, between runs, or not at all?* Learn what repeats along the run (Chapter 10), correct between runs what persists (Chapter 7), and leave the rest to feedback, which a correction can only make worse.
6. *How many results are in flight?* Count the runs between an action and its measurement, and either estimate from the record of each run or detune the gain (Chapter 9).
7. *Where does the plant switch, and what stops it chattering or bumping?* Choose a band, a dwell, a latch, tracking or a persistence time for each switch (Section 1.4.7), and check that a shared utility does not pull independent switches into step (Chapter 12).

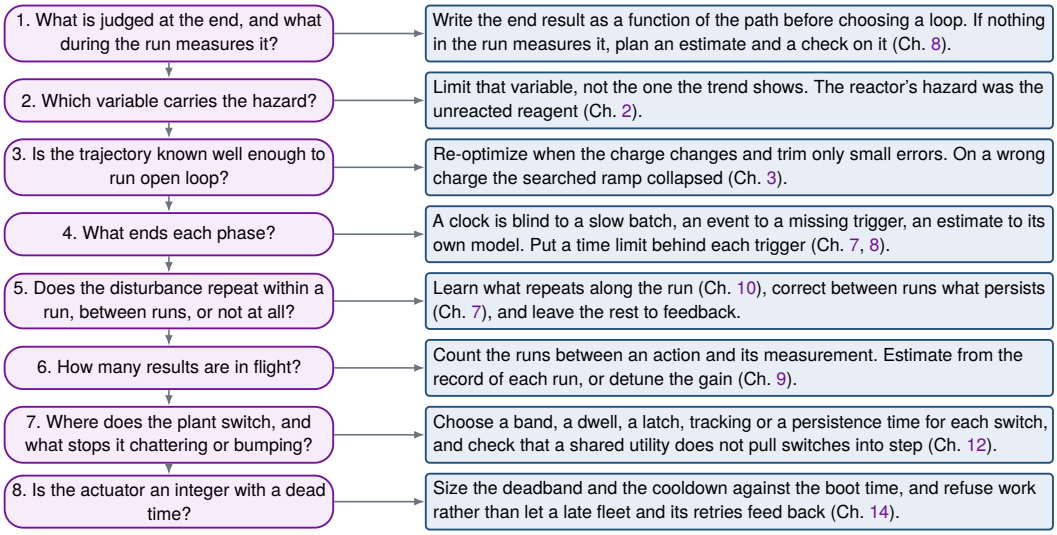


Figure 16.2 The eight questions of this section in order, with the list's answer and the chapters where it decided the result. Notice that what is judged and what carries the hazard are settled before any loop is chosen, and that the last two questions are about switches rather than loops.

8. *Is the actuator an integer with a dead time?* Size the deadband and the cooldown against the boot time, and refuse work rather than let a late fleet and its retries feed back (Chapter 14).

A predictive controller earned its place where it planned around a limit the rules met by cutting. On warm cooling water it dosed the reactor with no cut where the one-step optimizer cut 21 times, for a batch 5.8 min longer (Table 2.6). On four freeze-dryer shelves it never lost the chamber pressure, which the endpoint-aware policy lost for 36.7 min, and moved the shelf 31 K against 482 K (Table 5.9). Elsewhere it traded one measure for another or lost. The crystallizer's batch fell from 187.0 to 126.9 min while fines rose from 13.8 to 15.6 % (Table 4.9). On the batch column it made 14.2914 mol against 14.3393 mol for a ramp searched on the run's model, with a reflux travel of 173.47 against 5.81 (Table 3.8). On the refrigeration pack it used 21.5184 kWh in three hours against the scheduler's 21.5574 kWh, with 24 rack starts against 8 (Table 12.8), and on the web service it billed $11.93 against the PI loop's $11.06 on the morning ramp (Table 14.7). Volume 9 treats that choice in full. Figure 16.2 draws the eight questions in order.

16.3 Where Practice Is Heading

Batch optimization moves to measurements. A review of batch control names the main difficulty as run-end outputs, such as final quality, that cannot be measured during the run, with process models that tend to be poor. It places the research push in measurement-based optimization that uses data from the current and previous batches [1]. The batch column and the coater of this volume show both halves of that statement. A ramp searched on the model was the best policy until the charge changed, and an estimate of thickness was right until the efficiency moved outside it. In the editor's view the recipe written once and run unchanged will give way to recipes that are re-planned from each batch's own measurements, with the safety layer left as it is.

