INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 19 Control of Biological Systems and Human Performance

CHAPTER 15

Summary and Outlook

Pages 373–387

Edited by John D. Hedengren

How to cite this chapter. "Summary and Outlook," in *Control of Biological Systems and Human Performance*, John D. Hedengren, Ed., vol. 19 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 15, pp. 373–387.

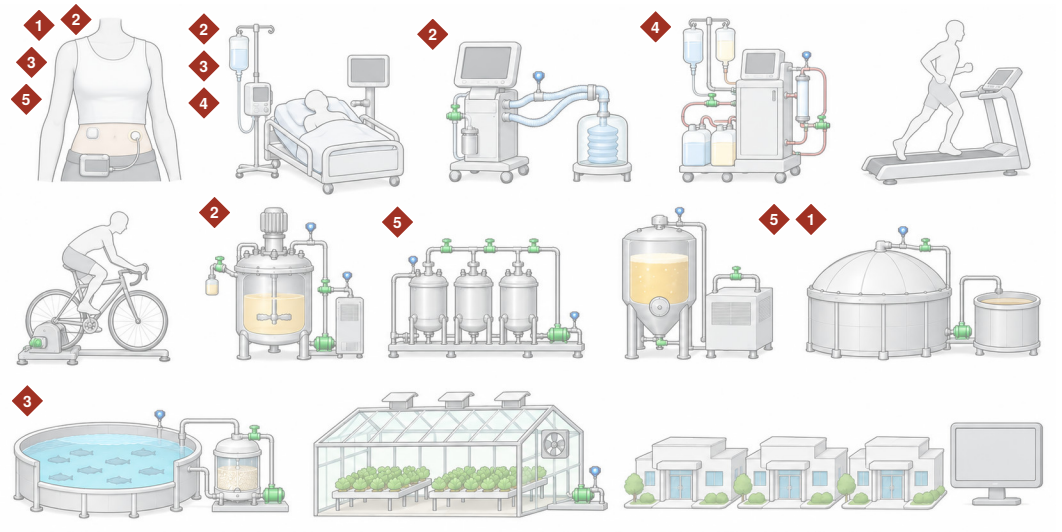
The page numbers in this file are those of the printed volume.

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v19-biological-human/

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

CHAPTER 15

Summary and Outlook



Open gaps 1 unmeasured variability 2 virtual populations 3 one-way actuators 4 measurement checks 5 objectives

Notice that the insulin pump carries four gaps and the anesthesia bed three, and that no gap marks the athletes, the fermenter, the glasshouse or the city. Illustration, not to scale.

IN PLAIN TERMS **What Thirteen Living Plants Taught**

The main idea. Each chapter of this book put an automatic loop on something alive: a patient, an athlete, a culture of cells, a farm or a city. The loop sets a pump, a valve or a pace by itself. No two were alike, and what mattered was often measured late or not at all.

The challenge. A pump that adds a drug cannot take it back. A reading can look right while the body drifts away, and a score that treats a low and a high alike can pick the wrong loop.

Why it matters. The loops that worked measured the effect, checked their readings, kept a margin and were judged on what the plant is for.

Thirteen chapters put a loop on a patient, an athlete, a culture, a farm and a city. Section 1.4.7 read each plant in three parts: what varies between individuals or during the run, the limit the plant may not cross, and the structure that keeps it inside. This chapter adds the measured results to that reading and compares how each loop handled variability and limits it could not exceed. It then gives a short guide for choosing a structure, looks at where practice is heading, reports how the volume was made with AI agents and what they got wrong, and lists the problems that remain open.

15.1 What the Chapters Showed

Tables 15.1 and 15.2 extend Table 1.1. Each row names what varied and the limit that bound, a structure that held the plant, one that did not, and the measured result with the table of the chapter that reports it. Every number is the chapter's own, and Figure 15.1 draws one pair from each row.

The columns of the two tables make five points.

Table 15.1 How Each Loop Handled Variability and Its Limits, Chapters 2 to 8

Chapter	What varied, the limit	What held it	What failed	Measured result
2, insulin pump	Exercise raises insulin sensitivity unseen. The pump stops at zero, and a low costs more than a high.	Direct-acting PI with back-calculation. A predictive controller told of the meal.	A symmetric score, which ranks the exercise day above the large meal. The predictive controller not told.	In band 78.9 to 100 %. Exercise low 47.8 mg/dL at an IAE of 118.5 against 210.3 (Table 2.5). Meal IAE 50.94 told, 98.98 not told, 83.67 PID (Table 2.7).
3, anesthetic depth	Sensitivity 0.70 to 2.0 times the model's. A depth band, a pressure floor, and a way down only by elimination.	A depth loop that moves the target of a model-inverting pump. A signal-quality hold. A vasopressor loop.	Model inversion alone. A controller that trusts the monitor through cautery.	Index spread 18.3 open, 1.1 closed (Table 3.7). Cautery low 31.8 against 48.9 and 333 mg against 178 (Table 3.6). Challenge in band 17 to 91 % (Table 3.8).
4, ventilator	Compliance, resistance, effort and leak. A plateau limit of 30 cmH_2O.	Identification on every breath under a protective cap. Leak compensation with adaptive support.	Volume control on a stiff lung. Pressure control on the same lung. Volume control with a leak.	Plateau up to 31.5 at 420 mL, or 179 mL at a held pressure. With the cap 393 to 396 mL at 29.9 to 30.0 (Table 4.5). Hour in band 27.34 against 99.74 % (Table 4.6).
5, kidney circuit	Pump tubing, wear, scale drift and fouling. Filter and access pressures, and the patient's true balance.	A weighed balance with each scale checked against its pump. Overrides. Effluent headroom.	Pump counts alone. A balance loop on a drifting scale. A drift under the check's tolerance.	466.3 mL behind in a day on counts. Error held within 28.4 while the patient falls 304.3 mL behind (Table 5.6). 35 g/h never flagged (Table 5.4). Challenge −893.0 to −51.2 mL (Table 5.7).
6, runner	Runner, heat and fuel. Lactate, heart rate and core temperature, and an hour to run.	Two PI loops on a low selector with pace feedforward. A predictive controller with a rollout terminal value.	The selector without feedforward. A terminal value written as a frozen product.	9.584 km selector, 9.947 predictive, both compliant, 8.676 without feedforward (Table 6.8). 210 of the 363 m by 3 min (Table 6.9). Under-fueled 8.35 against 9.09 km (Table 6.10).
7, cyclist	Rider, and a work capacity that drains. Heart rate and lactate limits.	A power loop with two overrides on margin-scaled reset. A predictive controller on the grade underfoot.	The treadmill's fixed reset band. A target above critical power. The whole route given to the predictive controller.	Circuit 28.90 km PID, 29.05 predictive, 28.97 with the route, 28.80 and not compliant at 600 W (Table 7.8). Fixed band 18.50 km (Table 7.11).
8, antibody culture	An appetite that doubles, a cooled culture, probe drift. Glucose 1 to 6 g/L, osmolality.	A continuous feed with an hourly Raman trim. An osmolality cap.	A daily bolus. A capacitance feed without the Raman. A glucose target of 4.5 g/L.	Titer 4.14 by bolus, 34.8 h outside, 3.43 without the Raman, 5.03 with it (Table 8.5). Challenge 4.25 to 5.21 g/L (Table 8.6).

