INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 3 Feedforward, Cascade, and Ratio Control

CHAPTER 13

Summary and Outlook

Pages 313–325

Edited by John D. Hedengren

How to cite this chapter. "Summary and Outlook," in *Feedforward, Cascade, and Ratio Control*, John D. Hedengren, Ed., vol. 3 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 13, pp. 313–325.

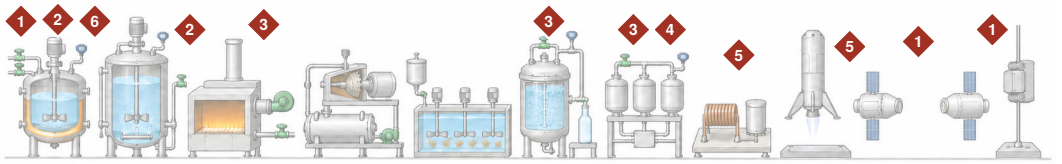
The page numbers in this file are those of the printed volume.

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v03-feedforward-cascade/

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

CHAPTER 13

Summary and Outlook



Open gaps 1 feedforward model range 2 inner-loop speed rules 3 trims and instruments 4 streams in a blend 5 undershoot, authority 6 structure from data

Notice that the tank with two upsets carries three gaps and that the trim gap runs from the heater to the formulation skid, while the vehicles met the gaps of model range and actuator authority. Illustration, not to scale.

IN PLAIN TERMS **What Eleven Plants Taught About Acting Early**

The main idea. Every chapter of this book tried to correct an upset before it reached the thing being controlled, from a water tank to a spacecraft. Some plants measured the upset and acted at once. Others closed a small, fast loop around the part the upset hit first.

The challenge. An early action is only as good as the meter and the model behind it. A meter that drifts, or a model used far from where it was made, can push the plant the wrong way.

Why it matters. Asking what each early action reads, and where its model holds, separated the designs that worked from those that failed.

Chapter 1 reduced the volume to one rule: a structure helps against a disturbance only if it measures the disturbance, or a variable on its path, earlier than the controlled variable does. Eleven chapters then put that rule on eleven plants, from a tank with two upsets to a climber whose weight fades to nothing on the way to orbit. This chapter sets the results side by side, draws a short guide for choosing a structure, 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.

13.1 What the Chapters Showed

Tables 13.1 and 13.2 return to the path map of Table 1.3 with the measured result of each pairing. Each row names the disturbance path that made the plant hard, a structure that met it, one that did not, and the measured result with the table or section of the chapter that reports it.

The columns of the two tables make five points.

The entry point decided the result before any tuning did. On the tank of Chapter 2 each structure removed the upset on its own path and left the other untouched. The cascade divided the pressure error by 4.7 and did not improve the demand error, and the outlet feedforward did the reverse. Several chapters kept a disturbance that no structure reached: the probe lag of Chapter 3, the condenser of Chapter 5, the deceleration authority of a rocket with too small an engine. Those limits were stated, and no tuning moved them. Figure 13.1 sets one like pair of numbers from each row side by side.

A feedforward is only as good as its model, at the bandwidth paid for. The orbital feedforward was exact, and a mean motion 15.8 % low cost 22.6 % more propellant on a slow loop and 2.7 % on one twice as fast. The tank's trim was exact only at the mean supply pressure, which left 0.0444 m of

Table 13.1 Which Structure Matched Which Disturbance Path, Chapters 2 to 7

Chapter	Difficulty	What worked	What failed	Measured result
2, tank level	Supply pressure enters at the valve. Outlet demand enters at the vessel.	A flow cascade for the pressure, feedforward from the outlet flow for the demand, and both together.	Feedforward from the inlet flow, which the valve itself moves. An inner gain past its stability limit.	Worst error 0.2989 single loop, 0.0136 both structures, 0.3590 m inlet feedforward (Table 2.4). Cascade divides the pressure error by 4.7, feedforward the demand error by about 250 (Section 2.6).
3, fermenter oxygen	Agitator errors at the actuator. A feed step on the culture. Drive 15.4 s, holdup 12.4, probe 20 s.	A speed cascade settling in 2.9 s, with a carbon-balance feedforward.	The same cascade with its inner loop detuned to 12.3 s. Any structure against the probe lag.	IAE 10.6 single, 3.5 cascade, 2.1 with feedforward (Table 3.4). Detuned 72.3 against 23.8 (Section 3.6). Break-even at 7.4 s or 13.4 s by inner reset time (Figure 3.4).
4, fired heater	A 2 s fuel valve and a 12 s fan. Unmeasured heating value. A drifting oxygen analyzer.	Cross-limiting selectors. Oxygen trim with a carbon monoxide override.	Air slaved to the measured fuel on a load increase. A trim at ±25 % following its analyzer.	λ_b 0.990 and 3.7 s rich against at least 1.136, for a dip of 5.8 against 4.6 °C (Table 4.5). Over the permit 118 against 0 s (Table 4.6).
5, waste-heat cycle	Lost source heat reaches the superheat through the evaporator. The condenser sets a limit no loop moves.	The flow cascade with feedforward from the source meters.	The superheat loop alone, independent or in cascade. Any setting on 55 °C server water.	IAE 276.7, 437.2 and 29.0 K s, expander isolated 54.9, 83.4 and 0 s (Table 5.5). Best net output of 192 settings −0.59 kW (Section 5.5).
6, water plant	The dose acts 193.6 min to 63 % on a 30 min storm. Organic carbon only from the laboratory.	A dose computed from raw turbidity, with a slow multiplying trim.	Feedback alone, which integrated past the coagulation optimum to its 140 mg/L stop.	Above 1.00 NTU 47.15 against 0 min, alum 415 against 306 kg (Table 6.3). Trim at 600 min: 2.04 against 1.23 NTU target (Section 6.5.3).
7, carbonation line	Filler rate measured. Syrup strength unmeasured. A water header dip at the water valve.	Throughput feedforward, a ratio on the measured water flow and a Brix trim.	The ratio on the water setpoint when the valve cannot follow. The trim switched off.	Trim off: 4.60 m^3 of 6.67 off specification (Table 7.5). Header dip: 10.80 °Bx against 11.33 °Bx on the setpoint basis (Section 7.6).

