INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 20 Motion Control, Robotics, and Autonomous Systems

CHAPTER 16

Summary and Outlook

Pages 380–394

Edited by John D. Hedengren

How to cite this chapter. "Summary and Outlook," in *Motion Control, Robotics, and Autonomous Systems*, John D. Hedengren, Ed., vol. 20 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 16, pp. 380–394.

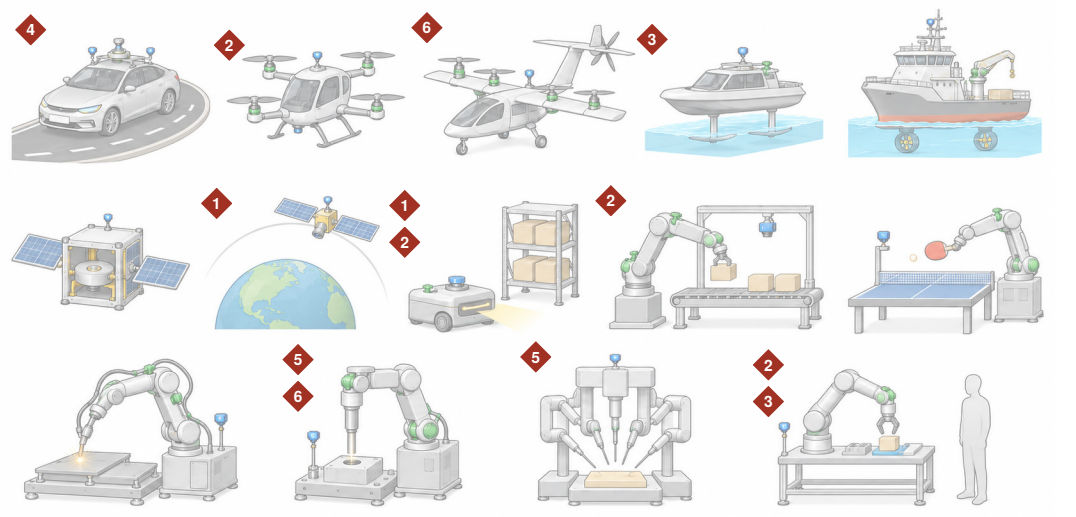
The page numbers in this file are those of the printed volume.

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v20-motion-robotics/

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

CHAPTER 16

Summary and Outlook



Open gaps 1 trusted uncertainty 2 verified told values 3 proof of infeasibility 4 model mismatch 5 soft contact 6 event-keyed benchmarks

Notice that the pallet robot, the peg and the cobot carry two gaps each, and that no gap marks the vessel, the small satellite, the table tennis arm or the welding arm. Illustration, not to scale.

IN PLAIN TERMS **What Fourteen Moving Machines Taught**

The main idea. Each chapter of this book let a machine move on its own: a car, an aircraft, a boat, a satellite, a robot in a warehouse and arms that pick, weld and operate. None could see all it needed, so each acted on what it believed.

The challenge. A number written down in advance can be wrong. A guess takes time to improve, and some choices must be made before the result is known. When a machine cannot stop in the space it can see, the safe answer is to slow down or not go.

Why it matters. The machines that worked were told what was checked, kept a margin and were judged against the truth.

Fourteen chapters put a moving machine to work on a task it could not see in full: a car, an air taxi, an aircraft handing its weight to a wing, a foiling boat, a vessel on station, two satellites, a pallet robot and six arms. Section 1.4.10 named four ways a controller closes the gap between what it believes and what is true. It can be told, it can estimate, it can bound the unknown and keep a margin, or it can commit to a prediction before the outcome is known. This chapter adds the measured results to that reading and compares how each system completed its task under uncertainty. 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.

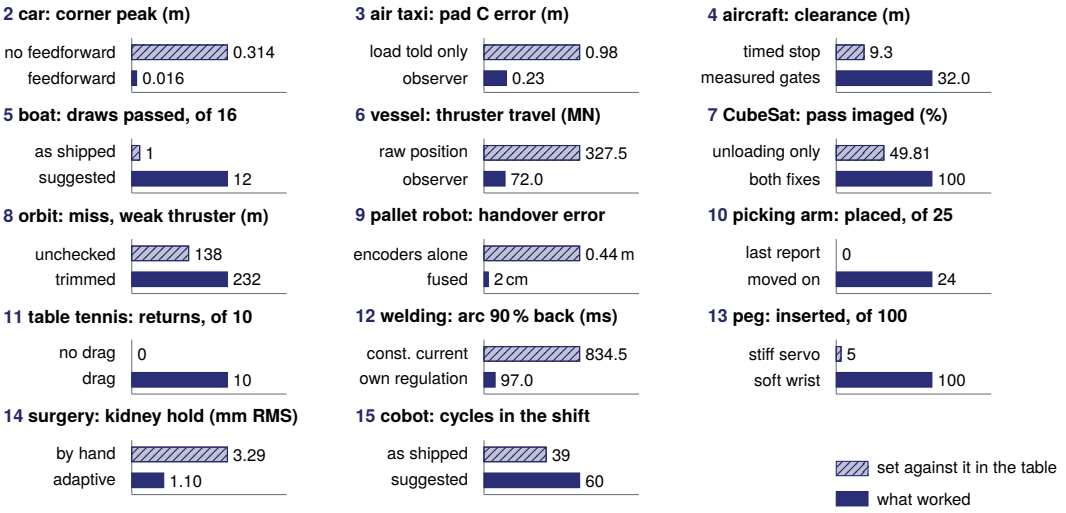


Figure 16.1 One pair of results from each chapter, with the numbers as the "Measured result" column of Tables 16.1 and 16.2 prints them. Each panel has its own scale from zero. A larger value is better in the panels of Chapters 4, 5, 7, 8, 10, 11, 13 and 15, and a smaller one in the others. The pallet robot's pair is printed in two units and drawn on one axis in meters, and a result of zero draws no bar. The cobot's suggested shift also ends with no violation against 47. Notice that in most rows the better result came from a change of structure, of what the controller was told or of what it estimated, not from tuning alone.