Table 15.2 How Each Loop Handled Variability and Its Limits, Chapters 9 to 14

Chapter	What varied, the limit	What held it	What failed	Measured result
9, insulin refolding	Folding constants and the harvest load. An unfolded ceiling, hold clocks and a 72 h campaign.	A charge scheduled on an assay-corrected estimate at a loose ceiling.	A tight ceiling. A late harvest. A plant that folds 20 % slower than the controller assumes.	5.56 against 5.47 g/h for a bolus. At 0.35 g/L a lot folds 79.6 %, 665 g expire and the campaign fails (Table 9.7). Slow folding 3.68 against 5.01 (Table 9.6).
10, lager fermenter	Yeast viability, mixing that stops with the gas. A recipe within 1 °C, diacetyl, shared glycol.	Time-proportioned PI on three probes. A rest ended on the corrected estimate. A glycol coordinator.	One mid-height probe. A rest ended on the day-old laboratory result. The estimate without its correction.	Batch done on day 12.57 with the cone 2.73 °C off, then day 10.47 within 0.52 (Table 10.6). Uncorrected, the beer ships at 0.123 mg/L (Table 10.5).
11, digester	Feed, temperature and a hidden inhibitor. A titration ratio of 0.4 and a fold at 287.6 m^3/d.	A sampled PI on the daily titration under a supervisor.	A constant feed. A ratio target of 0.28. A predictive feed under hidden inhibition.	Constant feed sours at \$166,751, the supervisor passes at \$428,198 (Table 11.6). Predictive feed 6 high reports against 1 and less revenue (Table 11.5).
12, fish farm	A biofilter that must grow, and the fish biomass. Nitrite, ammonia and oxygen.	A feed governor, an oxygen soft sensor and an economic ration.	The feeding table. Fingerlings stocked on new media.	Hours over 57.0 to 20.7 (Table 12.4). Challenge 69.8 h over, 1,359 killed and 0 of 11 seeds, then 6.6 h, none killed and 11 of 11 (Table 12.5).
13, glasshouse	Weather and its forecast. A daily mean temperature and an hourly humidity limit.	A planner over the forecast together with the humidity override.	Either one alone. Fixed setpoints.	Shipped \$4292, three targets failed. Planner alone \$3780 with 715 min humid. Override alone \$6312. Both \$5594 and pass (Table 13.6).
14, city	Reported share, testing and a variant. Hospital beds.	Estimators anchored by the census. A scenario tree, or a nominal plan with a 20 % margin.	PI on the census. A plan with a 5 % margin. The shipped threshold.	PI peaks at 1,749 beds for a target of 900 (Table 14.4). Shipped and 5 % margin pass 0 of 7 records, tree and 20 % margin 7 of 7 (Table 14.5).

Variability was handled by feedback on the effect, when the effect was measured. Model inversion in Chapter 3 delivered the modeled concentration to every patient and left the depth index spread over 18.3 units. Closing a loop on the depth monitor held the same patients within 1.1 units, because the loop moved the commanded concentration from 1.90 to 5.03 µg/mL. That is the P-1 example of Section 1.4.4 at full scale. The athlete chapters handled the individual with a schedule on one known number, threshold pace or critical power, and Chapter 4 identified the lung on every breath. Where the variation reached no measurement, nothing caught it in time. Exercise in Chapter 2 took glucose from 74.9 to 47.8 mg/dL before the controller next executed (Section 2.7), and a hidden inhibitor in Chapter 11 turned a predictive feed into the loser of the comparison.

