INTERACTIVE CONTROL ENGINEERING
Principles and Applications

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

CHAPTER 10

Learning From the Last Layer: A Powder Bed Fusion Simulator for Iterative Learning Control

Pages 243–259

Edited by John D. Hedengren

How to cite this chapter. "Learning From the Last Layer: A Powder Bed Fusion Simulator for Iterative Learning Control," in *Batch, Run-to-Run, and Hybrid Control*, John D. Hedengren, Ed., vol. 8 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 10, pp. 243–259.

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

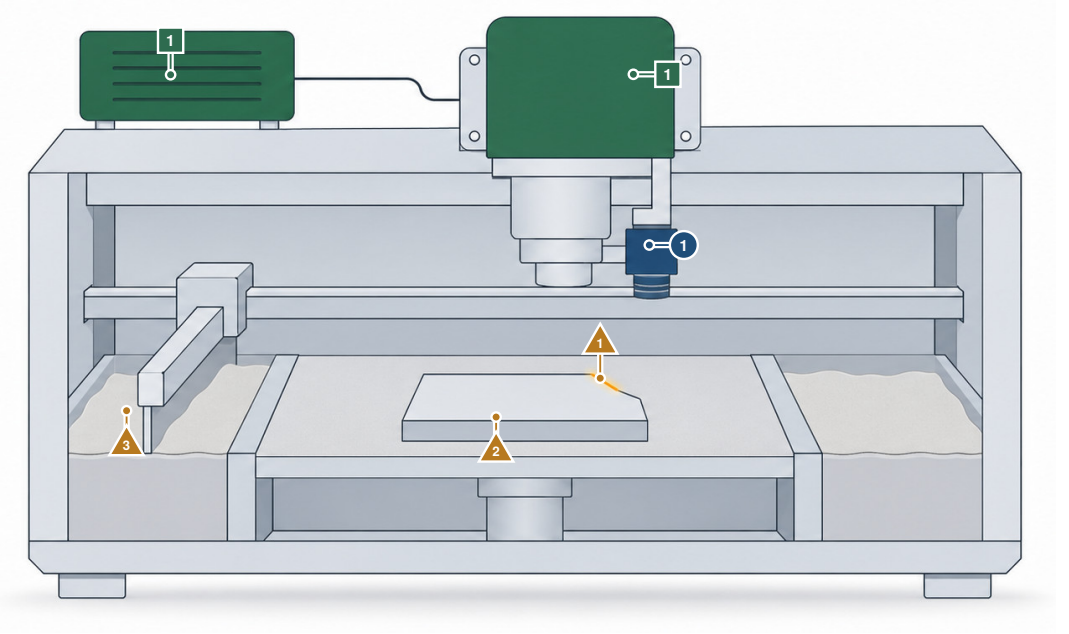
The app for this chapter
https://apmonitor.com/apps/laser-powder-bed/

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

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

CHAPTER 10

Learning From the Last Layer: A Powder Bed Fusion Simulator for Iterative Learning Control



The controller moves 1 the laser and the scan head that set the power along the path

The controller reads 1 the photodiode that watches the melt pool

What disturbs it 1 a sharp corner of the outline 2 heat that builds up in the part 3 a layer spread short of powder

The laser traces the same outline on every layer, so the pool overheats at the same corner each time and a controller can learn where. Illustration, not to scale.

IN PLAIN TERMS **A Laser That Learns From the Last Layer**

The main idea. A laser melts metal powder one thin layer at a time. It traces the same outline on every layer. At each sharp corner the small pool of molten metal runs too hot and too wide, in the same places every time.

The challenge. A controller that watches the pool and turns the power down reacts too late, because a corner is over in about a thousandth of a second. Remembering where the corners were on the last layer lets it lower the power just before each one.

Why it matters. A pool that is too wide or too narrow leaves tiny holes in metal parts made for aircraft and medical use.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	the first page, the history and the first fact in Section 10.2, the second fact and both try-its in Section 10.5, and the key points
Study	30 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	15 min	the first page, the results (Section 10.5), Technology in Practice (Section 10.6), Limitations (Section 10.7) and the two open challenges

The control question A powder bed fusion laser overheats the melt pool at the same corners on every layer, faster than a loop with sensor and laser latency can react. What can feedback do about it, what does learning from the last layer add, and what goes wrong when the layers stop repeating?