16.1 What the Chapters Showed

Tables 16.1 and 16.2 extend Table 1.2. Each row names what was uncertain, a structure that completed the task, one that did not, and the measured result with the table or section of the chapter that reports it. Every number is the chapter's own, and every result is scored against the true state that the app logs beside the controller's estimate. Figure 16.1 draws one pair of results from each row.

The columns of the two tables make five points. The first four follow the four ways of Section 1.4.10.

Telling worked when the sheet was right. A value given to the controller beforehand acted at once. The air taxi told its load climbed 0.05 m behind where fixed tuning fell 1.78 m behind, the picking arm told its payload made no hits where the arm model alone made 11, and curvature feedforward cut the corner peak of Chapter 2 nearly twentyfold. A told value carries no error signal of its own, so a wrong one is believed in full. The pallet robot planned with its empty braking hit a cart at 0.74 m/s despite an emergency stop. The cobot planned with a brochure figure gained no cycle and produced six violations, because the loaded arm took 640 ms rather than 200 ms to stop (Section 15.4).

Estimating worked when the estimator had a state for the unknown. The vessel's observer separated wave motion from drift, and a vessel model with a bias estimate but no wave states already captured 92 % of the reduction in thruster travel (Section 6.6). The satellite's filter with a bias state held the camera to 0.054° through the Sun's blind window, where the same filter without it drifted 2.35°. Only an estimator with a spin state read a 40 rps backspin before the bounce in Chapter 11 (Table 11.4). An estimate takes time to form, and the chapters measured it. After the squall of Chapter 6 the bias estimate covered 63 % of the change in 237 s, and a gain lowered for a heavy sea took 17.35 instead of 5.05 s to notice a real step.

Table 16.1 How Each Machine Completed Its Task Under Uncertainty, Chapters 2 to 8

Chapter	Uncertain	What worked	What failed	Measured result
2, lane keeping	Speed, curvature, a misaligned linkage, the rack's lag.	A lateral gain divided by speed, with curvature feedforward. A predictive controller against the misalignment.	No curvature feedforward. The predictive controller on a 0.45 s rack that its 0.10 s model misjudges.	Corner peak 0.314 against 0.016 m (Table 2.5). Recovery to $1/e$ in 0.84 and 0.85 s at 13.0 and 30.6 m/s (Section 2.6). IAE 0.326 against 8.850 m s on the misalignment (Table 2.7), unstable above 26.4 m/s on the slow rack (Section 2.6.7).
3, air taxi	Load, balance, wind among buildings.	Load-sheet feedforward. A disturbance observer with a 0.67 s time constant.	Fixed tuning with a heavy load. Feedforward without the observer at the gap's jet.	Heavy climb 1.78 m behind with fixed tuning, 0.05 m when told (Table 3.7). Pad C error 0.99, 0.98 and 0.23 m, only the observer passes (Table 3.8), on 20 of 20 seeds against 5 for each other mode (Section 3.6).
4, lift-plus-cruise aircraft	Airspeed, lift margin, mass, a failed motor or pitot.	Rotor stop gated on measured airspeed and lift margin. An airspeed derived from lift beside the pitot.	A timed schedule tuned for 250 kg in calm air. The pitot alone.	Stop at 37.9 against 55.7 m/s, clearance 9.3 against 32.0 m (Table 4.8). Unchecked pitot: into the wood at 38.5 s. A failed motor leaves 4.8 % of margin: divert (Table 4.7).
5, foiling boat	Waves and sensor loss, inside a band 0.80 m tall between slam and ventilation.	Contouring. A cascade on the measured heave rate. Slowing off the foils in a sea too large.	Platforming in a 0.5 m following sea. Every setting tried in a 1.4 m head sea.	Cascade contouring clean on 17 of 17 draws, the PID's on 7 (Section 5.6). Challenge passed on 1 of 16 draws as shipped, 12 as suggested (Table 5.6). None of 33 795 runs flew the head sea past 12.49 s.
6, vessel on thrusters	Wave motion inside the measurement, the environmental force, a heading the job leaves free.	An observer with bias and wave states, bias feedforward, the minimum-thrust heading.	The raw position handed to the loops. A slow integral in place of the feedforward.	Thruster travel 327.5 against 72.0 MN for a mean offset of 0.55 against 0.51 m (Table 6.4). Heading: 2297 to 738 kW. Squall without feedforward: 7.7 m off, past the disconnect radius with the waves (Section 6.6).
7, CubeSat attitude	Gyro bias, stored momentum, a failed wheel.	A filter with a bias state together with magnetic unloading. The eigenaxis slew.	Either fix alone. Any tuning tried after a wheel fails.	Pass imaged 24.12, 49.81, 20.64 and 100 %: shipped, unloading only, bias state only, both (Table 7.6). Blind-window error 2.35 against 0.054°. Failed wheel 81.56 to 84.75 % (Table 7.5).
8, satellite in low orbit	Drag, the debris covariance, thruster output.	A predictive planner on the mean altitude. Collision probability, not miss distance. A burn checked and trimmed.	A deadband deciding every 6 h. A plan that trusts a weak thruster.	Storm: 6.08 h below the band against none. A pass predicted 1.6 km away carries $P_c = 1.8 \times 10^{-4}$. Weak thruster: 138 m unchecked, 232 m trimmed (Table 8.7). Campaign: 3 unresolved passes as shipped, 0 suggested (Table 8.8).