Corrections between runs spread beyond the wafer fab. Run-to-run control was developed for semiconductor processes, combining response surfaces, statistical process control and feedback, and one review proposed a control chart on the output that acts as a deadband [2]. Run-to-run control for

batch processes has its own survey [3], and learning control, repetitive control and run-to-run control have been surveyed together [4]. A survey of iterative learning control treats stability, performance, learning transients and robustness, and describes a field still active after two decades [5]. In additive manufacturing, learning schemes now use data from previous layers and previously printed parts and are designed for plants that vary from one iteration to the next [6]. A review of metal additive manufacturing found that factory-ready machines had not yet incorporated monitoring that detects discontinuities during the build [7]. The powder bed of Chapter 10 shows why both matter: learning needs a measurement of every layer, and it fails where the repetition breaks.

Some batches become continuous. An article on modernizing pharmaceutical manufacturing states that continuous processing can address agility, flexibility, cost and robustness, supported by quality by design and process analytical technology, and that the US Food and Drug Administration supports its implementation with science- and risk-based approaches [8]. The antibody capture of Chapter 11 is already a cyclic schedule fed without a stop. In the editor's view the freeze-dryer and the vial filler will remain batch and lot steps for longer, and the work there is the estimate of what no probe measures and the record that ties each result to its vial.

Switching is designed, not tuned. A tutorial on hybrid dynamical systems sets out models, stability and control for plants that combine flows and jumps, with emphasis on robustness to perturbations and measurement error [9]. Synchronized display cases have been modeled as a hybrid system whose synchronized state is a periodic trajectory, and the analysis found that more measurement uncertainty makes synchronization less prevalent [10]. Predictive hybrid control has been applied to a supermarket refrigeration benchmark [11], and the autoscaling of elastic cloud services has a review of its own techniques [12]. In the editor's view the devices of this volume (band, dwell, latch, tracking, flush, budget) will stay, because they are simple enough to verify, and optimization will be placed above them, as the autoclave's predictive controller had to be.

Instruments that reach the product. A white paper on monitoring practice in pharmaceutical freeze-drying finds thermal-conductivity gauges and capacitance manometers both commonly used, though often not in the most useful way. It names wireless product temperature, laser absorption spectroscopy in the vapor duct and heat flux measurement as tools that may become best practice. Wireless sensors had appeared in the few years before it, avoid the lead wires that make probing awkward in manufacturing, and serve at laboratory and production scale [13]. A review of crystallization control finds most in situ image analysis quantitative only at low solid concentrations, typically at most 5 to 7 percent, and expects research on systems that work at higher ones [14]. The freeze-dryer of Chapter 5 reads two probe vials and two gauges, and the crystallizer of Chapter 4 reads an analyzer of the liquid. In the editor's view both would gain most from a measurement of the product itself.

Feedforward from the part's own geometry. Adjusting the laser power once per layer is a common strategy in powder bed fusion, and some commercial machines close photodiode feedback within a layer. A study of 2025 notes that such feedback stays reactive, cannot anticipate known geometry such as overhangs, and complicates certifying parts for critical applications when it changes parameters in real time. Its feedforward on each scan vector, built on a light finite-difference thermal model and a reduced melt pool model and tested on a complex part, cut geometric inaccuracy in key dimensions by 62 percent and porosity by 16.5 percent on average [15]. The learning of Chapter 10 reaches the corners by another road. It learns them from the last layer, where such a model predicts them before the first.

Clients that limit their own retries. A handbook of site reliability engineering describes a retry budget of up to three attempts per request and a budget per client that allows a retry only while