A loop that holds its measurement need not hold the plant. The balance loop of Chapter 5 kept its measured error within 28.4 mL while the patient fell 304.3 mL behind, because it followed a drifting scale exactly. The ventilator of Chapter 4 read 1521 to 1830 mL behind a mask leak while the lung received 103 to 414 mL (Section 4.6). The anesthesia controller that trusted the monitor through cautery let the true index fall to 31.8. The remedy in each case was a second view of the same

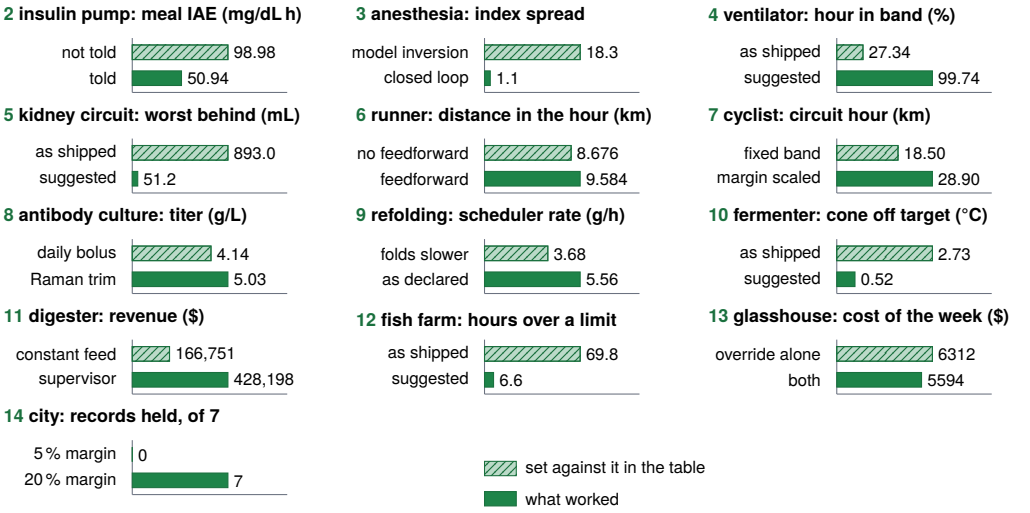


Figure 15.1 One pair of results from each chapter, with the numbers as the "Measured result" column of Tables 15.1 and 15.2 prints them. Each panel has its own scale from zero. A larger value is better in the panels of Chapters 4, 6, 7, 8, 9, 11 and 14, and a smaller one in the others. The kidney circuit's worst errors of −893.0 and −51.2 mL are drawn as amounts behind, and the city's 0 of 7 records draws no bar. Notice that in most rows the better result came from a change of structure or measurement, not of tuning alone.

quantity: the pump counts, the leak estimate, the signal-quality hold, three probes in the brewery tank and the census beside the case counts of Chapter 14.

One-way actuators were held by margin and by acting early. A pump that only adds leaves the way down to the plant. In Chapter 2 the pump sat at zero for 0.84 h while the insulin already given went on acting. In Chapter 3 the index fell to 35.5 a quarter of a minute after a closing, and only lightening six minutes ahead stopped the dip at 43.3 (Section 3.6). The governor of Chapter 12 and the humidity override of Chapter 13 act before their limits, not at them. The margin must cover the spread of the plant. On the epidemic's variant a 5 % margin breached every record and a 20 % margin held every one.

The objective decided what counted as held. A symmetric IAE in Chapter 2 ranked a run that fell to 47.8 mg/dL above one that never fell below 64 mg/dL, and the ranking reversed only at a low-side weight of 7.9. In Chapter 9 the best fold of a lot cost the campaign. In Chapter 11 the mean margin to the fold fell from 54.7 to 32.7 m^3/d as the revenue rose from $406,506 to $440,312 (Section 11.6). The glasshouse planner gave the cheapest week and failed the humidity limit. A score has to say which side of a limit matters and what the plant is paid for.

Prediction paid where it had a deadline and an anchored estimate. The predictive controller of Chapter 6 won the hour by 363 m, and that of Chapter 7 covered 29.05 against 28.90 km. More than half of the runner's lead was won in the first three minutes, a share the regulator recovers by shortening its reference ramp. In Chapter 14 no threshold, PI or constant dial in a sweep of 1,957 held the beds at $R_0 = 4$ within the budget, while the predictive controller did on every estimator (Section 14.7). Prediction lost where the model missed the biology. Without the meal announced, the predictive insulin controller scored 98.98 against the PID's 83.67, and when the Raman was lost in

Chapter 8 the model-based feed on the predictive controller's own estimate made 5.06 against its 5.01 (Table 8.7). The estimate carried most of the value.

15.2 Choosing a Method

The chapters suggest an order of questions for a loop on a living plant. Each question names the chapter where its answer decided the result, and Figure 15.2 draws the same questions as a flowchart.

1. *What varies between individuals, and does any measurement reveal it?* If the effect is measured, close a slow outer loop on it and let it correct the model's target (Chapter 3). If one number is known per individual, schedule on it (Chapter 6, Chapter 7). If nothing reveals it, carry margin, because no loop will learn it in time (Chapter 2).
2. *Is the measurement the plant?* Check each measurement the loop holds against a second view of the same flow or state, with a stated detection limit. In Chapter 5 a drift of 35 g/h passed the check and left the patient 379.3 mL behind.
3. *How old is the sample, and how fast does the plant grow?* The uptake between samples must fit inside the band, Equation (1.7). A culture held warm in Chapter 8 eats more in a day than its band is wide, so no daily bolus holds it (Section 1.4.5). Correct the model at the time of the draw.
4. *Can the actuator take back what it gave?* If not, size the margin for the spread of individuals plus the slow way back, Equation (1.9), and act on the event that will arrive, not on the limit it will reach (Chapter 3, Chapter 12).
5. *Which limit binds, and on which loop?* Give each limit an override on a selector, and test the selector's rules one by one. Removing the margin scaling of Chapter 7 cost 10.4 km of the hour.
6. *What is the plant paid for?* Score the two sides of a limit separately, Equation (1.10), judge the campaign rather than the vessel (Chapter 9) and the season rather than the day (Chapter 11).
7. *Is there a deadline, a forecast or an announcement?* Then a plan can earn its place (Chapter 6, Chapter 13, Chapter 14), provided its estimate is anchored by a measurement and its margin is sized for the futures it cannot see.