Table 16.2 How Each Machine Completed Its Task Under Uncertainty, Chapters 9 to 15

Chapter	Uncertain	What worked	What failed	Measured result
9, pallet robot	Position, braking with a load, space it cannot see.	A fused estimate, braking planned for the payload, slowing for unseen space.	The empty robot's braking with a pallet on. Encoders alone. Unseen space treated as clear.	Cart hit at 0.74 m/s against a stop 0.25 m short. Handover 0.44 m past against 2 cm. Worker hit at 0.72 m/s against 0.92 m kept. Shift fails at 11.1 s, passes 3 of 3 in 120.6 s (Table 9.6).
10, picking arm	Report age, belt speed, payload.	Each report moved on by the encoder travel since its stamp. Meeting the product at belt speed. The payload in the arm model.	Aiming at the last report, at any latency. Waiting at rest. The arm model alone.	Placed 24 of 25 against 0. At rest 0 placed and 60 failed, moving with it 30 placed. Arm model alone: 11 hits, torques 47.1 and 24.3 N m. Challenge 0 to 98 placed (Table 10.4).
11, table tennis arm	Drag, spin, the time left to swing.	A flight model with drag, a commit 200 to 275 ms after the serve, a minimum-jerk swing, a filter with a spin state.	A predictor without drag. Commits too early or too late. Stepped commands. A slow camera.	Returns 0 and 10 of 10 without and with drag. Commits at 50, 200 and 450 ms: 16, 19 and 11 of 20. Slow camera: 0 of 10 in all 144 designs (Table 11.4). Match 0–11 shipped, 11–9 suggested (Table 11.5).
12, arc welding cell	Torch height, a seam that wanders ahead of the arc.	The arc's own regulation. The current read as a height. A preview filed by distance.	A constant-current arc. A look-ahead configured 20 mm short.	After a 3 mm step the arc is 90 % regrown in 97.0 ms, at constant current in 834.5 ms. Short look-ahead: 1.451 mm against 1.455 mm predicted, 372 mm unfused (Section 12.6). Graded seam 1198 to 0 mm defective (Table 12.6).
13, peg in a hole	Hole position, friction, a late force signal.	Compliance: a 3 N/mm spring, a soft wrist, contact detection and a search.	A 200 N/mm servo against a 0.4 mm offset. A force loop behind 80 ms of delay.	Damaged at 136.1 N against inserted at 13.7 N. At 80 ms: 42.2 N true, 13.4 N read. Of 100 parts: stiff 5, impedance 66, sequence 87, sequence with a soft wrist 100 at 30.2 N (Table 13.5).
14, surgical instrument	Tremor, delay, organ motion, tissue creep.	Motion scaled 3:1 with a 6 Hz filter. An adaptive harmonic predictor.	A 2 Hz filter. A force hold at 2.5 N on creeping tissue.	Artery contacts 6, 0 and 3 for 1:1 at 30 Hz, 3:1 at 6 and 3:1 at 2 Hz. Kidney 3.29 by hand, 1.10 mm RMS adaptive. Tear at 13.3 s (Table 14.5). Delay times motion 86 % of the error, 91 % of it removed (Table 14.6).
15, collaborative arm	The person's motion, latency, braking with a load.	A separation budget on the true braking. Restart hysteresis. Power and force limiting at the chest.	A brochure braking figure. No hysteresis. Any speed at the table.	Brochure: 6 violations, the worst 506 mm, against none. Stops 17 against 2 in 90 s (Table 15.5). Shift 39 cycles with 47 violations against 60 with none (Table 15.6).

Bounding cost time, and sometimes the bound said do not go. Slowing for space the robot cannot see kept 0.92 m from a worker at a cost of 8.7 s. The cobot's speed law cost 4 cycles in a minute of a person coming and going, and every 100 ms of latency from 100 to 400 ms cost 45 to 60 cycles an hour. Where the margin could not be met, the right answer was to decline the task. The aircraft with a failed motor diverted. The foiling boat in a 1.4 m head sea was safest kept off its foils. At the cobot's table the protective distance at rest, 532 mm, exceeds the 214 mm the path leaves a person, so no speed override works there.

Committing traded certainty for time. The table tennis arm returned 16 or 17 of 20 serves when it committed early, 19 at 200 to 275 ms, and 10 to 12 from 350 ms on, when the swing no longer fit although the prediction was sharper (Section 11.4). The satellite in low orbit paid 1.20 cm/s to clear a pass with 8.8 h of lead, 2.10 with 5.8 h and 5.08 with less than an orbit (Section 8.7). In Chapter 14 the interaction of the surgeon's delay with the organ's motion held 86 % of the error of a hold, and an adaptive predictor of the breathing removed 91 % of it.

Structure decided before tuning, and the verdict belonged to the truth. In Chapter 6 the three PID loops accounted for almost none of the difference between a good and a poor system, and the heading cut the thruster power by 68 % with nothing in the control system changed. The cascade on heave rate flew 17 draws of a sea that the PID flew 7 times, and the satellite reached a full pass only with both fixes. Every app scored against a truth the controller never saw, and the gap was often silent. The pallet robot handed its load 0.44 m past the station with the truth outside its own ellipse. The vessel at a low observer gain wandered to 4.14 m while the trace reported that all was well.

16.2 Choosing a Method

The chapters suggest an order of questions for a new moving machine. Figure 16.2 draws the same questions as a flowchart.

1. *What does the task score, and against which truth?* Score the true state, and log it beside the estimate, as the audit of Chapter 15 does.
2. *What can the controller be told, and how is the sheet checked?* Feed forward the load, the payload and the curvature (Chapter 3, Chapter 10, Chapter 2). Take a braking figure from a test of the loaded machine, not from a brochure (Chapter 9, Chapter 15).
3. *Is part of the measured motion real and not an error?* Give the estimator a state for it, as the wave states of Chapter 6 and the platform-or-contour choice of Chapter 5 do, and a bias state for any drifting sensor (Chapter 7).
4. *How old is each piece of information?* Move a report forward by the motion since its stamp (Chapter 10), key a preview by distance rather than time (Chapter 12), and predict across a delay that multiplies motion (Chapter 14).
5. *Can the machine stop or abort inside what it can see?* Set the speed from believed braking and visible space, gate a mode change on a measured margin (Chapter 4), and add hysteresis where a bound flickers (Chapter 15). Decide in advance what the machine does when the margin is not there.
6. *Must it commit before it knows?* Find the window in which the prediction is good enough and the action still fits (Chapter 11), and plan early when lead time is cheap (Chapter 8).
7. *Does contact or a shared actuator turn an error into a force or a limit?* Make the contact compliant before raising any gain (Chapter 13), and allocate among actuators with their stored momentum and free variables in view (Chapter 7, Chapter 6).