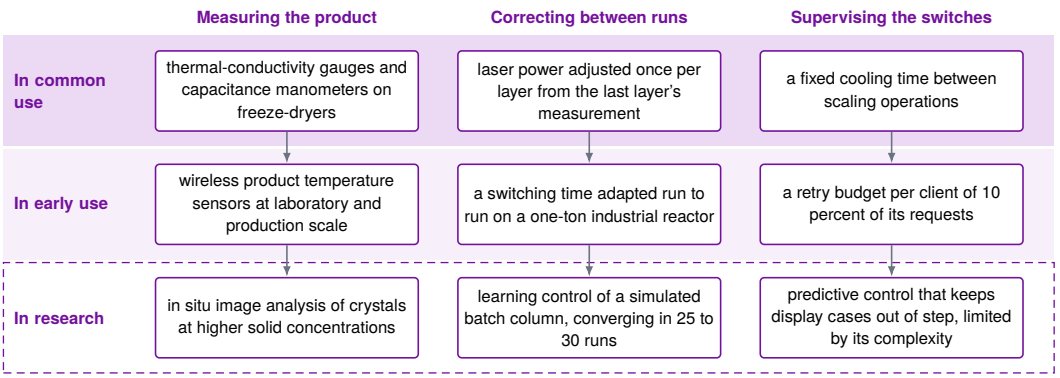


Figure 16.3 Developments in three threads. The top lane holds what the cited source calls commonly used, a common strategy or a common solution adopted in practice, the middle lane what it reports in service at production scale or at one industrial plant, and the bottom lane what it reports from simulation or names as a need, in the editor's reading. Sources by thread: measuring the product [13, 14], correcting between runs [15, 18], supervising the switches [16, 17, 19]. No date is drawn, because the sources give none for these steps.

retries are under 10 percent of its requests. In its worst case, a datacenter that rejects most requests, the limit per request alone lets the load grow to just below three times, and the budget per client holds the growth to about 1.1 times [16]. A survey of autoscalers finds a fixed cooling time between scaling operations a common solution adopted in practice, set to at least the time to acquire and boot a machine [17]. The web service of Chapter 14 carries both devices. Its retry budget of 10 % left no storm minute on any seed, and its cooldown of 300 s reversed no action on the evening shoulder. Figure 16.3 sets these developments and those above in three threads.

> ⧖ HISTORY **2017: Run-to-Run Optimization on a One-Ton Reactor**
>
> The history boxes of Chapter 1 end in 1995, with run by run control in chip making. A chapter of 2017 reported the idea on a one-ton industrial reactor. Each batch ran isothermal and then adiabatic, and the time of the switch was adapted from one batch to the next. The reaction time fell significantly after only two batches, without off-specification product, and ended 35 percent shorter than the isothermal policy of industrial practice. A back-off from the final temperature limit was kept for safety, and the policy became standard practice for the industrial partner [18]. A review of 2018 found the quality loop of injection molding the least developed of its three loops, its learning and predictive solutions good but not widely used [20]. In the editor's reading the chapters of this volume stand between those two reports.

16.4 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 [21]. 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 [22]. 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.

16.4.1 How This Volume Was Made

Each of the fourteen 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. At least nine of the fourteen reports are self-reviews by the agent that wrote the paper. Six apps of the later rounds passed usability, design and technical-accuracy reviews before their papers. 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. The fourteen logs written beside the papers hold 104 entries with a severity. Of these, 32 are wrong behavior or output of an app, 49 are documentation, 13 are cosmetic and 10 are other kinds, such as a test allowance or a scan period. The other eight apps account for 91 entries. The six that had a technical review before their papers account for 13. Their reviews before the paper had already logged and fixed findings in each app's own record. Five cases show the errors that matter for this volume's topic.

A batch graded off specification when it was run well. The batch column's scorecard compared the receiver's true average with the specification at a tolerance of 10^{-6}. The tail cut closed the drum on the analyzer's running average, which lags the column, so a well-run batch landed 3.2×10^{-4} below 0.990 and the policy the app ships as the winner, 15.735 mol at 0.98968, was graded a failure. The app's own tests already used a tolerance of 10^{-3}. The paper agent found it while writing the paper, and one shared test now grades all five places that judge the drum.

Switches that were not bumpless. On the same column the composition loop kept integrating during the total-reflux startup, when the plant could not answer it, and collection opened at a reflux ratio of zero. The distillate fell to 0.9462. With the loop tracking the ratio it will open on, the dip is to 0.9896. On the crystallizer, supersaturation control entered with a proportional kick of 5.5 °C on its setpoint, and the result depended on how the learner reached it. With 6 g of seed two routes finished at 285 and 293 min. Treated as a transfer, the batch moves the valve 66 % instead of 153 %. The crystallizer's log records that its entries were found by reading the source line by line against the numbers the app quotes.