15.3 Where Practice Is Heading

Automated dosing at the bedside. Automated insulin delivery has reached clinical care. A joint consensus report of the European and American diabetes associations records that regulators have approved the first systems, that they have been adopted into care and that a community of people with type 1 diabetes builds its own, and it reviews the systems with a focus on their safety [1]. A meta-analysis of twelve randomized trials with 1284 patients found that index-guided closed-loop delivery of intravenous anesthesia used less propofol at induction, held the target depth more often and shortened recovery [2]. A survey of the field finds that anesthesia control has moved from single controllers to strategies of several components without yet reaching an integrated system [3]. Physiological closed-loop ventilation has moved from gas exchange toward preventing lung injury, early commercial systems are appearing, and many clinicians do not yet know them, which limits their adoption [4]. In the editor's view the lessons of Chapters 2 to 4 mark the edge of this work: an unannounced meal, a one-way pump, and a monitor that can be wrong while it looks right.

Evidence from simulation. Regulators have started to accept evidence obtained by modeling and simulation, and a framework for the credibility of such models asks for a stated context of use, a risk analysis that sets the acceptance thresholds, and verification, validation and uncertainty quantification [5]. The apps of this volume are verified and not validated, and each chapter says so. In the editor's view that distinction, made routinely and stated with every result, is the first thing a simulator used for clinical or manufacturing decisions must carry.

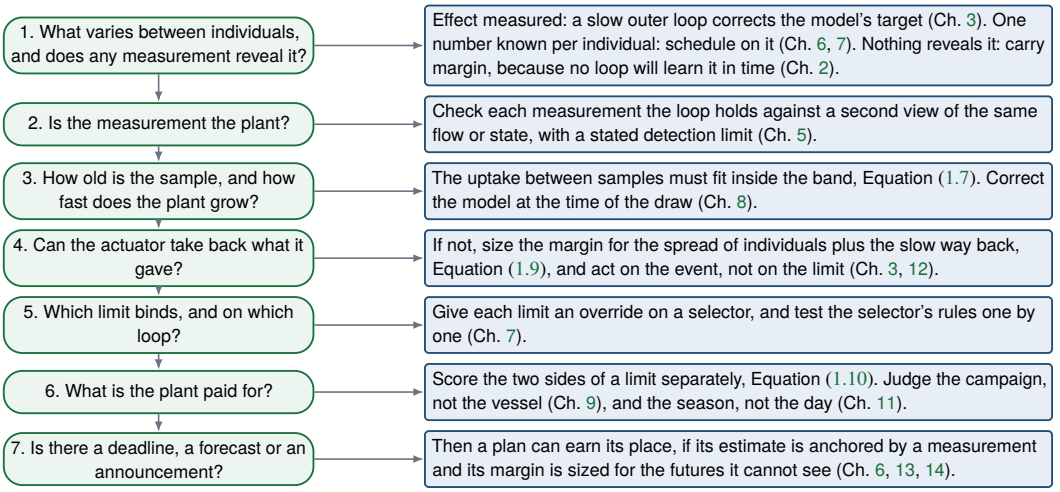


Figure 15.2 The seven questions of this section in order, with the list's answer and the chapters where it decided the result. Notice that the questions about the individual and the measurement come first, and that only the last question asks for a plan.

> ⧖ HISTORY **1978 to 2017: From Separate Pumps to One Machine**
>
> The first history box of Chapter 1 left the kidney circuit in 1978, balanced by gravity with scales and electronic clamps. A review written forty years after the first continuous treatment records what followed. Peristaltic pumps soon replaced those simple mechanisms, and veno-venous circuits took over with machines that carry a blood pump. The first such circuits joined separate pumps that could not communicate and operate together, with the possibility of amplifying errors. Integrated machines with four pumps and a pre-assembled circuit then made the therapy possible in almost every intensive care unit. The delivered dose still often differed from the prescription [6]. In the editor's reading the arc ends where Chapter 5 begins: one machine, one balance, and a check on every pump.

Models and spectra in biomanufacturing. A review of process models in biopharmaceutical manufacturing finds that the sector lags other industries in digitalization and names models that turn process data into decisions as key enablers under quality by design and process analytical technology [7]. Raman spectroscopy has been integrated into real-time release testing, continuous manufacturing and statistical process control, with a growing number of in situ applications in bioprocesses [8]. The hourly Raman trim of Chapter 8, worth 5.03 against 3.43 g/L of titer, is a small instance of that direction.

Autonomous growing and estimated epidemics. In the second autonomous greenhouse challenge, five teams grew cherry tomatoes for six months with climate and crop strategies set by their own algorithms, and every team outperformed the greenhouse operated by people [9]. For epidemics, guidance on estimating the effective reproductive number names two challenges common to all methods, the generation interval and the reconstruction of infections from observations made long after transmission, and warns that naive treatment of observation delays introduces bias [10]. That is the problem Chapter 14 solves first, before any policy acts. In the editor's view, control of living plants will move toward loops whose estimators carry most of the design, with the controller downstream of them.