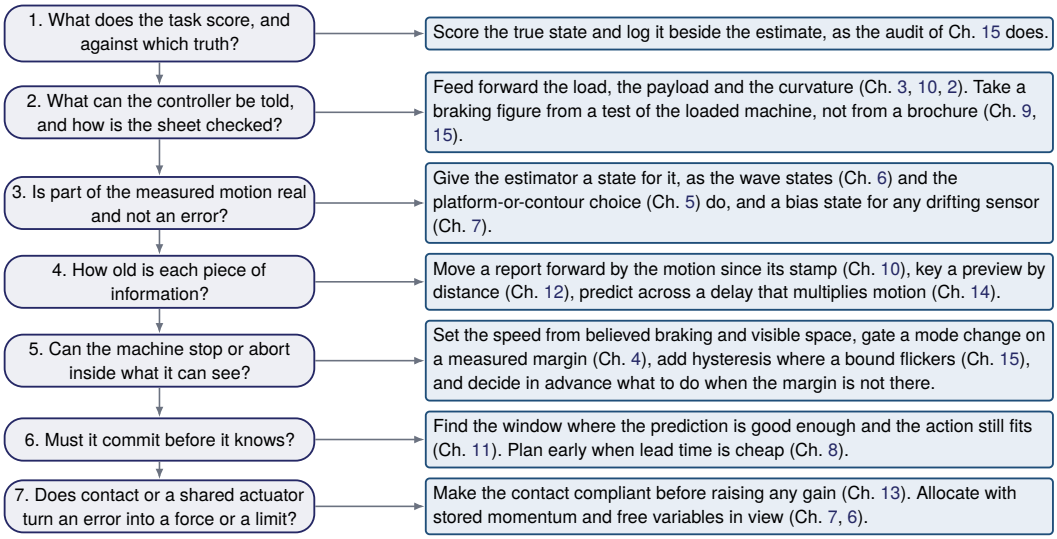


Figure 16.2 The seven questions of this section in order, with the list's answer and the chapters where it decided the result. Notice that the truth and the told values come first, and that the questions of timing and commitment follow the questions of what the machine knows.

A predictive controller earned its place where it carried a constraint or a model the simpler law lacked. In Chapter 2 it removed the standing error of a misaligned linkage, and in Chapter 8 it never let the mean altitude reach the band through a storm. Its gains were often smaller than they first looked. On the foiling boat it rode case 6's sea at 0.74 m/s^2 against the shipped cascade's 1.30, but a cascade retuned by the same rule rode at 0.90 and rode the 0.65 to 1.0 m following seas more smoothly (Table 5.5). The satellite's predictive law reached the next site in 39.6 against 52.1 s, a gain that came from staying inside the wheels' limits rather than from looking ahead (Section 7.6.11). Volume 9 treats that choice in full.

16.3 Where Practice Is Heading

Evidence that cannot come from driving alone. A statistical study found that demonstrating the safety of autonomous vehicles by test driving would take hundreds of millions, and in some cases hundreds of billions, of miles, and concluded that driving alone is not a practical way to assess it [1]. An analysis of autonomous vehicle safety names validating inductive learning against novel inputs and reaching the dependability of a full fleet as open challenges, and suggests that the largest may be a single design and deployment process that joins every specialty's safety concerns [2]. In the editor's view, evidence for moving machines will come increasingly from simulators scored against a logged truth, with seeds and disturbances published so that a result can be checked.

Learned controllers trained in simulation. An autonomous racing drone trained by deep reinforcement learning in simulation, combined with data collected in the physical world, won several races against each of three human champions and set the fastest recorded race time [3]. A survey of foundation models in robotics finds that models pretrained on internet-scale data may improve perception, decision-making and control, and names the scarcity of robot data, safety guarantees, uncertainty quantification and real-time execution as open challenges [4]. A review of safety filters, which monitor a controller's commands and intervene to keep the system safe, proposes one framework for them and argues that deployment needs model-based methods together with data-driven

ones, whose guarantees are less well understood [5]. The speed law of Chapter 15, the envelope of Chapter 9 and the gates of Chapter 4 are filters of that kind.

> ⧖ HISTORY **1999 to 2025: Gyros and Stars in a Small Cube**
>
> The first history box of Chapter 1 left the lunar spacecraft's gyros on a stable member, realigned in flight by sighting stars. A NASA survey of small spacecraft technology records what followed. In 1999 two universities developed the CubeSat, an educational platform built on a cube 10 cm on a side, and CubeSats are now a common form of small spacecraft that can weigh only a few kilograms. Three-axis control with satellite navigation had flown for decades on spacecraft of about 100 kg, and only in the past few years has it reached the smallest classes. Fiber-optic and micromachined gyros carry the attitude between star tracker updates, which typically arrive a few times a second [6]. In the editor's reading the arc ends where Chapter 7 begins: the same stars and the same gyro bias, in a cube.

Machines among people. A survey of human-robot collaboration in industry describes collaborative robots entering the market and names guaranteed safe interaction as the first challenge, ahead of intuitive interfaces for the operators who program them [7]. A review of autonomous mobile robots in intralogistics reports their introduction in manufacturing, warehousing, cross-docks, terminals and hospitals, with decisions moving from a central unit to the robots themselves [8]. In the editor's view the believed braking and the truth audit of Chapter 15 will matter more as such fleets share aisles with people.

Autonomy in the operating room and in orbit. A framework for regulatory, ethical and legal discussion identifies six levels of autonomy for medical robots [9]. A review of surgical autonomy places the bulk of commercial platforms at level 0, with no decision autonomy [10]. In orbit, collision avoidance has become an essential part of satellite operations. The European Space Agency released conjunction messages collected from 2015 to 2019 for a competition to predict the final collision risk of an encounter from its earlier messages [11]. The planner of Chapter 8 flies on that kind of prediction, and the editor's view is that both domains will keep a person in charge of the decision while software proposes it.