the combined upset. Near the geostationary altitude a gravity error of 0.1 % of g_0 exceeds the whole effective gravity, so the model error is the disturbance. The water plant's model used the morning's laboratory organic carbon, and the settled water came to rest at 2.13 against 1.21 NTU.

The speed of an inner loop is a measured number, not a ratio from a rule. The fermenter's cascade stopped paying at an inner time constant of 7.4 s with one inner reset time and at 13.4 s with another. Below about 15 %/A the converter's outer loop inherited the inner lag. On the rocket a separation of 3.0 failed with a fast outer loop, while the same ratio reached by slowing the inner loop still landed. On the tank the inner loop's stability limit moved with the supply pressure inside each cycle. The textbook ratio of three to five (Section 1.4.5) is a starting point, and each of these chapters found the true boundary only by sweeping.

Table 13.2 Which Structure Matched Which Disturbance Path, Chapters 8 to 12

Chapter	Difficulty	What worked	What failed	Measured result
8, nanoparticle skid	Four feeds in ratio. Weak lots no flow meter sees. Size reported a minute late.	A supervisor choosing loop and flow, a size trim, a conductivity probe and a window guard.	Coordinated ratio control alone. The supervisor on the metered blend or without the guard.	Rate 441.7 against 600.0 mL/min, passing 6 of 6 seeds. Metered blend 776.7 mL off specification, no guard 1928.5 mL (Table 8.5).
9, boost converter	A right-half-plane zero that moves 17-fold with load. An overload.	An inner current loop with its clamp, and load-current feedforward.	The single voltage loop pushed harder. An inner gain below about 15 %/A.	Overload peak 9.47 against 11.32 A. Excursion 58.6 against 161.9 % (Section 9.6.4). Feedforward: recovery 43 against 804 ms, dip 7.82 against 5.69 V (Section 9.6).
10, rocket landing	The outer loop commands an angle. Zero at 2.80 rad/s. Weight carried at a tilt.	Bandwidth separation 8.7, gravity feedforward with $1/\cos\theta$, no position integral.	Separation 3.0 (a scripted run). The feedforward switched off before a Reset.	Position IAE 123.0 to 82.4 m s while the attitude IAE rose from 0.264 to 4.607 rad s, a hard arrival (Table 10.7). With a 10° tilt, descent IAE 6.83 against 12.94 m (Table 10.5).
11, orbital rendezvous	Gravity gradient and Coriolis terms couple the axes. The coupling is known.	Four model terms that cancel the coupling at the controller's mean motion.	The same gains without them. A mean motion 15.8 % low on a slow loop.	Docked at 1785 against 3516 s, along-track excursion 0.0074 against 11.475 m (Table 11.5). Wrong model: propellant penalty 22.6 at 2 mrad/s, 2.7 % at 4 (Section 11.6).
12, elevator climber	The weight fades to zero at 35 786 km and reverses. The climber's own moves excite a ribbon mode.	Feedforward of the effective gravity. ZV and ZVD shaping.	No feedforward at lift-off. A constant g_0 past the geostationary altitude.	Lift-off at 1.775 against 3.56 min. A false 107.9 kN drives the climber to 169.1 m/s (Table 12.6). Ringing after a move 21.36 against 2.31 m (Section 12.6.2).