Learning objectives

- Explain why a melt pool loop with 300 µs of latency removes the slow heat drift but not the corner overheating, and why more gain makes the build worse.
- Compute the melt pool width the model supports at a corner and the heat that accumulates in a part from layer to layer.
- Explain how the learning update with a lead and a Q-filter converges, and why it diverges without the filter.
- Compare keeping and forgetting a learned profile when the outline changes, and judge how much feedback to keep under the learning.

Prerequisites Sections 1.4.6 (the learning update and its convergence factor), 1.6.2 (the PID of the feedback loop) and 1.6.1 (seed, speed and Reset), and Chapter 7 for a correction made between cycles.

The app https://apmonitor.com/apps/laser-powder-bed/

See also Volume 3, Chapter 6 for feedback that acts after the disturbance has passed and the feedforward that replaces it, and Volume 3, Chapter 2 for a loop structure that decides performance before any tuning.

10.1 Introduction

Feedback is taught as the remedy for disturbances, and for most disturbances it is. A disturbance that arrives faster than the loop can respond is the exception, and a disturbance that arrives at the same place in every repetition of a task is the exception that can be exploited. Iterative learning control was formulated for exactly that case: a robot repeating the same motion corrects its input from the error it saw on the previous repetition, and the error that repeats shrinks from one trial to the next [1, 2]. The theory that decides whether the learning converges, and what a low-pass Q-filter does to that answer, is mature [3]. Section 1.4.6 sets out the update and its convergence factor on the batch kettle.

Laser powder bed fusion is repetitive by construction. Each layer the laser traces the part cross-section, and features such as corners, thin walls and edges built over loose powder heat the melt pool at the same position on every layer. Too narrow a pool leaves unfused powder between tracks, too wide a pool becomes a vapor keyhole that collapses into pores, and both defects stay in the part [4, 5]. Coaxial photodiodes and cameras watch the melt pool, and closed-loop power control on their signal was demonstrated more than a decade ago [6]. A corner, however, lasts about a millisecond, which is comparable to the latency of the sensor, the filter and the laser modulation together.

A common assumption, which the simulator is designed to test, is that feedback fixes every repeatable error if tuned hard enough. The broader class is control of repetitive processes: batch reactors run the same recipe, wafer scanners repeat the same exposure, rolling mills pass the same product, and additive machines repeat nearly the same path, and in each the choice between reacting and remembering decides what is achievable [2]. Chapter 7 corrects the pack pressure of a molding machine from each weighed part, the same choice made one number per cycle. Spatial learning controllers indexed by position rather than time have been proposed specifically for additive manufacturing [7], and layer-to-layer control of interlayer temperature in powder bed fusion has been demonstrated more recently [8].

This chapter sets out a powder bed fusion simulator in which feedback, learning, the learning gain, the Q-filter and the response to a change in part geometry can be compared on the same build. Iterative learning control, its Q-filter and melt pool feedback are established methods. What the simulator adds is a contour-per-layer model whose geometry effects repeat exactly along arc length, with a learning controller indexed the same way. The scenarios move from a fixed-power build through a learned one to a build that changes shape.

Section 10.2 reviews melt pool monitoring and why it led to control. Section 10.3 sets down the melt pool, the geometry along the path, the heat buildup and the photodiode, with the assumptions. Section 10.4 gives the feedback loop on the photodiode and the learning controller with its lead and Q-filter. Section 10.5 reports the measured builds: what feedback can and cannot do, what learning does, why the Q-filter matters, overhangs, heat and broken repetition, and the graded challenge. Section 10.6 sets the loops on machine hardware, Section 10.7 states the limitations, Section 10.8 summarizes, and the exercises follow in Section 10.9.

10.2 Melt Pool Monitoring and Control

Powder bed fusion began as a prototyping process and became a production one for metal parts over the 2000s, and with that change the question of process quality moved from appearance to porosity, microstructure and properties [5]. The melt pool is where those are decided. Its width and depth follow the energy delivered per unit length of track, and above a threshold the pool changes from conduction mode to keyhole mode, with a sharp increase in depth and in porosity [4].

Monitoring grew first. How a coaxial photodiode and camera came to close a loop on laser power is told in the history box below. Quality control schemes for powder bed machines built on the same sensors [9], and reviews of in-situ monitoring later cataloged photodiodes, pyrometers, cameras, acoustic sensors and layer imaging, together with the difficulty of relating a sensor signal to a defect [10]. The same signals define the practical limit of melt pool feedback, because corners and overhangs change the pool within a time comparable to the sensing, filtering and actuation delay.