Weight measured on board. United States guidance on aircraft weight and balance lets an operator use an onboard system that measures the actual weight and the location of the center of gravity. Such an operator may not need to narrow the loading envelope for variation in passenger weight and seating, although it must still allow for the system's own tolerances and keep a load manifest [12]. The load sheet of Chapter 3 is the told value that such a system checks. In the editor's view the braking figures of Chapter 9 and Chapter 15 still wait for the same step, from an assumed value to one measured in service.

Learned models with physics inside. A framework for robotic welding joins physical laws to a model learned from data. It predicted weld instability with an accuracy of up to 98 % over the next 50 ms and 86 % up to 1 s ahead, and it reduced the dependence on costly data, which is scarce for unstable welds and for low-volume work on complex paths [13]. The cell of Chapter 12 carries a physical model of the arc and no learned one. In the editor's view a learned model that keeps the physics is easier to bound under the mismatch of gap 4 of Section 16.5 than one that does not.

Mobile robots tested with their loads. The published performance standards for autonomous industrial vehicles address navigation, obstacle description and detection, environmental conditions, vehicle configuration and terminology. In NIST experiments heavy payloads affected the vehicles when they accelerated, on slippery floors and at the transition from floor to ramp, and the report recommends that the capabilities standard combine ramps and other conditions with the payload the

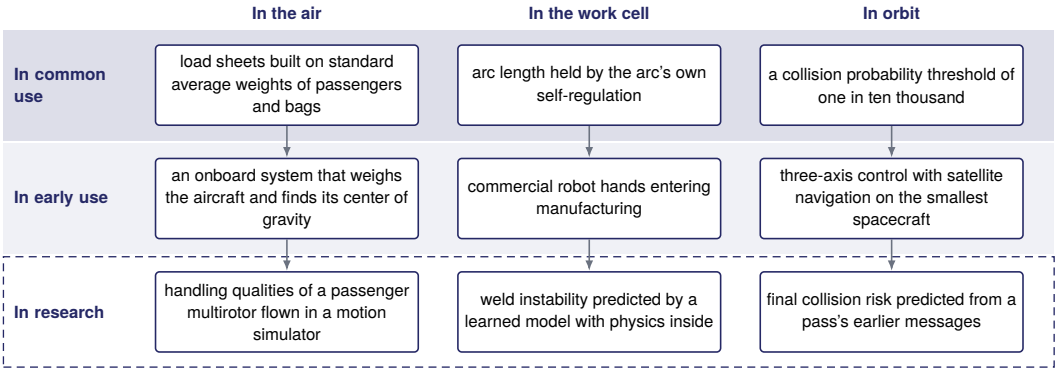


Figure 16.3 Developments in three threads. The top lane holds what the cited source calls standard, most common or widely accepted, the middle lane what it reports in service or entering use, and the bottom lane what it reports from a simulator, a laboratory or a competition on recorded data, in the editor's reading. Sources by thread: in the air [12, 15], in the work cell [13, 16, 17] and in orbit [6, 11, 18]. Notice that each thread moves from a value assumed or held by a rule toward one measured or predicted in service.

vehicle carries [14]. The braking that the pallet robot of Chapter 9 planned without its load is the same combination seen from the controller's side. Figure 16.3 sets developments of this section in three threads.

16.4 Engineering With AI Agents

An AI agent is a language model that works toward a goal in a loop. It plans, calls tools such as a compiler, a simulator or a browser, reads the results and continues [19]. 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 [20]. An agent may draft the hazard analysis or the test procedure. A person owns the decision to accept it, to start up and to keep running.

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

16.4.1 How This Volume Was Made

Each of the fourteen apps was built and tested by agents. An agent then wrote the paper on it, keeping verification notes that trace each number to a run. Six of the fourteen referee reports call themselves self-reviews, and the other eight are written in a referee's voice. None names a reviewer independent of the writing. The defect logs hold 133 entries with a severity. Of these, 31 are wrong output or behavior, one of them marked a blocker, 85 are documentation or content, 15 cosmetic and 2 something else, a test and a check allowance. Some logs remove repaired entries or list a round's 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 controller or a comparison met uncertainty in a way its text did not say.

Controllers that were told the truth. Two logs record a controller that read the run's hidden configuration. The satellite's pointing laws took the inertia of the run and its filters the tracker noise of the run. With the moments of inertia raised two and a half to three times, a slew still settled as the declared spacecraft does, and with a tracker twenty times noisier than declared the camera still imaged the whole pass. The foiling boat's flap map was trimmed on the run's mass, passenger and weed, so its trim flap moved from 0.496° to −0.027° with a passenger moved aft and to 2.097° at 1900 kg, before any measurement had changed. In a volume about uncertainty, a controller that reads the true plant has none to handle. The boat's was found by the paper agent of a later round and the satellite's by an audit of what each controller reads. Both were measured the same way: move a hidden value before Run and ask whether the controller moves while the measurements do not. No shipped run moved when the leaks were closed, 19 of 19 and 17 of 17 hashed sets, so only a learner who changed a value could have met them. With the leak closed, the satellite images 78.0 % of the pass with the noisier tracker.

An advantage with two causes. The vessel's texts credited the whole of a 1.34 against 0.63 m gap in mean offset to the raw measurement chasing the sea. Reading the controller against its own comments showed that the raw setting also switched the environment feedforward off. A copy of the controller with that condition removed put about 0.4 m of the 0.71 m gap on the missing feedforward, with the travel unaffected. The same consistency pass found that every automatic scenario started

with its bias estimate and integrals at zero, so the vessel drove off station and left the watch circle at 91.6 s in scenario 4. After both repairs the offsets are 0.55 and 0.51 m, and the chapter rests the case for the wave filter on the travel, 327.5 against 72.0 MN, where it belongs.

Comparisons that were not like for like. The peg-in-hole challenge drew each part's offset and tilt from the same random stream as the sensor noise, so the number of noise samples before part k depended on the controller. Between the stiff and the suggested runs 1 of 100 parts was the same. The paper agent traced the draw and gave the parts a generator of their own, and ten further seeds confirmed the ranking. In the picking cell the draft compared grasp errors of 84.7 and 55.0 mm at latencies of 300 and 0 ms, which would make the latency worth 29.7 mm where the belt predicts 45 mm. That mismatch led the self-review to the code, where the figure was the distance to the nearest product, not to the one the arm was sent for.