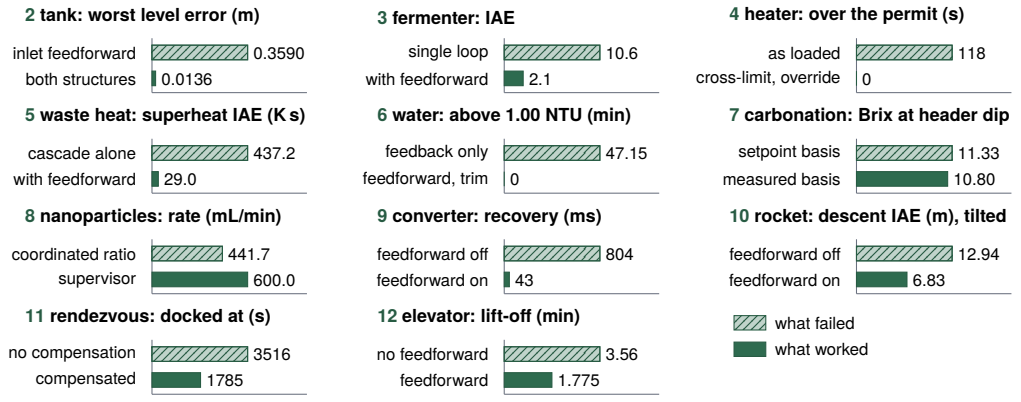


Figure 13.1 One like pair of numbers from each row, as the "Measured result" column of Tables 13.1 and 13.2 prints them. Each panel has its own scale from zero. The hatched bar is the structure or setting that failed and the solid bar the one that worked; a value of zero has a label and no bar. Notice that the carbonation panel is the one where the two bars nearly match: the setpoint basis moved the filler only from 10.80 to 11.33 degrees Brix, and that was enough to put product off specification.

An inner loop or a selector often earns its place by what it limits. The converter's single loop recovered faster than the cascade, 135 against 196 ms, with a smaller first dip. The cascade's value was a current held to 9.47 A in an overload and an outer loop that could be pushed. Cross-limiting on the heater bought no temperature at all. It deepened the outlet dip from 4.6 to 5.8 °C and removed every second of fuel-rich firing. Such structures are judged on the constraint they protect, not on the IAE.

A slow trim corrects what was not measured, slowly, and follows its instrument. The water plant's trim, with a 300 min integral time on a 190 min plant, had not converged six hours later. The heater's oxygen trim followed a drifting analyzer until the burners went rich, and only the carbon monoxide override held the permit. The mRNA skid's size trim made the rate, and the conductivity probe made the quality. In each plant the trim needed its own protection: a clamp, an override, or a second measurement that did not share the first one's fault.

13.2 Choosing a Structure

The chapters suggest an order of questions for a new disturbance. Each names the chapter where its answer decided the result, and Figure 13.2 draws the same order.

1. *Where does the disturbance enter, and which measurement sees it first?* Place it on the map of Section 1.4.2. A lag in the measurement itself is met by no structure (Chapter 3).
2. *Is it measured, and can its effect be computed?* Then feed it forward from a balance or from physics (Chapter 6, Chapter 11). Estimate the model error and ask whether the feedback loop is fast enough to absorb it.
3. *Does the measured signal depend on the manipulated variable?* Then it is not a disturbance measurement. The inlet feedforward of Chapter 2 put a pole at $z = -1$ and reversed the valve on 70.0 % of its moves.
4. *Does the upset enter between the controller output and a faster measurable variable?* Close an inner loop there, and sweep its speed against the outer lags before trusting it (Chapter 3, Chapter 10).
5. *Is one direction of a ratio error hazardous?* Cross-limit, and accept the slower response in the safe direction (Chapter 4).
6. *Is part of the disturbance unmeasured or the model uncertain?* Add a slow multiplying trim, with a clamp and an independent override (Chapter 4, Chapter 7).
7. *Is the plant nonminimum phase or short of actuator authority?* Then a feedforward trades undershoot for speed as a higher gain does (Chapter 9), and a capability shortfall needs a change of target or equipment, not of structure (Chapter 10).

A predictive controller earned its place where it carried a forecast or a hard constraint. Told of a turndown, it held the waste-heat cycle's superheat at 5.80 K where the cascade fell to 2.43 K (Table 5.7). The converter's explicit MPC reached the band 1.9 to 6.8 times sooner than the cascade (Section 9.6.6), and the orbital MPC docked 462.5 s sooner on 14.9 % less propellant (Table 11.8). It lost where its model was wrong. Under fouling the waste-heat MPC never settled, an IAE of 5583 against 335 K s for the cascade, and the landing MPC tipped over with a slow gimbal where the cascade landed (Sections 5.5 and 10.7). In the steady crosswind a bias feedforward added to the cascade cut the offset from 5.771 to 0.0308 m, as small as the MPC's, without a horizon. Volume 9 treats that choice in full.