A headline that hid where it came from. The freeze-dryer's endpoint-aware policy stoppers 1.38 h before a recipe tuned on the same batch. The referee asked where the time came from, and 1.07 h of it lies in secondary drying, the weaker half of the model (Section 5.6). The paper and the app now say so. The vial filler's first draft claimed that the record-based estimate has an advantage when the gain must be high. The lot-step test gave 2093 against 2094 accepted, and the claim was rewritten. The molding paper stated the variance law $2\sigma^2/(2-k)$ before it was measured, and a 195-part run was added to test it.

A predictive controller less safe than the selector. The first autoclave MPC kept only the max loop above its plan and set its ceilings without a margin for the reaction heat. A technical review ran 46 worst cases and found it worse than the lead/lag selector with feedforward on the selector's limit in 12 of them, with one thermocouple 3.15 °C over for 4.8 min. The cause was the model's heat of reaction. The fix was chosen by a rule written before any candidate ran and amended once after the first results. It keeps the feedforward and the lead clamp above the plan, adds a reaction margin to the ceilings and moves the optimized recipe from 281.0 to 291.3 min (Table 6.9). A headline of 281.0 min would otherwise have credited the MPC with minutes it earned by planning to the limit on its own model.

Controllers that read the plant's truth. In three apps a controller took values it could not measure on a real plant, and a fourth log documents one more. The reactor's calorimeter, its heat transfer law, the MTSR limiter and the optimizer took fifteen constants from the run's own parameters, so a weaker jacket entered in Parameters looked safe. The limiter dosed in 354.0 min with the true MTSR at 80.51 °C, against 222.1 min and 104.09 °C when the same jacket came from the disturbance slider. The autoclave's cure observer knew a learner's changed resin, and with it a thick part peaked at 202.6 °C instead of the 215.2 °C it reaches on the declared values. The juggler's impact switch reported where the ball truly hit, up to 4.9 mm from the camera's estimate. The capture train's surge hold reads the true tank volume, documented and left as a plant-side interlock. The reactor and autoclave fixes moved no shipped run, and the juggler's moved two cases. The reactor and autoclave entries record the probe behind the shared section's advice to read code for what it reads. A hidden value is changed, and the question is whether the controller moves while the measurements do not.

The book round. The fourteen chapter records list 207 rows of numeric verification, each a stated value or group of values set beside its reproduction. No reported number failed to reproduce, apart from one sweep. On the coater three last digits of a spread sweep differ when the setpoint is set before

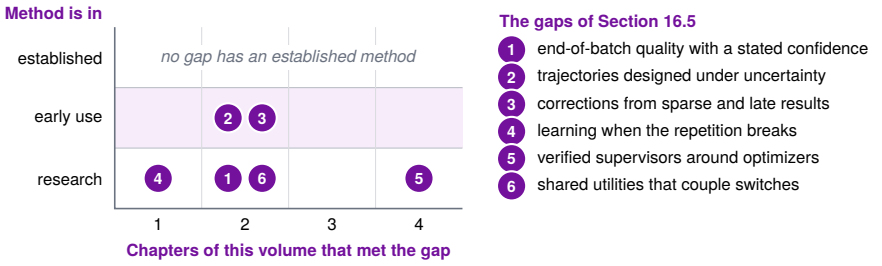


Figure 16.4 The six gaps of this section. Across is the number of chapters that each item names as evidence: two for gap 1 (Chapters 8 and 5), two for gap 2 (3 and 4), two for gap 3 (9 and 7), one for gap 4 (10), four for gap 5 (6, 13, 2 and 11) and two for gap 6 (12 and 14). Up is the maturity of the method available, the editor's reading of the sources cited in each item and in the Open challenge box. Notice that the gap met by the most chapters, the verification of the logic around an optimizer, has the least mature method.

Run rather than during the opening purge, a path dependence the exercise now states as a range. The book round's defect list holds 13 entries for this volume, five of them check counts that grew after the papers. The molding paper called both loops direct acting, against the convention of Section 1.6.2, and the autoclave paper used "maximum selector" loosely. Both are corrected in the chapters and have since been fixed in the papers. The agent writing Chapter 1 first stated that a deadband around an integer actuator turns the PI loop into a thermostat with a band of $\pm 7.5\,^{\circ}\mathrm{C}$. The simulation gave $\pm 2.6\,^{\circ}\mathrm{C}$, because the integral moves too, and the formula was dropped.