Mode changes that stranded a plan. In Chapter 8 a reboost committed before a switch to manual, whose window passed in manual, never fired and blocked every later reboost. Left alone for three days the satellite spent 61.16 h outside its band, against 0.53 h in automatic. The fix closes such a burn as expired and plans again. The book round then found a second path to the same state: pressed a third time while earlier burns are pending, the manual burn button queues a burn that never starts and never expires, so ten presses fly 10 cm/s instead of 50. That defect has since been repaired, and pending manual burns now fly in full.

The book round. The fourteen chapter records list 202 rows of numeric verification, each a stated value or group of values set beside its reproduction, and all fourteen report no drift. The writer of Chapter 1 read each app's action setting and found eleven sentences in eight chapters that named a loop's action against the convention of Section 1.3. Two causal claims of the papers were corrected by run. The surgical force loop holds its estimate, not the force, and the true force dips below what the dissection needs. The cobot's chatter persists with the scan error at its smallest setting. The pass on Chapter 4 found that its graded challenge passes on the timed schedule once the altitude integral time is 60 s, because the climb is slower and the gust, fixed at 64 s, passes before the timed stop. The app now times the gust from the start of the transition, and the chapter's review exercise shows the timed schedule failing again with that integral time. Four chapter passes led to five corrections of Chapter 1. Ten of the fourteen passes logged an app defect, most of them stale or rounded numbers in briefings or code comments.

Every error above was caught by a check that could fail: a hidden value moved with the measurements held, a counterfactual copy of the code, a difference that did not match the belt speed, a mode switched in the middle of a plan, a causal claim tested on the variable it names. None was caught by rereading the text that contained it. Four working rules follow. Give the controller the declared plant and never the run's. Apportion every advantage by counterfactual before explaining it. Compare designs on the same draws, with disturbances timed to the event they test. Exercise every mode change a user can make, in every order.

16.5 Gaps in Motion Control and Autonomy

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

1. *Uncertainty a safety layer can trust.* The pallet robot handed its load over with the truth outside its own ellipse, and the satellite in low orbit carried a probability set by an ellipse 971 by 150 m.

Speed laws and collision thresholds consume these bounds as if they were right. A handbook of conjunction assessment notes that probability methods almost always need a covariance for each object. The catalog's covariances have had twenty years of work on their realism, those in operators' own predicted orbits are much more uneven, and its appendix gives test statistics that check predicted covariances against precise positions from onboard satellite navigation [18]. The gap is hard because the truth that would check a stated uncertainty is seldom measured in operation. Practice needs estimators whose stated uncertainty is checked against the truth in operation, with a defined response when it is not.

2. *Told values that are verified.* The cobot's brochure braking, the pallet robot's empty braking and the air taxi's load sheet were each believed in full. The picking arm given an assumed belt speed placed 10 products and failed 40 after the speed changed (Table 10.4). Online estimates of mass, braking and friction that a safety monitor may use are not routine. The Open challenge box after this list takes the gap further.

3. *Proof that a task cannot be done.* The foiling boat's head sea defeated 33 795 settings, and a search tuned to one draw flew 59.07 s on it and the minute on none of fifteen others (Section 5.6). The cobot's table shows what such a proof looks like: at rest the protective distance of 532 mm exceeds the 214 mm the path leaves a person, so no speed override works there (Chapter 15). Practice sets limits from capacity and geometry. Guidelines for vessels on dynamic positioning ask for a consequence analysis that raises an alarm when the thrusters left after the worst failure could not hold the vessel in the prevailing weather [21], and a hydrofoil specification of 1977 makes the effective strut length at least 1.4 times the significant wave height of the 90th-percentile sea [22]. Neither bounds what a controller could achieve. The gap is hard because a negative result must hold for every setting, not only those tried. Methods that certify what no controller in a class can achieve, and so set an operating limit, remain thin.

4. *Predictive and learned controllers under model mismatch.* The lane-keeping predictive controller lost stability on a rack its model misjudged, and its limit fell from 32.3 to 16.9 m/s as the rack lag grew from 0.40 to 0.60 s (Section 2.6.7). Learned policies share the exposure, and neither comes with a stated margin against mismatch. A learned table tennis player trained in simulation and moved to the real robot without further training kept a gap between the two: its return rates on the real table stayed below the more than 80 % seen in simulation, and its builders corrected its choices online from real play [23]. The gap is hard because the mismatch that matters is the one the model leaves out, so the model cannot measure it. Progress would be a margin against mismatch, stated and tested on the real plant, beside every predictive or learned design.

5. *Contact with soft and moving bodies.* The surgical model is lumped and rigid, the peg-in-hole model omits the arm's dynamics and a damage criterion that depends on duration, and no setting tracked the beating heart within the vessel's width. An autonomous robot that joined small bowel in pigs detected the breathing motion and the deformation of the tissue, notified its operator to start a replanning step, and synchronized its tool with the breathing [24]. The gap is hard because tissue changes its shape and position under the tool, and neither change is measured directly. Models of deforming tissue and of damage that a controller can be designed against are still research.