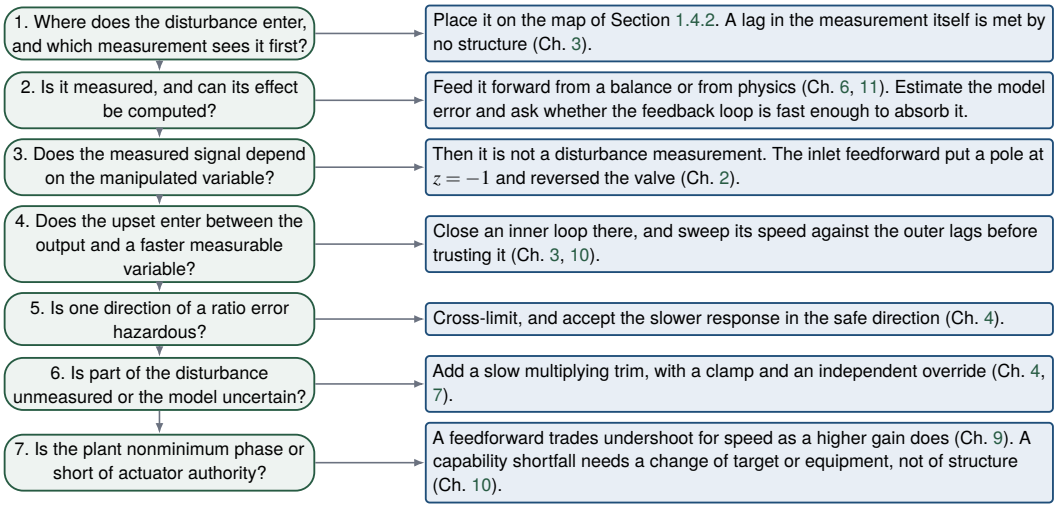


Figure 13.2 The seven questions of this section in order, with the list's answer and the chapters where it decided the result. Notice that the entry point and the measurement come first and the trim comes late: a slow trim corrects only what the structure above it has already placed correctly.

13.3 Where Practice Is Heading

Feedforward is tuned by rule and from data. Simple tuning rules for feedforward compensators have been published for process control [1], along with the optimal performance of low-order compensators and the practical considerations in using them [2]. A 2024 monograph is devoted to feedforward control and closes with a summary of its tuning rules [3]. In motion control, feedforward is tuned iteratively from measured data [4], and because accurate data are inexpensive to collect on such machines, models and controllers are identified and learned from experiments [5]. The orbital and elevator chapters show what that work protects against: a feedforward whose model is slightly wrong at the operating point that matters.

Unmeasured disturbances are estimated and fed forward. Disturbance-observer-based control lumps disturbance and model uncertainty into one total disturbance, estimates it and compensates it, and has been applied across industrial sectors for four decades [6]. Active disturbance rejection control is one member of that family, built around an observer of the total disturbance [7]. In the editor's view this is the natural successor to the slow trim of this volume: the crosswind of Chapter 10, the heating value of Chapter 4 and the organic carbon of Chapter 6 are all disturbances that an estimator could feed forward, provided its estimate is checked as carefully as an analyzer's.

Trims run on inferred quality. Data-driven soft sensors infer a quality from routine process measurements [8]. The mRNA skid of Chapter 8 trims on such an estimate between laboratory results, and it missed one of twelve seeds when neither the estimate nor the last reported size saw the true size leave its window. In the editor's view the trims of the next decade will run more often on inferred values, and the work of the engineer moves toward deciding which second measurement can catch an inference that is wrong.

Inner loops meet optimization. Predictive control has been applied to power converters and drives because it handles constraints and nonlinearities directly, at a high computational cost [9]. Powered-descent guidance has been posed as a convex program, solved by methods with deterministic

convergence, with a view to running it onboard [10]. In the editor's view cascades will remain the inner layer of such systems, as in Chapter 9 and Chapter 10, because they bound a current or an angle with a few lines of code that can be read and verified.

> ⧖ HISTORY **1956 to 2020: From a Cascade to a Hierarchy of Loops**
>
> The timeline of Chapter 1 places the first quantitative analysis of cascade control in 1956. A thesis on hierarchical control traces its field to those first results and calls cascade control a precursor of hierarchical control. The theory of multilayer control that followed in the 1960s rested on aggregation and on multiple time scales [11]. A 2020 review of predictive control in power electronics shows the same division built into hardware. In its example the fast controller runs in a gate array and takes its reference from an outer loop that runs at a lower frequency on a processor [12]. In the editor's reading, the tank of Chapter 2 and the converter of Chapter 9 stand at the two ends of that arc, and the sweeps of Section 13.1 show where the separation of time scales stops holding.

Feedforward is computed from an inverted plant model. On a waste-heat recovery cycle for a heavy-duty diesel engine, a dynamic feedforward obtained by inverting a nonlinear model was added to the commonly used gain-scheduled PID that holds the superheat. On an engine test bed the largest superheat error fell to about 1.9 K, against about 10 K without the feedforward, and the feedforward computation was shown robust enough for an embedded implementation [13]. The cycle of Chapter 5 feeds forward from its source meters at the same point. In the editor's view the step from a static balance to an inverted dynamic model is where feedforward in these plants is heading, and gap 1 of Section 13.5 is the price of that step.