Every error above was caught by a check that could fail: a second measurement, a rerun along another path, a decomposition of a headline, a sweep of worst cases, a hidden value moved to see what else moved. Most sat at a boundary that this volume is about, an endpoint, a switch or a correction from a record. Four working rules follow. Judge the end result on the truth and the estimate separately. Test every switch along every route into it. Decompose every comparative claim before it is printed. Keep the protections above any optimizer, and let a person decide what changes.

16.5 Gaps in Batch, Run-to-Run and Hybrid Control

Six gaps belong to this volume's topic, and the gaps common to agent-era practice follow them in Section 16.5.1. In the editor's view none is closed. Each item gives the chapters' evidence, then the state of the literature with its sources. Why each gap is hard and what progress would look like are the editor's view. Figure 16.4 places the gaps by the chapters that met them and by the maturity of the method available, and the same numbers mark the plants in the picture that opens this chapter.

1. *End-of-batch quality with a stated confidence.* The coater's load cell closed one gap and could not see a wrong charge, and the freeze-dryer's learned resistance failed on a gauge error of 6 %. A chapter on batch control warns that estimates of infrequently measured product properties can never completely replace the measurements themselves [18]. A retrospective of a coated particle fuel program asks for measurement techniques deployable in an industrial setting, since its buffer density is measured by interrupting the coating and gives only an average [23]. The gap is hard because the error that matters comes from a disturbance no instrument in the batch sees. An estimator that states, during the batch, which disturbance it cannot see and how large an error it would miss is not available.

2. *Trajectories designed under uncertainty.* The column's ramp collapsed on a wrong charge, and the crystallizer's predictive batch bought an hour with more fines. Optimization on a model computes inputs optimal for the model, often fails to reach plant optimality, and in practice gives suboptimal and often infeasible operation. In crystallization and polymerization an off-specification batch

often cannot be corrected at all [18]. The gap is hard because each batch is one draw from the spread and cannot be repeated to average it out. A routine method that searches a recipe over the spread of charges and kinetics, and reports what the spread costs, remains an open engineering problem.

3. *Run-to-run correction with sparse and delayed results.* With one group of vials in four weighed, no correction passed the challenge, not even ideal information. The molding correction widened the spread of what did not persist by 15 %. A review of injection molding finds the loop that feeds part quality back the least developed of three, and its learning and predictive solutions good but not widely used for lack of standardization [20]. The gap is hard because the weighing plan, the delay and the gain are set by different people for different reasons. Design rules that tie them to the fraction of product released untested are thin.

4. *Learning when the repetition breaks.* A learned profile kept at a new outline left 13.2 % of the first layer defective, and one reset left 23.2 %. A study of powder bed control finds feedback reactive to geometry it could know in advance, and real-time changes of parameters a complication for certifying parts [15]. The gap is hard because what was learned belongs to the shape it was learned on. Methods to transfer it to a new shape, a new machine or a new powder lot, with a guarantee during the first layers, are needed.

5. *Verified supervisors around optimizers.* The autoclave's first predictive controller was less safe than the selector in 12 of 46 worst cases, and the pasteurizer's flush, the reactor's latch and the capture train's ready signal were all checked by simulation alone. The Open challenge box at the end of this section treats this gap.

6. *Shared utilities that couple switches.* Four thermostats locked in step through one suction pressure, and a capped web fleet held its clients in a retry storm for two hours. A thesis that kept two display cases apart by predictive control found the method limited by its growing complexity and computation time on larger systems [19]. A workshop paper on metastable failures, in which a bad state persists after its trigger is removed, calls a systematic approach to robustness against unknown ones an open problem [24]. The gap is hard because the coupling appears only when many units act together. Design rules that predict it from the plant's structure, before it appears in operation, are needed for plants with many switched units on one header or service.