6. *Benchmarks whose disturbances follow the event.* Two chapters had to build their tests this way. The transition challenge meets its gust 36.9 s into the transition (Chapter 4), and the peg-in-hole challenge draws its parts from a random stream of their own, so the sensor noise an attempt consumes cannot change the next part (Chapter 13). Section 16.4.1 tells how both needs were found. A NIST roadmap for robot dexterity lists grasping objects while in motion among the capabilities that need test methods, and its first tests were to assume static conditions, with no picking from a conveyor belt [17]. The gap is hard because a disturbance fixed in time rewards

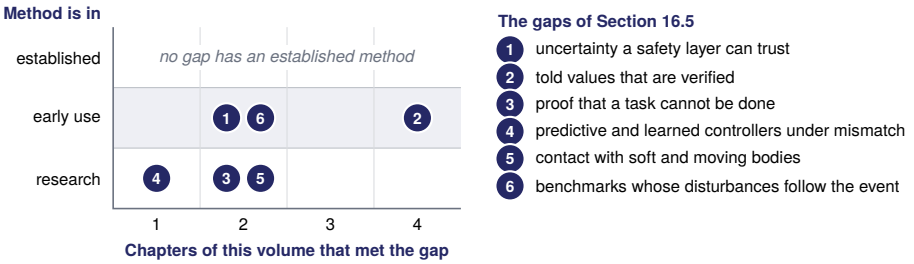


Figure 16.4 The six gaps of this section. Across is the number of chapters that each item names as evidence: two for gap 1 (Chapters 8 and 9), four for gap 2 (3, 9, 10 and 15), two for gap 3 (5 and 15), one for gap 4 (2), two for gap 5 (13 and 14) and two for gap 6 (4 and 13). Up is the maturity of the method available, the editor's reading of the sources cited in each item and in the Open challenge box. Notice that the gaps about stated uncertainty, told values and benchmarks have methods in early use, while proof of infeasibility, model mismatch and soft contact are still research.

a design that happens to arrive early or late. Shared benchmarks for moving machines need disturbances keyed to the maneuver, common random draws for every design and several seeds per verdict.

OPEN CHALLENGE **A Braking Figure That Stays True**

Four chapters name gap 2, and three of them believed a braking figure or a weight that nothing checked. A NIST study of collaborative cells notes that many installations treat the stopping time and distance as constants, though they should be measured again from time to time as the brakes wear. Not every integrator or user maintains a robot regularly, and the study asks for simple, inexpensive tests of joint brake wear [25]. In the editor's view the gap is hard because braking changes with load, pose and wear while the monitor needs one number it can trust. Progress would be a braking figure measured in service and checked against each stop the machine makes.

16.5.1 Gaps Common to Agent-Era Practice

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

1. *Verifying agent-written control code and models before they reach a plant.* Tests show that a controller behaves on the cases tested. For controllers learned from data, guarantees of stability and constraint satisfaction rest on assumptions about the model and its uncertainty, and open problems remain [26]. In the editor's view, code written by an agent has no comparable body of method yet. Practice needs a routine path from an agent's draft to code a plant can trust: independent derivation, tests of what the code reads and clamps, and trials with hardware in the loop.

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

3. *Safety cases and standards for agents in commissioning and operation.* Functional safety standards define a lifecycle of specified requirements, verification and assigned human roles [20]. In the editor's view, how an agent that writes configuration, runs commissioning checks or advises an operator is qualified, limited and audited under them is not yet settled.

4. *Security of agent access to control systems.* An agent with credentials to a control system is a new path into it. Industrial security standards state the requirements a control system must meet at each security level [27]. An agent that reads documents, messages or web pages can also be steered by text placed there. In the editor's view, least privilege, read-only access by default and human approval of every write are a starting point and not a solution.

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

6. *Transfer from simulation to hardware.* A controller developed on a simulator inherits what the simulator omits. The robotics literature documents the gap between simulated and real performance and the methods used to narrow it, such as randomizing the simulated plant and adapting on the real one [28]. Process plants add slow dynamics, rare upsets and little freedom to experiment, so each transfer still needs a staged plan from the simulator to a bench rig to the unit.

7. *Keeping engineers' judgment sharp when agents do the routine work.* Automation removes the routine work through which skill is built and leaves the person the hard cases, an irony described for process operators in 1983 [29]. The same now applies to engineers. An engineer who has never tuned a loop by hand or closed a balance on paper is poorly placed to review an agent that has. Deliberate practice, of the kind the exercises in this book provide, has to be designed into education and into work.

KEY POINTS

- A machine coped with what it could not see when it was told what had been checked, or estimated what drifts.
- A told value acted at once, and a wrong one was believed in full.
- An estimator needed a state for each unknown, and time to form it.
- Where no margin could be met, the right answer was to decline the task.
- Structure decided before tuning, and every verdict belonged to the truth.

References

[1] N. Kalra and S. M. Paddock, "Driving to safety: How many miles of driving would it take to demonstrate autonomous vehicle reliability?" *Transportation Research Part A: Policy and Practice*, vol. 94, pp. 182–193, 2016.

[2] P. Koopman and M. Wagner, "Autonomous vehicle safety: An interdisciplinary challenge," *IEEE Intelligent Transportation Systems Magazine*, vol. 9, no. 1, pp. 90–96, 2017.

[3] E. Kaufmann, L. Bauersfeld, A. Loquercio, M. Müller, V. Koltun, and D. Scaramuzza, "Champion-level drone racing using deep reinforcement learning," *Nature*, vol. 620, no. 7976, pp. 982–987, 2023.

[4] R. Firoozi, J. Tucker, S. Tian, A. Majumdar, J. Sun, W. Liu, Y. Zhu, S. Song, A. Kapoor, K. Hausman, B. Ichter, D. Driess, J. Wu, C. Lu, and M. Schwager, "Foundation models in robotics: Applications, challenges, and the future," *The International Journal of Robotics Research*, vol. 44, no. 5, pp. 701–739, 2025.

[5] K.-C. Hsu, H. Hu, and J. F. Fisac, "The safety filter: A unified view of safety-critical control in autonomous systems," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 7, pp. 47–72, 2024.

[6] 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

[7] V. Villani, F. Pini, F. Leali, and C. Secchi, "Survey on human–robot collaboration in industrial settings: Safety, intuitive interfaces and applications," *Mechatronics*, vol. 55, pp. 248–266, 2018.

[8] G. Fragapane, R. de Koster, F. Sgarbossa, and J. O. Strandhagen, "Planning and control of autonomous mobile robots for intralogistics: Literature review and research agenda," *European Journal of Operational Research*, vol. 294, no. 2, pp. 405–426, 2021.