Cascades are closed over networks. Cascade control systems, with at least two nested loops, are widely used in industrial processes, power plants and chemical processes. When the loops are closed through a real-time network, delays and packet loss in both loops can reduce performance or cause instability. Event-triggered execution has been proposed to lighten the load on the network, and cyber-attacks change the closed-loop dynamics [14]. In the editor's view this carries the inner-loop question of Section 13.1 onto the network. The inner gain of Chapter 2 was already capped by the delay of its own sampling, and a network adds delays that vary from one scan to the next.

Vehicles find their own relative state. A survey of small spacecraft technology reports a vision-based payload on a mission launched in 2023 that gave angles-only relative navigation of an object without prior knowledge of it, with a relative position error of 0.5 % of range. The same survey notes that much of this navigation must be done without commands from the ground, that guidance for proximity work is still primarily research based, and that many of these missions had difficulties with their control systems and actuators [15]. The feedforward of Chapter 11 cancels a coupling computed from the relative state and the mean motion, so in the editor's view the quality of that on-board estimate sets how far the feedforward can be trusted. Figure 13.3 sets developments of feedforward and cascade practice in three threads.

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 [20]. 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.

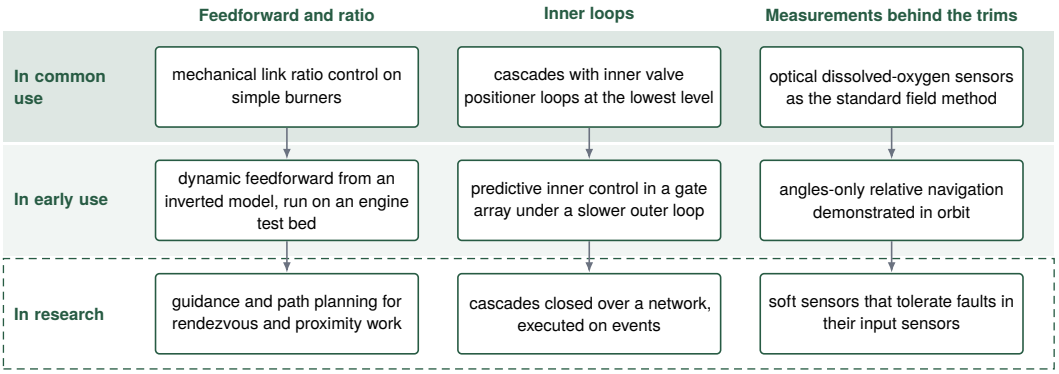


Figure 13.3 Developments in three threads. The top lane holds what the cited source calls standard, customary or common practice, the middle lane what it reports in service at some sites, on a test bed or in orbit, and the bottom lane what it reports from research, as a proposal or as a need, in the editor's reading. Sources by thread: feedforward and ratio [13, 15, 16], inner loops [12, 14, 17] and measurements behind the trims [15, 18, 19]. Notice that each thread moves from a fixed device toward a computation that must be checked while the plant runs.

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 [21]. 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 each number to a run, and ten of the eleven papers received a referee report, at least four of them a self-review by the agent that wrote the paper. Defects went into a log for each app. For this book one agent per chapter reran every number an exercise uses. The editor set the scope of each round and what each agent could change, and statements that runs contradicted were left for the editor. The record holds the following errors and the checks that caught them.

The defect logs. Ten of the eleven logs give each entry a severity, 102 entries in all. Of these, 32 are wrong behavior of an app (a wrong output, a wrong lesson, or a controller reading what it should not), 46 are documentation, and 19 are cosmetic or display. The water plant's log has no severity field. It records that eight quantitative claims in the app's teaching text were wrong when the model was driven, and that the model was right in every case.

A feedforward that knew too much. The fermenter's carbon-balance feedforward read a hidden flag that told it when a quarter of the sparger had plugged. No instrument reports that, and the challenge credited the feedforward with knowledge no plant has. On the converter, the cascade, its preload and its feedforward read the inductor resistance of the run in place of the declared value. With a changed resistance its duty differed from that of a controller fed the same measurements. Both were found by reading the code for what it reads, both now use what a plant would have, and the converter's fix moved no shipped run.

A teaching number that made a slow inner loop look safe. The fermenter's teaching text placed the cascade's break-even at an inner time constant of 10 s, 0.8 times the oxygen holdup, where the sweep gives 7.5, 0.6 times the holdup. The error ran in the direction that makes a slow inner loop look acceptable. The same text gave the drive lag as eight seconds where the code computes 15.4 s. The water plant's text said its trim closes the organic-carbon offset "completely" in about two hours. In the run it had moved the dose from 31.32 to 33.42 mg/L by 600 min, short of target. Each was caught by driving the shipped model and comparing.