OPEN CHALLENGE Verifying the Logic Around an Optimizer

Gap 5 runs through four chapters, each with protections checked by simulation alone. The run-to-run reactor policy kept a back-off from its temperature limit for disturbances during the run that the method cannot handle [18]. A juggling robot was given an open-loop stroke because no verification process for complex closed-loop network policies was known to exist [25]. Hybrid systems theory states robust stability for plants that flow and jump [9]. In the editor's view the logic is hard to verify because interlocks, switches and an optimizer act on one plant at once. Progress would be a check of the combined logic, run before commissioning, that fails when a protection is removed.

16.5.1 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 [26]. 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 [22]. 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 [27]. 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 [28]. 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 [29]. 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

- The path a batch followed decided its product, not how closely the loop tracked that path.
- An endpoint is only as good as the estimate behind it. The guard that worked read a second, independent measurement.
- A correction between runs pays only for what persists, and learning fails where the repetition breaks.
- Simple devices outside the controller kept every switch safe, and an optimizer earned its place only with the protections kept above it.

References

[1] D. Bonvin, B. Srinivasan, and D. Hunkeler, "Control and optimization of batch processes," *IEEE Control Systems Magazine*, vol. 26, no. 6, pp. 34–45, 2006.

[2] E. Del Castillo and A. M. Hurwitz, "Run-to-run process control: Literature review and extensions," *Journal of Quality Technology*, vol. 29, no. 2, pp. 184–196, 1997.

[3] K. Liu, Y. Chen, T. Zhang, S. Tian, and X. Zhang, "A survey of run-to-run control for batch processes," *ISA Transactions*, vol. 83, pp. 107–125, 2018.

[4] Y. Wang, F. Gao, and F. J. Doyle, "Survey on iterative learning control, repetitive control, and run-to-run control," *Journal of Process Control*, vol. 19, no. 10, pp. 1589–1600, 2009.

[5] D. A. Bristow, M. Tharayil, and A. G. Alleyne, "A survey of iterative learning control," *IEEE Control Systems Magazine*, vol. 26, no. 3, pp. 96–114, 2006.

[6] Z. Afkhami, D. J. Hoelzle, and K. Barton, "Robust higher-order spatial iterative learning control for additive manufacturing systems," *IEEE Transactions on Control Systems Technology*, vol. 31, no. 4, pp. 1692–1707, 2023.

[7] S. K. Everton, M. Hirsch, P. Stravroulakis, R. K. Leach, and A. T. Clare, "Review of in-situ process monitoring and in-situ metrology for metal additive manufacturing," *Materials & Design*, vol. 95, pp. 431–445, 2016.

[8] S. L. Lee, T. F. O'Connor, X. Yang, C. N. Cruz, S. Chatterjee, R. D. Madurawe, C. M. V. Moore, L. X. Yu, and J. Woodcock, "Modernizing pharmaceutical manufacturing: From batch to continuous production," *Journal of Pharmaceutical Innovation*, vol. 10, no. 3, pp. 191–199, 2015.

[9] R. Goebel, R. G. Sanfelice, and A. R. Teel, "Hybrid dynamical systems," *IEEE Control Systems Magazine*, vol. 29, no. 2, pp. 28–93, 2009.

[10] J. Leth, R. Wisniewski, J. G. Rasmussen, and H. Schiøler, "Stochastic analysis of synchronization in a supermarket refrigeration system," *Electronic Proceedings in Theoretical Computer Science*, vol. 174, pp. 40–50, 2015.

[11] N. L. Ricker, "Predictive hybrid control of the supermarket refrigeration benchmark process," *Control Engineering Practice*, vol. 18, no. 6, pp. 608–617, 2010.

[12] T. Lorido-Botran, J. Miguel-Alonso, and J. A. Lozano, "A review of auto-scaling techniques for elastic applications in cloud environments," *Journal of Grid Computing*, vol. 12, no. 4, pp. 559–592, 2014.

[13] S. Nail, S. Tchessalov, E. Shalaev, A. Ganguly, E. Renzi, F. Dimarco, L. Wegiel, S. Ferris, W. Kessler, M. Pikal, G. Sacha, A. Alexeenko, T. N. Thompson, C. Reiter, J. Searles, and P. Coiteux, "Recommended best practices for process monitoring instrumentation in pharmaceutical freeze drying—2017," *AAPS PharmSciTech*, vol. 18, no. 7, pp. 2379–2393, 2017.

[14] Z. K. Nagy and R. D. Braatz, "Advances and new directions in crystallization control," *Annual Review of Chemical and Biomolecular Engineering*, vol. 3, no. 1, pp. 55–75, 2012.