> ⧖ HISTORY **2010 to 2012: A Photodiode Finds the Hot Corners**
>
> A university group at Leuven built a control system for laser powder bed fusion around the melt pool. A camera and a photodiode looked down the laser's own optical path. The camera gave the pool's shape, the photodiode signal followed the pool's area, and that signal was fed back to the laser power, which improved the surface roughness of the parts. The group later saw the photodiode reading rise where the pool overheated at acute corners of the scan path. By mapping the signal against the beam position on each layer, it found deformation from thermal stress and hot zones over overhangs. A 2015 review of what real-time control of the process still needs tells this sequence first among the control schemes it surveyed [11].

Qualification drives the interest in control. Metal parts for aerospace and medical use must be shown to be free of critical defects, and doing so part by part through inspection is costly, which has made process monitoring and control central to qualification strategies [12]. Iterative learning control fits the problem because the path repeats, and because a learned correction is a record of what the machine did to each feature. Spatial formulations indexed by position for additive processes [7] and layer-to-layer feedback on part temperature [8] are two recent answers, and the simulator described here uses the same idea on melt pool width along a contour.

> ✦ DID YOU KNOW? **Powder Holds the Heat In**
>
> Loose metal powder conducts heat far worse than solid metal, potentially less than 1 percent as well. Near an overhang or a fine feature the powder keeps heat in the melt pool, which overheats, turns unstable and grows [13]. The model's downskin edge over loose powder adds 35 percent to the width the pool supports (Table 10.2).

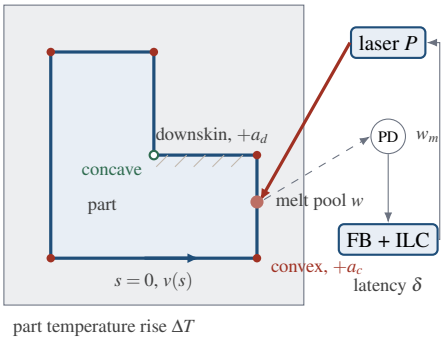


Figure 10.1 One contour pass of laser powder bed fusion. The laser of power P traces the part outline from $s = 0$ at scan speed $v(s)$, slowing at each corner. Convex corners (filled dots) add heat a_c, the concave corner (open dot) removes it, and the downskin edge over loose powder (hatched) adds a_d. A coaxial photodiode PD reads the melt pool width w as w_m, and the feedback and learning controllers set the power after a latency δ. The part warms by ΔT as layers accumulate.

Table 10.1 Variables at the Nominal Operating Point

Symbol	Role	Nominal	Range
w	CV, melt pool width	100 µm	SP 70–130
P	MV, laser power	200 W	0–370
$P_{ff}(s)$	MV, learned correction	0 W	−200 to 170
w, s, E	states	—	—
ΔT	per-layer state, part	0 K	0–150
$g(s)$	DV, geometry gain	1	0.93–1.44
$v(s)$	DV, scan speed	800 mm/s	400–800
σ_{sp}	DV, spatter (1σ)	3 %	0–20
ν	DV, photodiode noise	2 µm	0–20
f_{pow}	DV, powder layer	1	0.85 once

10.3 Process Model

The laser traces the outline of a 10 mm square part at 800 mm/s and a nominal 200 W, one pass per layer, with the melt pool width targeted at 100 µm. A pass takes about 52 ms. The recoat between layers is not simulated in time, so the clock counts scanning time only. The arrangement is shown in Figure 10.1.

The variables are listed in Table 10.1. Three states are integrated: the melt pool width w, the arc position s along the outline and the energy E delivered on the current layer. The part temperature rise, the learned profile and the layer scores change once per layer.

10.3.1 Melt Pool Width

The width follows a first-order lag behind the width the local conditions support, Equation (10.1).

$$\tau_p \frac{\mathrm{d}w}{\mathrm{d}t} = w_{ss} - w$$
$$w_{ss} = w_0\, g(s)\,(1 + \alpha \Delta T)\, f_{pow} \left(\frac{P}{P_0}\right)^n \left(\frac{v_0}{v(s)}\right)^m \tag{10.1}$$

with $\tau_p = 0.3$ ms, $w_0 = 100$ µm at $P_0 = 200$ W on bulk metal, $n = 0.7$, $m = 0.35$ and $\alpha = 1.5 \times 10^{-3}\,\mathrm{K}^{-1}$. The plant gain at nominal is $\hat{G} = n w_0 / P_0 = 0.35$ µm/W. Each pass starts with the pool already formed, as if the laser were switched on over a lead-in outside the part.