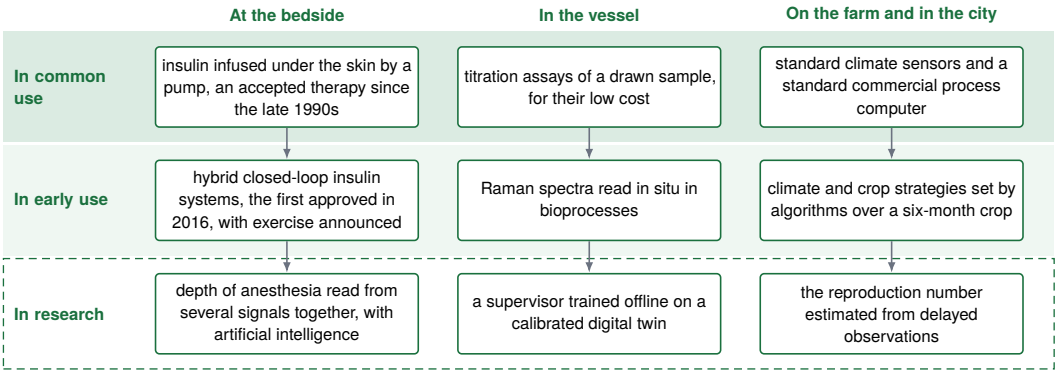


Figure 15.3 Developments in three threads. The top lane holds what the cited source calls standard, accepted or common, the middle lane what it reports in service or grown on a real crop, and the bottom lane what it reports from simulation or as a direction still to be shown, in the editor's reading. Sources by thread: at the bedside [11, 14, 15], in the vessel [8, 12, 13] and on the farm and in the city [9, 10, 16]. Notice that each thread moves from a reading taken by hand or by a standard instrument toward a decision made by a model.

Depth read from more than one signal. A review of depth-of-anesthesia monitoring notes that the indices in use reflect only the statistical features of cortical electrical activity. It expects multimodal monitoring, brain network analysis and artificial intelligence to improve the accuracy and the individualization of anesthesia management [11]. The depth loop of Chapter 3 closes on one index, and its signal-quality hold exists because that index can be wrong while it looks right. In the editor's view a loop on several signals will need the same check on each of them.

Online acids for the digester. A review of sensors for anaerobic digestion finds that simple titration assays remain common in research and in industry, largely for their low cost, while offline sampling introduces delays and limits the temporal resolution. Microbial electrochemical sensors have reached an accuracy comparable to gas chromatography for the volatile acids, and electrochemical sensors remain constrained by their longevity, reproducibility and biofouling [12]. The supervisor of Chapter 11 acts on one titration a day. In the editor's view an online acid reading would shorten that sample period only if it carries its own check against the titration.

Supervisors trained on a digital twin. A study of a full-scale digester notes that such plants are often operated conservatively because of delayed process feedback, nonlinear stability constraints and limited observability. It trained a supervisory policy offline on a calibrated mechanistic digital twin, without experiments on the live plant, and found a modest potential gain of 2.7 percent in biogas production, subject to substantial uncertainty [13]. The comparison of Chapter 11 gives the same caution, since its predictive feed lost to the PI loop under a hidden inhibitor. In the editor's view a supervisor trained in simulation inherits the credibility of its model, which is gap 2 of Section 15.5. Figure 15.3 sets the developments of this section in three threads.

15.4 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 [17]. 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 [18]. 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.

15.4.1 How This Volume Was Made

Each of the thirteen 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. Seven of the thirteen referee reports call themselves self-reviews or an internal review by the drafting agent, and the other six write in the voice of a reviewer. None names a reviewer independent of the writing. Four carry later referee passes on predictive-control or estimator material, three of them in the insulin pump's paper. The defect logs hold 157 entries with a severity. Of these, 33 are wrong output or behavior, 86 documentation or

content, 24 cosmetic or display and 14 something else, such as a completion rule or an entry checked and found not to be a defect. Some round notes list fixes without a severity, so the counts are a floor. The editor set the scope of each round and what each agent could change. Four cases show what went wrong and what caught it. In each, a loop met variability or a limit in a way its text did not say.

Controllers that knew the individual. Four entries in four logs record a controller that read the run's hidden parameters or the plant's true state. The runner's predictive controller started every plan from the runner's hidden state and stepped the run's own parameters. With the readings held fixed, 24 of 50 parameters scaled by 1.25 moved its plan, and a hidden core temperature raised by one percent moved its predicted heart rate. In the rider's predictor 35 of 61 parameters and an unmeasured headwind moved the plan, and under the fingerstick it read the true lactate in place of the reading. The insulin pump's controller took its starting bias from the run's balances, and the antibody culture's estimators started from the broth itself. In a volume about variability between individuals, a controller that reads the individual's true parameters has no variability to handle. Each fix gave the controller a declared individual and its own estimate. The runner's predictive hour fell from 9.991 to 9.947 km, and two of its hours now break the lactate ceiling, one of them because the controller is no longer told that the room is warm. The probe that found them is the one the shared section describes: move a hidden value with the measurements held, and ask whether the controller moves.

A lead measured where it was not earned. The runner app's texts said that almost the whole lead of the predictive controller was earned in the last five minutes. The paper agent, measuring every number the app asserts, sampled both runs once a second and found 215 of a final 406 m already won by three minutes, because the regulator walks a five-minute reference ramp up from rest. The same texts said the terminal value was what made the predictive controller work. Deleting it changed the flat hour by one meter. Its product form carried a kink from the fuel store, so on the under-fueled hour the belt alternated between 6.4 and 11.9 km/h and covered 8.33 km, behind the regulator. A smooth rollout fixed the alternation, and the chapter now reports the terminal value as measured, within 31 m of none at all. In the same app a claim that the sustainable pace falls through the hour was true in direction and 0.5 % in size on the hour it was attached to.

Protective layers that leaked. The anesthesia app's signal-quality hold froze the reading through cautery but let the monitor's filter go on integrating the artifact. On release the controller received a reading 24.7 index units above the patient, and its propofol target stepped from 1.81 to 2.83 µg/mL just after the cautery it was meant to ignore. The predictive-control round found it because one comparison was dominated by the minute after the hold released. The antibody culture's predictive controller, as first built, ran its pump on the declared stock and not the label, and its plan was not offset-free, so a weaker lot took the glucose under 1 g/L for 22.4 h. A technical review found these and a third fault, and Table 8.8 reports the controller before and after. The cyclist's predictive controller never started the bicycle on level ground, since its predictor turned zero cadence into zero power. The circuit hid the defect, because its start rolls downhill.