A stability rule that was not a stability condition. The tank app's help text placed the inner flow loop's limit where a product of gains passes one, and called a gain of 1.5 the boundary. Writing the sampled recursion from the code gave the Jury limit of 1.925 at 12 bar, and runs found the loop well behaved at 1.5 and broken at 2.0. The same derivation showed the inlet feedforward sitting exactly on the unit circle at the shipped gain. The chapter now carries both results.

A metric that flattered the tuning. On the converter the overshoot tile read 49.3 % where the logged output peaked at 58.2 % of the step. The tile latched the starting value one sample after the step, when an inverse-response plant had already moved the wrong way. The paper agent found it by checking the tile against the exported trajectory. The rocket's feedforward-off comparison gave two answers, a descent IAE of 6.668 or 12.94 m, depending on whether Reset came before or after the switch. Its chapter reports both paths.

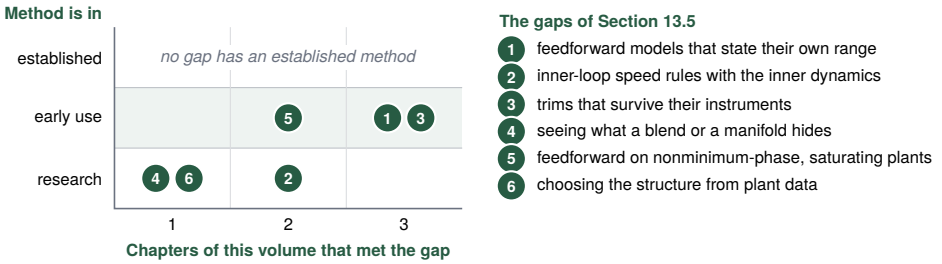


Figure 13.4 The six gaps of this section. Across is the number of chapters that each item names as evidence: three for gap 1 (Chapters 2, 11 and 12), two for gap 2 (3 and 2), three for gap 3 (4, 7 and 8), one for gap 4 (8), two for gap 5 (9 and 10) and one for gap 6 (2). Up is the maturity of the method available, the editor's reading of the sources cited in each item and in the Open challenge box. Notice that the two gaps met by the most chapters, feedforward models that state their range and trims that survive their instruments, have methods in early use, while rules for inner-loop speed, measurement per stream and choosing a structure from data are still research.

An explanation that was wrong. A draft of the heater paper said the trim clamp did nothing in the challenge because the trim had "not yet" reached 8 %. Measured, the drift drove the trim down, to only –0.3 % when carbon monoxide crossed the permit. The self-review caught it by measuring the trim instead of reasoning about it.

The book round. The eleven chapter records list 184 rows of numeric verification, each a stated value beside its reproduction. Nine chapters showed no drift. Two showed one number each: the water plant's storm IAE of 1553.5 now runs to 1555.6, and the elevator's constant-g_0 transfer docked 0.353 m from the mark where the chapter said 0.375. Both were listed for the editor, in a list of six entries for this volume that also held app texts with stale numbers and an orbital sentence implying a compensation the challenge does not need. A later repair round corrected both numbers in the papers and the chapters, and the orbital sentence now reads correctly in the chapter. The agent writing Chapter 1 narrowed two of its own first drafts, "every chapter enters the feedforward as a PID bias" among them, after checking them against the sources and its script.

Every error above was caught by a check that could fail: a run beside a quoted number, a recursion written from the code, a hidden value read where a measurement belonged, a tile set against the trajectory. None was caught by rereading the text that contained it. For feedforward and cascade work the record adds one rule to those of the shared section. Ask of every feedforward what it reads, and confirm that a plant could measure it.

13.5 Gaps in Feedforward, Cascade and Ratio Practice

Six gaps belong to the topic of this volume, and the gaps common to agent-era practice follow in Section 13.5.1. 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 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 plants in the picture that opens this chapter.

1. *Feedforward models that state their own range.* The tank's trim was exact at one supply pressure and the orbital terms at one mean motion, and a constant-g_0 weight on the elevator holds only near the ground. A textbook of feedback systems notes that feedforward, because it matches two signals, requires good process models, without which its corrections may have the wrong size or be badly timed [17]. A standard for models and simulations requires a record of each validated model's domain of validation and allows use outside it only within the model's limits [22].

The gap is hard because the operating point that leaves the model's range is often the one the feedforward exists for, as at the geostationary altitude. A feedforward that reports when the operating point has left the region where its model holds, and backs off, is not standard practice.

2. *Inner-loop speed rules that include the inner loop's own dynamics.* The fermenter's break-even moved from 7.4 to 13.4 s with the inner reset time alone, and the tank's inner limit moved within each pressure cycle. An open textbook states as non-negotiable that the inner loop respond faster than the outer, and that cascade control cannot function if the relation is reversed, but gives no ratio [16]. The gap is hard because the inner integral, the sample period and a moving operating point each change the closed inner loop that the outer design assumes. A design rule that accounts for all three would replace a sweep.