[9] G.-Z. Yang, J. Cambias, K. Cleary, E. Daimler, J. Drake, P. E. Dupont, N. Hata, P. Kazanzides, S. Martel, R. V. Patel, V. J. Santos, and R. H. Taylor, "Medical robotics—regulatory, ethical, and legal considerations for increasing levels of autonomy," *Science Robotics*, vol. 2, no. 4, p. eaam8638, 2017.

[10] A. Attanasio, B. Scaglioni, E. De Momi, P. Fiorini, and P. Valdastri, "Autonomy in surgical robotics," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 4, pp. 651–679, 2021.

[11] T. Uriot, D. Izzo, L. F. Simões, R. Abay, N. Einecke, S. Rebhan, J. Martinez-Heras, F. Letizia, J. Siminski, and K. Merz, "Spacecraft collision avoidance challenge: Design and results of a machine learning competition," *Astrodynamics*, vol. 6, no. 2, pp. 121–140, 2022.

[12] Federal Aviation Administration, "Aircraft weight and balance control," Advisory Circular 120-27F, 6 May 2019, https://www.faa.gov/documentLibrary/media/Advisory_Circular/AC_120-27F.pdf, 2019.

[13] J. Liu, F. Jiang, S. Tashiro, S. Chen, and M. Tanaka, "A physics-informed and data-driven framework for robotic welding in manufacturing," *Nature Communications*, vol. 16, p. 4807, 2025.

[14] S. Yoon, R. Bostelman, and A. Virts, "Experiments to test the A-UGV capabilities standard," National Institute of Standards and Technology, Gaithersburg, MD, NIST Interagency/Internal Report 8407, 2021.

[15] S. Withrow-Maser, J. Aires, A. Ruan, C. Malpica, and S. Schuet, "Evaluation of heave disturbance rejection and control response criteria on the handling qualities evaluation of urban air mobility (UAM) eVTOL quadrotors using the vertical motion simulator," in *Vertical Flight Society Aeromechanics for Advanced Vertical Flight Technical Meeting*, San Jose, CA, 2022, nASA Technical Reports Server 20210026423, https://ntrs.nasa.gov/citations/20210026423.

[16] R. B. Madigan, T. P. Quinn, and T. A. Siewert, "Control of gas-metal-arc welding using arc-light sensing," National Institute of Standards and Technology, Boulder, CO, NIST Interagency/Internal Report 5037, 1995.

[17] J. Falco, J. Marvel, and E. Messina, "A roadmap to progress measurement science in robot dexterity and manipulation," National Institute of Standards and Technology, Gaithersburg, MD, NIST Interagency/Internal Report 7993, 2014.

[18] National Aeronautics and Space Administration, "NASA spacecraft conjunction assessment and collision avoidance best practices handbook," National Aeronautics and Space Administration, Washington, DC, Special Publication NASA/SP-20205011318, Dec. 2020. [Online]. Available: https://ntrs.nasa.gov/citations/20205011318

[19] 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.

[20] 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.

[21] International Maritime Organization, "Guidelines for vessels with dynamic positioning systems," MSC/Circ.645, Maritime Safety Committee, London, 1994, adopted 6 June 1994.

[22] D. R. Stark, W. E. Farris, and A. O. Harang, "Hydrofoil ship control and dynamics specifications and criteria, volume III," Boeing Marine Systems, for the David Taylor Naval Ship Research and Development Center, Document D321-51313-1, Mar. 1977, contract N00600-75-C-1107. Copy in the library of the International Hydrofoil Society, https://foils.org/control/.

[23] D. B. D'Ambrosio, S. Abeyruwan, L. Graesser, A. Iscen, H. Ben Amor, A. Bewley, B. J. Reed, K. Reymann, L. Takayama, Y. Tassa *et al.*, "Achieving human level competitive robot table tennis," in *2025 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2025, pp. 74–82.

[24] H. Saeidi, J. D. Opfermann, M. Kam, S. Wei, S. Leonard, M. H. Hsieh, J. U. Kang, and A. Krieger, "Autonomous robotic laparoscopic surgery for intestinal anastomosis," *Science Robotics*, vol. 7, no. 62, p. eabj2908, 2022.

[25] J. A. Marvel and R. Norcross, "Implementing speed and separation monitoring in collaborative robot workcells," *Robotics and Computer-Integrated Manufacturing*, vol. 44, pp. 144–155, 2017.

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

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

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

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

16.6 Exercises

16.1 Discussion · 8 min

An agent proposes to delete the load sheet of Chapter 3 and instead raise the altitude loop's integral action "so that the loop learns each passenger's mass by itself and a wrong sheet can never mislead it". Using Table 3.5 and the chapter's case 4, say what the proposal would gain and what it would lose, and what you would build instead.

16.2 Discussion · 8 min

Section 16.4.1 reports a loophole in the challenge of Chapter 4 and parts drawn differently for each controller in Chapter 13. Write the grading rules for a benchmark that compares agent-built and conventional controllers of moving machines, and say which failure each rule prevents.

16.3 Discussion · 8 min

An earlier version of the vessel app of Chapter 6 reported mean offsets of 1.34 and 0.63 m with the raw measurement and with the wave filter, and its texts credited the whole difference to the controller chasing the sea. Today Table 6.4 gives 0.55 and 0.51 m. Explain how the gap was apportioned, and write the business case for the wave filter that the chapter's numbers support.

16.4 Discussion · 8 min

A plant manager asks for two changes to raise throughput: switch off the pallet robot's slowing for unseen space in Chapter 9, which costs 8.7 s a run, and replace the collaborative arm's measured braking in Chapter 15 with the maker's brochure figure, which is faster. Discuss both requests with the chapters' numbers, and say what you would offer instead.

16.5 Discussion · 7 min

An agent reports a setting that flies the 1.4 m head sea of Chapter 5, where a search tuned to one draw flew 59.07 s of the minute and its best settings failed fifteen others. Say what evidence you would ask for before accepting the claim, and what the chapter's result means for an operating limit on the boat.

16.6 Discussion · 9 min

An agent hands you the code of a new attitude controller for a small satellite and a passing test report. List the checks you would run before the code reaches a hardware test, and say for each which defect of this volume it would have caught.