Causes the runs did not support. The book round tested each chapter's causal sentences by run and corrected six in three chapters without adding a number. The anesthesia paper said the end of a case is shortened by anticipation rather than by tuning. Twice the depth gain cut the time below 45 from 12.95 to 4.35 min, about what the opioid cut gives, while the lowest index stayed at 35.5. Only anticipation raises the bottom of the dip, and the chapter now says so. The kidney circuit's paper said no tuning helps a loop that follows a drifting scale. A slower loop left the patient 244.8 mL behind at six hours instead of 304.3, and caught up with the ramp by twelve. The brewery paper blamed an undershoot partly on the glycol left in the jacket. Runs over the probe lag moved the minimum by 0.18 °C and runs over the glycol holdup by 0.02 °C.

The book round. The thirteen chapter records list 177 rows of numeric verification, each a stated value or group of values set beside its reproduction, and all thirteen report no drift. The writer of Chapter 1 read the apps' action flags and found five chapters that named a loop's action against the convention of Section 1.3. Its own first outer loop for P-1 compared the measured effect site with the setpoint, so it integrated the effect-site lag and made the population patient overshoot. Comparing with the model's prediction made the loop act on model error only, which is the lesson of Chapter 3. Chapter 1 was later corrected from three chapter passes, on the order of an override, the meaning of a lag and the notation. Seven of the thirteen chapter passes logged an app defect, most of them stale numbers in briefings. One was behavior: at the smallest glycol holdup its slider allows, the brewery's integration step is too long and the run breaks down once a valve opens. One chapter was finished in two runs after its first agent was cut off, and the second agent checked the first one's editorial pass before going on.

Every error above was caught by a check that could fail: a hidden value moved while the measurements were held, a lead decomposed over time, a term deleted to see what it was worth, a protective layer released on purpose, a causal sentence tested against the variable it names. None was caught by rereading the text that contained it. Four working rules follow. Give a controller only what the plant measures and the design declares. Decompose every claimed advantage before explaining it. Test each protective layer on the event it guards and on its release. Run every "no tuning helps" with two tunings.

15.5 Gaps in the Control of Living Systems

Five gaps belong to this volume's topic, and the gaps common to agent-era practice follow them in Section 15.5.1. 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 15.4 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. *Variability that no measurement reveals.* Exercise multiplies insulin sensitivity by 1.6 in Chapter 2 with no change in any measured quantity, and the hidden inhibitor of Chapter 11 misleads a forecast that has no state for it. A review of closed-loop insulin systems finds that excursions after meals remain a challenge because insulin under the skin is absorbed with a delay, that users usually announce exercise to the algorithm in advance, and that heart rate or accelerometry may detect activity sooner than the glucose monitor alone [15]. Full-scale digesters are often run conservatively, in part because of limited observability [13]. The gap is hard because an unmeasured parameter shows itself only in the harm it does. Practice needs estimators of an individual's unmeasured parameters that stay safe while they learn, and a way to bring signals such as activity or a second analyzer into them.

2. *Virtual populations with stated credibility.* Every chapter runs one declared individual: a patient with one covariate set and a slider in Chapter 3, one virtual subject in Chapter 2, a scripted lung in Chapter 4 and a culture with teaching constants in Chapter 8. A margin sized for the spread of individuals needs that spread, drawn from data, with a statement of the context of use in which the model may be trusted. The Open challenge box after this list takes the gap further.

3. *Margins for one-way actuators.* No standard method sizes the margin a one-way pump needs in front of a hard limit from the spread of individuals and the slow way back, and the chapters set it by run. In Chapter 2 the pump sat at zero for 0.84 h while the insulin already given went on acting. The anesthesia dip at the end of a case is raised only by feedforward from an event nobody instruments, and the fish farm's governor cannot stop the nitrite of meals already eaten. Insulin given under the skin acts late, because it must first be absorbed, with peaks up to 120 min after a bolus of regular insulin [14]. The gap is hard because the margin must be set before the

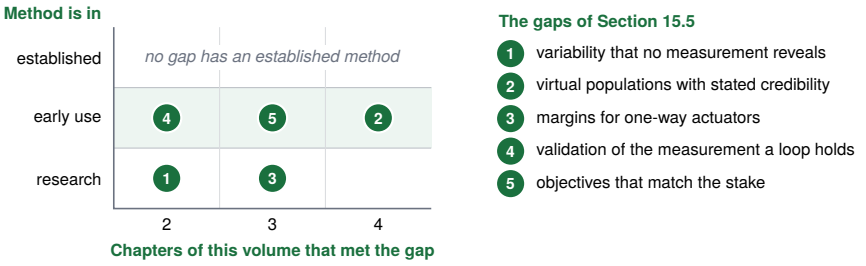


Figure 15.4 The five gaps of this section. Across is the number of chapters that each item names as evidence: two for gap 1 (Chapters 2 and 11), four for gap 2 (2, 3, 4 and 8), three for gap 3 (2, 3 and 12), two for gap 4 (3 and 5) and three for gap 5 (2, 9 and 11). Up is the maturity of the method available, the editor's reading of the sources cited in each item and in the Open challenge box. Asymmetric targets are agreed for glucose alone, so gap 5 sits in early use. Notice that the two gaps about what a loop cannot see or cannot take back, gaps 1 and 3, are still research.

individual is known, and it is paid for on every ordinary day. A method that sizes the margin from the spread of individuals and the slow way back, and states the risk it leaves, would replace setting it by run.