3. *Trims that survive their instruments.* The heater's trim followed a drifting analyzer, the carbonation analyzer read dissolved air as carbon dioxide, and the mRNA soft sensor and the last reported size both stayed in the window while the true size left it. Clamps and overrides were set by hand in each chapter. The Open challenge box after this list takes the gap further.

4. *Seeing what a blend or a manifold hides.* In Chapter 8 13 500 mL of the 15 409 mL a passing run accepted was made while a channel was outside its own window, and one flow meter read its setpoint throughout. An industrial review of lipid nanoparticle manufacturing notes that scaling out by parallel mixing lines risks mixer-to-mixer variability, with inconsistent flow rates, pressure imbalance or channel clogging, often leading to non-uniform product [23]. The gap is hard because a meter on the shared line sees only the sum, and a meter on every stream multiplies the instruments as the streams multiply. Ratio control of many parallel streams needs measurements or estimates per stream.

5. *Feedforward on nonminimum-phase and saturating plants.* The converter's feedforward deepened the dip from 5.69 to 7.82 V, and the rocket's engine could not supply the deceleration the profile asked for. A review of predictive control in power electronics names solution in real time, within a few microseconds, as the main challenge, and notes that converter models are often linearized at a loss of accuracy [12]. The gap is hard because the undershoot belongs to the plant, so a faster correction buys recovery only with a deeper first dip. Design methods that respect undershoot and actuator authority from the start are thinner than those for a minimum-phase loop.

6. *Choosing the structure from plant data.* This volume chose each structure by reading the disturbance path from a model. The tank of Chapter 2 offered two flow meters, and the one its own valve moves made a feedforward that reversed the valve on 70.0 % of its moves. A survey of industrial predictive control listed which variables should be used for control, and how to design a control system for an entire plant, among the needs industry perceived [24]. The gap is hard because routine records hold the plant under its existing structure, so the path a disturbance would take under another structure is seldom in the data. A method that proposes the feedforward signal and the inner variable from plant records, and says which entry point the records cannot resolve, would turn the path map of Table 1.3 into a tool.

⚠ OPEN CHALLENGE A Trim That Knows When Its Instrument Is Wrong

Three chapters met gap 3: an oxygen analyzer that drifted, a carbon dioxide analyzer that counted dissolved air and a size estimate that missed its window. A critical review of soft sensors for bioprocesses names erroneous inputs from sensor faults among the greatest challenges, finds that tolerance to sensor faults remains the greatest, and notes that a departure of a prediction from the hardware reading indicates a sensor fault [19]. In the editor's view the gap is hard because a drifting instrument still reports a plausible value, so nothing marks the reading the trim should stop following. Progress would be a trim that carries a second, independent measurement and a stated bound on the harm a failed instrument can cause.

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 [25]. 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 [21]. 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 [26]. 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 [27]. 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 [28]. 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 structure helped only against a disturbance that it saw earlier than the controlled variable did.
- A feedforward was only as good as its model, and the speed an inner loop needed was found by sweeping.
- Inner loops and selectors earned their place by the limits they kept, and slow trims needed guarded instruments.

References

[1] J. L. Guzmán and T. Hägglund, "Simple tuning rules for feedforward compensators," *Journal of Process Control*, vol. 21, no. 1, pp. 92–102, 2011.

[2] M. Hast and T. Hägglund, "Low-order feedforward controllers: Optimal performance and practical considerations," *Journal of Process Control*, vol. 24, no. 9, pp. 1462–1471, 2014.

[3] J. L. Guzmán and T. Hägglund, *Feedforward Control*. Berlin: De Gruyter, 2024.

[4] F. Boeren, T. Oomen, and M. Steinbuch, "Iterative motion feedforward tuning: A data-driven approach based on instrumental variable identification," *Control Engineering Practice*, vol. 37, pp. 11–19, 2015.

[5] T. Oomen, "Advanced motion control for precision mechatronics: Control, identification, and learning of complex systems," *IEEJ Journal of Industry Applications*, vol. 7, no. 2, pp. 127–140, 2018.

[6] W.-H. Chen, J. Yang, L. Guo, and S. Li, "Disturbance-observer-based control and related methods: An overview," *IEEE Transactions on Industrial Electronics*, vol. 63, no. 2, pp. 1083–1095, 2016.

[7] J. Han, "From PID to active disturbance rejection control," *IEEE Transactions on Industrial Electronics*, vol. 56, no. 3, pp. 900–906, 2009.

[8] P. Kadlec, B. Gabrys, and S. Strandt, "Data-driven soft sensors in the process industry," *Computers & Chemical Engineering*, vol. 33, no. 4, pp. 795–814, 2009.

[9] S. Vazquez, J. I. Leon, L. G. Franquelo, J. Rodriguez, H. A. Young, A. Marquez, and P. Zanchetta, "Model predictive control: A review of its applications in power electronics," *IEEE Industrial Electronics Magazine*, vol. 8, no. 1, pp. 16–31, 2014.