Table 10.2 Process Parameters, With Hidden Ones Marked

Symbol	Meaning	Value	Class
τ_p	melt pool time constant	0.30 ms	visible
w_0, P_0	nominal width, power	100 µm, 200 W	visible
n, m	power, speed exponents	0.70, 0.35	hidden
v_0, ρ	scan speed, corner fraction	800 mm/s, 0.50	visible
a (convex, concave)	corner heat, chill	15, 12 %	visible
a_d	downskin heat	35 %	visible
α	width per kelvin	1.5×10^{-3} K^{-1}	hidden
c_E, κ	heat per joule, kept	1.36 K/J, 0.90	visible
P_{max}	laser power limit	370 W	visible
δ	sensor and control latency	300 µs	visible
—	defect window	85–125 µm	visible

10.3.2 Geometry Along the Path

The outline is parameterized by arc length s and the effects of geometry are fixed functions of s, identical on every layer of the same outline, Equations (10.2) and (10.3).

$$g(s) = 1 + \sum_c \pm a\, b(s - s_c) + a_d \mathbb{1}_d(s) \tag{10.2}$$

$$v(s) = v_0 \left[1 - \tfrac{1-\rho}{2} \left(1 + \cos \frac{\pi (s - s_c)}{d_c} \right) \right] \tag{10.3}$$

The corner bump b rises as a Gaussian of width 0.3 mm into the corner and decays exponentially over 1.2 mm after it. Convex corners add $a = 15\,\%$ and concave corners remove 12 %, the downskin edge adds $a_d = 35\,\%$ with 0.3 mm smooth edges, and the mirrors slow the spot to $\rho = 0.5$ of full speed within $d_c = 0.5$ mm of every corner, Equation (10.3) applying for $|s - s_c| < d_c$. The square has four convex corners, and the L-shaped bracket has five convex corners, one concave corner and a 5 mm downskin edge, both with a 40 mm perimeter.

10.3.3 Heat Buildup and the Photodiode

Between passes the part keeps a fraction $\kappa = 0.9$ of its temperature rise and gains $c_E = 1.36\,\mathrm{K/J}$ per joule delivered, Equation (10.4).

$$\Delta T_{k+1} = \kappa \Delta T_k + c_E E_k \tag{10.4}$$

About 10 J per pass at 200 W settles a small part between 100 and 140 K warmer depending on the power. The photodiode reads the width with multiplicative spatter and additive noise, $w_m = w(1 + \sigma_{sp}\xi) + \nu$, where ξ is colored noise with a 0.2 ms time constant. Width outside the window from 85 to 125 µm is counted as defective, lack of fusion below and keyhole above. That defect share is a proxy for porosity and not a measurement of it.

Table 10.2 collects the parameter values.

10.3.4 Assumptions

One contour is scanned per layer. The interior hatch, which takes most of the energy and has its own repeating turnarounds, is not simulated.

The melt pool is a first-order lag with a power-law steady state. The sharp transition to keyhole mode is represented only by the upper width threshold [4].

Geometry and heat effects are fixed profiles along the path rather than the result of three-dimensional conduction, so they repeat exactly, which real layers only approximately do [5].

The photodiode is a width sensor with noise. Real signals depend on emissivity, plume and optics and need calibration before they can be read as a width [10].

The recoat takes no time and the part cools by a fixed fraction per layer.

10.4 Implemented Control

The engine of Section 1.6.1 integrates the melt pool with a fixed 20 µs Runge-Kutta step of fourth order, and the controllers execute at 10 kHz.