4. *Validation of the measurement a loop holds.* The kidney circuit's check missed every drift under its tolerance, and the monitor that is wrong during cautery reads like a light patient. Weighing is the most common way to balance fluid in continuous kidney therapy and the most reliable over long treatments, and its disadvantages include limited scale capacity, user errors and disturbances of the operating environment [19]. Muscle activity and electrical interference can pass into a depth index without being displayed as artifact, and raise its value spuriously [20]. The gap is hard because a fault that drifts slowly looks like the plant moving, and a tighter check raises more false alarms. Plausibility checks with stated detection limits, and redundant measurements designed into the equipment, are needed wherever a loop can hold a measurement while the plant moves.

5. *Objectives that match the stake.* A symmetric error, a vessel's yield and a day's revenue each chose the wrong controller in this volume. An international consensus on glucose monitoring already scores the two sides of the range apart, with targets of more than 70 % of readings in range, fewer than 25 % above 180 mg/dL and fewer than 4 % below 70 mg/dL, and it names raising the time in range while reducing the time below it as the primary goal [21]. The gap is hard because the price of each side is set outside control engineering, and most of these plants have no agreed price. Agreed asymmetric measures, campaign-level scores and a monitored margin to a fold, reported beside every result, would make a comparison mean what a clinician, a plant manager or a public-health officer needs it to mean.

⚠ **OPEN CHALLENGE** **A Population, Not One Individual**

Four chapters name gap 2, and every chapter runs one declared individual. A simulator of 300 virtual subjects with type 1 diabetes, built from real individuals' data to span the observed variability, was accepted by the Food and Drug Administration as a substitute for animal trials in the preclinical testing of closed-loop control [22]. The commercial models used to dose anesthetic drugs were built on selected populations that often left out children, obese and elderly patients [23]. In the editor's view the gap is hard because the spread that sizes a margin lies in the individuals that studies leave out. Progress would be a population drawn from pooled data and stated with the context in which it may be trusted.

15.5.1 Gaps Common to Agent-Era Practice

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

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

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

3. *Safety cases and standards for agents in commissioning and operation.* Functional safety standards define a lifecycle of specified requirements, verification and assigned human roles [18]. In the editor's view, how an agent that writes configuration, runs commissioning checks or advises an operator is qualified, limited and audited under them is not yet settled.

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

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

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

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

has. Deliberate practice, of the kind the exercises in this book provide, has to be designed into education and into work.

KEY POINTS

- A loop coped with the differences between individuals when it could measure the effect it was meant to hold.
- A loop that holds its measurement need not hold the plant.
- One-way actuators were held by margin and by acting early.
- The objective decided what counted as held.
- Prediction paid where it had a deadline and an anchored estimate.

References

[1] J. L. Sherr, L. Heinemann, G. A. Fleming, R. M. Bergenstal, D. Bruttomesso, H. Hanaire, R. W. Holl, J. R. Petrie, A. L. Peters, and M. Evans, "Automated insulin delivery: Benefits, challenges, and recommendations. A consensus report of the joint diabetes technology working group of the European association for the study of diabetes and the American diabetes association," *Diabetes Care*, vol. 45, no. 12, pp. 3058–3074, 2022.

[2] L. Pasin, P. Nardelli, M. Pintaudi, M. Greco, M. Zambon, L. Cabrini, and A. Zangrillo, "Closed-loop delivery systems versus manually controlled administration of total IV anesthesia: A meta-analysis of randomized clinical trials," *Anesthesia & Analgesia*, vol. 124, no. 2, pp. 456–464, 2017.

[3] M. Ghita, M. Neckebroek, C. Muresan, and D. Copot, "Closed-loop control of anesthesia: Survey on actual trends, challenges and perspectives," *IEEE Access*, vol. 8, pp. 206 264–206 279, 2020.

[4] P. von Platen, A. Pomprapa, B. Lachmann, and S. Leonhardt, "The dawn of physiological closed-loop ventilation—a review," *Critical Care*, vol. 24, no. 1, p. 121, 2020.

[5] M. Viceconti, F. Pappalardo, B. Rodriguez, M. Horner, J. Bischoff, and F. Musuamba Tshinanu, "In silico trials: Verification, validation and uncertainty quantification of predictive models used in the regulatory evaluation of biomedical products," *Methods*, vol. 185, pp. 120–127, 2021.

[6] C. Ronco, "Continuous renal replacement therapy: Forty-year anniversary," *The International Journal of Artificial Organs*, vol. 40, no. 6, pp. 257–264, 2017.

[7] H. Narayanan, M. F. Luna, M. von Stosch, M. N. Cruz Bournazou, G. Polotti, M. Morbidelli, A. Butté, and M. Sokolov, "Bioprocessing in the digital age: The role of process models," *Biotechnology Journal*, vol. 15, no. 1, p. 1900172, 2020.

[8] K. A. Esmonde-White, M. Cuellar, C. Uerpmann, B. Lenain, and I. R. Lewis, "Raman spectroscopy as a process analytical technology for pharmaceutical manufacturing and bioprocessing," *Analytical and Bioanalytical Chemistry*, vol. 409, no. 3, pp. 637–649, 2017.

[9] S. Hemming, F. de Zwart, A. Elings, A. Petropoulou, and I. Righini, "Cherry tomato production in intelligent greenhouses—sensors and AI for control of climate, irrigation, crop yield, and quality," *Sensors*, vol. 20, no. 22, p. 6430, 2020.

[10] K. M. Gostic, L. McGough, E. B. Baskerville, S. Abbott, K. Joshi, C. Tedijanto, R. Kahn, R. Niehus, J. A. Hay, P. M. De Salazar, J. Hellewell, S. Meakin, J. D. Munday, N. I. Bosse, K. Sherrat, R. N. Thompson, L. F. White, J. S. Huisman, J. Scire, S. Bonhoeffer, T. Stadler, J. Wallinga, S. Funk, M. Lipsitch, and S. Cobey, "Practical considerations for measuring the effective reproductive number, R_t," *PLOS Computational Biology*, vol. 16, no. 12, p. e1008409, 2020.