[10] B. Açıkmeşe and S. R. Ploen, "Convex programming approach to powered descent guidance for Mars landing," *Journal of Guidance, Control, and Dynamics*, vol. 30, no. 5, pp. 1353–1366, 2007.

[11] D. Gromov, "Analysis of hierarchical structures for hybrid control systems," Ph.D. dissertation, Technische Universität Berlin, Berlin, 2010.

[12] P. Karamanakos, E. Liegmann, T. Geyer, and R. Kennel, "Model predictive control of power electronic systems: Methods, results, and challenges," *IEEE Open Journal of Industry Applications*, vol. 1, pp. 95–114, 2020.

[13] J. Peralez, M. Nadri, P. Dufour, P. Tona, and A. Sciarretta, "Organic Rankine cycle for vehicles: Control design and experimental results," *IEEE Transactions on Control Systems Technology*, vol. 25, no. 3, pp. 952–965, 2017.

[14] M. Fallahnejad, "A survey of recent breakthroughs in networked cascade control systems," *International Journal of Robust and Nonlinear Control*, vol. 36, no. 5, pp. 2531–2541, 2026.

[15] S. V. Weston, C. D. Burkhard, J. M. Stupl, R. L. Ticknor, B. D. Yost, R. A. Austin, P. Galchenko, L. K. Newman, and L. Santos Soto, "State-of-the-art small spacecraft technology," National Aeronautics and Space Administration, Ames Research Center, Moffett Field, CA, Technical Publication NASA/TP-20250000142, Jan. 2025. [Online]. Available: https://ntrs.nasa.gov/citations/20250000142

[16] T. R. Kuphaldt, *Lessons in Industrial Instrumentation*, version 2.33 ed. Open textbook, Creative Commons Attribution 4.0 license, ibiblio.org/kuphaldt/socratic/sinst, 2024.

[17] K. J. Åström and R. M. Murray, *Feedback Systems: An Introduction for Scientists and Engineers*, 2nd ed. Princeton, NJ: Princeton University Press, 2021.

[18] U.S. Geological Survey, "Dissolved oxygen," U.S. Geological Survey, Techniques and Methods, book 9, chap. A6.2, 2020.

[19] V. Brunner, M. Siegl, D. Geier, and T. Becker, "Challenges in the development of soft sensors for bioprocesses: A critical review," *Frontiers in Bioengineering and Biotechnology*, vol. 9, p. 722202, 2021.

[20] 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.

[21] 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.

[22] National Aeronautics and Space Administration, "Standard for models and simulations," Office of the NASA Chief Engineer, Washington, DC, NASA Technical Standard NASA-STD-7009B, 2024, approved 2024-03-05, standards.nasa.gov.

[23] J. Hong Hoang, M. Ott, E. Samaridou, M. Beck-Broichsitter, and J. Simon, “Industrial perspective on the manufacturing of lipid nanoparticles for nucleic acid delivery,” *Pharmaceutics*, vol. 18, no. 4, p. 489, 2026.

[24] S. J. Qin and T. A. Badgwell, “A survey of industrial model predictive control technology,” *Control Engineering Practice*, vol. 11, no. 7, pp. 733–764, 2003.

[25] 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.

[26] 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.

[27] 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.

[28] 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 feedforward for a level loop on a tank fed through a control valve. Its signal is the inlet flow transmitter, because “the inlet flow is the largest disturbance on the balance”. Using Chapter 2, explain what will happen, list the checks you would ask for before the feedforward is enabled, and say which structure the agent should have proposed for each of the tank’s two upsets.

13.2 Discussion · 8 min

A design note says “the inner loop of a cascade must be three to five times faster than the outer loop”. Using Chapter 3, Chapter 9 and Chapter 10, argue why the note is a starting point and not a design rule, and describe the measurement you would make on a new cascade instead.

13.3 Discussion · 8 min

The heater of Chapter 4 trims its air ratio on stack oxygen, the water plant of Chapter 6 trims its dose on settled turbidity, and the skid of Chapter 8 trims its buffer on a soft sensor of particle size. For each, say what the trim corrects that the feedforward or ratio cannot, what can make the trim do harm, and what protects against that harm in the chapter.

13.4 Discussion · 7 min

An engineer must choose between spending a month improving the model behind a feedforward and spending the same month making the feedback loop faster. Using Chapter 11 and Chapter 12, argue which choice pays in which circumstances.

13.5 Discussion · 10 min

Section 13.4.1 describes errors that the agents building this volume made. For the feedforward that read a hidden fault flag, the teaching number that made a slow inner loop look safe, the stability rule that was not a stability condition, and the overshoot tile, name the check that caught each. Then write a procedure, in five to eight steps, for an agent that will add a feedforward to a running loop on a plant, stating what the agent may read, what it may write and what a person must approve.

13.6 Discussion · 8 min

A vendor proposes to replace the cascades of a waste-heat power cycle and a rocket lander with model predictive controllers. Using Chapter 5 and Chapter 10, write the questions you would ask before agreeing, and say what evidence would change your mind in each direction.