[15] N. Kirschbaum, N. Wood, C.-E. Kim, T. U. Tumkur, and C. Okwudire, "Vector-level feedforward control of LPBF melt pool area using a physics-based thermal model," *Additive Manufacturing*, vol. 112, p. 104981, 2025.

[16] B. Beyer, C. Jones, J. Petoff, and N. R. Murphy, Eds., *Site Reliability Engineering*. O'Reilly Media, 2016.

[17] C. Qu, R. N. Calheiros, and R. Buyya, "Auto-scaling web applications in clouds: a taxonomy and survey," *ACM Computing Surveys*, vol. 51, no. 4, pp. 1–33, 2018.

[18] D. Bonvin and G. François, "Control and optimization of batch chemical processes," in *Coulson and Richardson's Chemical Engineering*. Elsevier, 2017, pp. 441–503.

[19] L. F. S. Larsen, "Model based control of refrigeration systems," Ph.D. dissertation, Department of Control Engineering, Aalborg University, Aalborg, Denmark, 2006, iSBN 87-90664-29-9. [Online]. Available: https://vbn.aau.dk/ws/files/549479206/thesis.pdf

[20] O. Ogorodnyk and K. Martinsen, "Monitoring and control for thermoplastics injection molding: A review," *Procedia CIRP*, vol. 67, pp. 380–385, 2018.

[21] 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.

[22] 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.

[23] P. A. Demkowicz, D. A. Petti, J. D. Stempien, J. Palmer, W. F. Skerjanc, D. M. Scates, G. Hawkes, T. Gerczak, and G. Helmreich, "Advanced gas reactor fuel development and qualification program retrospective and lessons learned," Idaho National Laboratory, Idaho Falls, ID, Report INL/RTP-25-06381, Revision 0, 2026, oSTI 17429367.

[24] N. Bronson, A. Aghayev, A. Charapko, and T. Zhu, "Metastable failures in distributed systems," in *Proceedings of the Workshop on Hot Topics in Operating Systems (HotOS '21)*, 2021, pp. 221–227.

[25] K. Ploeger, M. Lutter, and J. Peters, "High acceleration reinforcement learning for real-world juggling with binary rewards," in *4th Conference on Robot Learning (CoRL 2020)*, 2020, arXiv:2010.13483.

[26] 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.

[27] 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.

[28] 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.

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

16.6 Exercises

16.1 Discussion · 8 min

On the batch column of Chapter 3 a ramp searched on the model collected 15.73 mol, constant-composition control 15.29 mol and the best constant reflux 15.10 mol. An agent proposes to ship the open-loop ramp on every charge because it won. Argue for and against, and say what you would require before the ramp runs without feedback.

16.2 Discussion · 8 min

An agent tuning the correction between runs on a new filling line proposes a gain of 0.50 because "a faster correction removes drift sooner." Using Chapter 7 and Chapter 9, explain when a higher run-to-run gain helps and when it hurts, and list what you would ask the agent to measure first.

16.3 Discussion · 7 min

The coater of Chapter 8 ended its layers on a timer, on an integrated estimate and on the estimate corrected by a load cell. For each guard, name a disturbance it handles and one it misses, using the chapter's results, and propose a guard for a new batch that has no load cell.

16.4 Discussion · 8 min

An agent replaces the autoclave's lead/lag selector with a predictive controller and reports the optimized recipe 8.8 min faster. Using Section 16.4.1 and Chapter 6, describe the review you would run before accepting it, and say where the protections should sit relative to the optimizer.

16.5 Discussion · 6 min

The freeze-dryer's endpoint-aware policy stoppered 1.38 h before a recipe tuned on the same batch, and 1.07 h of that came from secondary drying. Explain why such a decomposition belongs in any comparison an agent writes, and propose two other checks of a comparative claim, using examples from Section 16.4.1.

16.6 Discussion · 7 min

In the reactor of Chapter 2, a weaker jacket entered in Parameters let the MTSR limiter dose in 354.0 min with the true MTSR at 80.51 °C, while the same jacket from the disturbance slider gave 222.1 min and 104.09 °C. Explain what this difference reveals, why no shipped run showed it, and design a test that would catch the same defect in another batch app.