[11] H. Yu, B. Gao, and Y. Li, "Depth of anesthesia monitoring: Current evidence, clinical impact, and future directions," *Journal of International Medical Research*, vol. 54, no. 6, p. 03000605261457781, 2026.

[12] M. Marycz, I. Turowska, S. Glazik, and P. Jasiński, "Artificial intelligence in anaerobic digestion: A review of sensors, modeling approaches, and optimization strategies," *Sensors*, vol. 25, no. 22, p. 6961, 2025.

[13] A. I. Yunus, S. Jalla, J. F. Bozeman, and Y. Chen, "Reinforcement-learning-enabled supervisory control of a full-scale anaerobic digestion facility using a calibrated digital twin," *Environmental Science & Technology*, vol. 60, no. 39, pp. 28 064–28 078, 2026.

[14] C. Cobelli, E. Renard, and B. Kovatchev, "Artificial pancreas: Past, present, future," *Diabetes*, vol. 60, no. 11, pp. 2672–2682, 2011.

[15] C. K. Boughton and R. Hovorka, "New closed-loop insulin systems," *Diabetologia*, vol. 64, no. 5, pp. 1007–1015, 2021.

[16] S. Hemming, F. de Zwart, A. Elings, I. Righini, and A. Petropoulou, "Remote control of greenhouse vegetable production with artificial intelligence—greenhouse climate, irrigation, and crop production," *Sensors*, vol. 19, no. 8, p. 1807, 2019.

[17] 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.

[18] 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.

[19] G. Villa, M. Neri, R. Bellomo, J. Cerda, A. R. De Gaudio, S. De Rosa, F. Garzotto, P. M. Honore, J. Kellum, A. Lorenzin, D. Payen, Z. Ricci, S. Samoni, J.-L. Vincent, J. Wendon, M. Zaccaria, and C. Ronco, "Nomenclature for renal replacement therapy and blood purification techniques in critically ill patients: practical applications," *Critical Care*, vol. 20, p. 283, 2016.

[20] A. A. Dahaba, "Different conditions that could result in the bispectral index indicating an incorrect hypnotic state," *Anesthesia & Analgesia*, vol. 101, no. 3, pp. 765–773, 2005.

[21] T. Battelino, T. Danne, R. M. Bergenstal, S. A. Amiel, R. Beck, T. Biester, E. Bosi, B. A. Buckingham, W. T. Cefalu, K. L. Close, C. Cobelli, E. Dassau, J. H. DeVries, K. C. Donaghue, K. Dovc, F. J. Doyle, III, S. Garg, G. Grunberger, S. Heller, L. Heinemann, I. B. Hirsch, R. Hovorka, W. Jia, O. Kordonouri, B. Kovatchev, A. Kowalski, L. Laffel, B. Levine, A. Mayorov, C. Mathieu, H. R. Murphy, R. Nimri, K. Nørgaard, C. G. Parkin, E. Renard, D. Rodbard, B. Saboo, D. Schatz, K. Stoner, T. Urakami, S. A. Weinzimer, and M. Phillip, "Clinical targets for continuous glucose monitoring data interpretation: Recommendations from the international consensus on time in range," *Diabetes Care*, vol. 42, no. 8, pp. 1593–1603, 2019.

[22] B. P. Kovatchev, M. Breton, C. Dalla Man, and C. Cobelli, "In silico preclinical trials: A proof of concept in closed-loop control of type 1 diabetes," *Journal of Diabetes Science and Technology*, vol. 3, no. 1, pp. 44–55, 2009.

[23] O. Vandemoortele, L. N. Hannivoort, F. Vanhoorebeeck, M. M. R. F. Struys, and H. E. M. Vereecke, "General purpose pharmacokinetic-pharmacodynamic models for target-controlled infusion of anaesthetic drugs: A narrative review," *Journal of Clinical Medicine*, vol. 11, no. 9, p. 2487, 2022.

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

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

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

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

15.6 Exercises

15.1 Discussion · 8 min

An agent proposes to remove the depth loop of Chapter 3 and give the target-controlled pump a richer population model with more covariates, "so that the monitor, which fails during cautery, is no longer needed". Using Table 3.7 and scenario 9 of the chapter, say what the proposal would gain and what it would lose.

15.2 Discussion · 8 min

The balance loop of Chapter 5 keeps its measured error within 28.4 mL while the patient falls 304.3 mL behind in six hours. Explain why retuning the loop cannot fix this, and discuss how you would choose the tolerance of the scale check of Table 5.4.

15.3 Discussion · 7 min

In Chapter 2 exercise takes glucose from 74.9 to 47.8 mg/dL before the controller next executes, and the pump then sits at zero for 0.84 h. A colleague proposes a higher gain "so the loop reacts to lows faster". Discuss the proposal, and say what does protect a one-way loop from a fast fall.

15.4 Discussion · 9 min

An agent reports that its predictive controller covers 9.947 km in the runner's hour of Chapter 6 against 9.584 km for the selector, "because it plans the whole hour". Using Tables 6.8 to 6.10 and Section 15.4.1, list the checks you would run before accepting the claim and its explanation.

15.5 Discussion · 8 min

Three chapters chose a different controller depending on the score: Chapter 2 with a symmetric IAE, Chapter 9 with the fold of one lot, and Chapter 13 with the cost of the week. Pick one of them, state what the score hid, and propose a score that a person responsible for that plant would accept.

15.6 Discussion · 8 min

On the epidemic challenge of Chapter 14, the nominal predictive controller with a 20 % margin passes all seven records for less than the scenario tree, and the chapter still prefers the tree. Using Table 14.5 and the gaps of Section 15.5, argue for one of the two as the policy to recommend, and say what evidence would change your choice.