10.4.1 Feedback on the Photodiode

The feedback loop is a PI controller on the filtered photodiode width, executed every $\Delta t_c = 0.1$ ms in the velocity form of Section 1.6.2, Equation (10.5).

$$\begin{aligned} e_k &= w_{SP} - \mathscr{F}_{\tau_f}[w_m]_k \\ \Sigma_k &= \Sigma_{k-1} + K_c\,(e_k - e_{k-1}) \\ c_k &= \Sigma_k + I_{k-1} \\ u_k &= \mathrm{clamp}(c_k, -P_b, P_{max} - P_b) \\ I_k &= I_{k-1} + \frac{K_c}{\tau_I}\, e_k\, \Delta t_c + \frac{\Delta t_c}{\tau_t}\,(u_k - c_k) \end{aligned} \tag{10.5}$$

where $P_b = P_0 + P_{ff}(s)$ is the power the learned profile asks for and the commanded power is $P = P_b + u_k$. The running sum Σ_k, the demanded output c_k and the integral I_k are the S_k, v_k and I_k of Equations (1.15) to (1.17), with no derivative and no bias. The loop acts to cut power when the pool is too wide. The defaults are $K_c = 1.0\,\mathrm{W/\mu m}$, $\tau_I = 1$ ms and a photodiode filter $\tau_f = 0.2$ ms. The clamp keeps the total power between zero and P_{max} around whatever the learned profile asks for, and while it acts the back-calculation, with tracking time $\tau_t = \tau_I$, removes a tenth of the clipped excess from I each scan. The command is held for the control period and reaches the melt pool after the 300 µs latency. In manual mode the operator sets the total power, and transfer in either direction is bumpless. Each run starts with the photodiode filter at the initial width and the delay line holding the nominal power, so the first scan commands 200 W and a power chosen before the run reaches the pool only after the latency.

10.4.2 Learning Along the Outline

The learned profile P_{ff} is stored in 0.2 mm bins along the outline. During a pass the controller averages the measured width and the feedback output in each bin. After the pass the profile is updated by Equation (10.6), the update of Equation (1.10) indexed by arc length instead of time, with the feedback effort added.

$$P_{ff}^{k+1}(s) = Q\left[P_{ff}^{k}(s) + \bar{u}_{fb}^{k}(s) + \frac{\gamma}{\hat{G}}\, e_k(s+\lambda)\right] \tag{10.6}$$

where e_k is the error of the bin-averaged measured width, $\gamma = 0.5$ the learning gain, $\lambda = 0.6$ mm the lead that times each correction ahead of the melt pool lag, the latency and the filter, and $\bar{u}_{fb}$ the average feedback output in each bin, included only when feedback is running. The lead is a whole number of bins, and its setting steps in 0.2 mm. Q is a zero-phase first-order low-pass applied forward and backward around the closed outline with length $\ell_Q = 0.2$ mm, so that $a = \exp(-0.2\,\mathrm{mm}/\ell_Q)$ per bin. The profile is clamped so that the feedforward power stays between zero and P_{max}. When the outline changes between layers, the profile is either kept or reset to zero. Figure 10.2 draws the update as it runs between two layers.

The lead of 0.6 mm is the sum of the melt pool time constant, the latency and the filter at full scan speed, as if all three were pure delays. The condition for monotone convergence at a spatial frequency ω is Equation (10.7), the factor of Equation (1.11) written along the path.

$$\left| Q(\omega)\left[1 - \frac{\gamma}{\hat{G}}\, G(\omega)\, e^{j\omega\lambda}\right]\right| < 1 \tag{10.7}$$

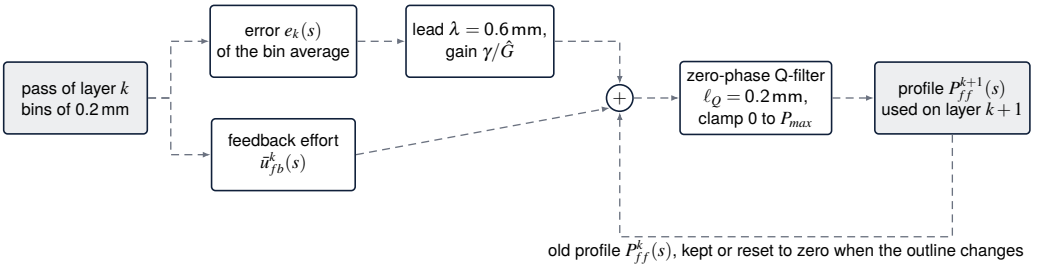


Figure 10.2 The learning update of Equation (10.6) as it runs between two layers. The error of each 0.2 mm bin is read a lead ahead and scaled by the learning gain, the feedback effort of the bin is folded in, and the sum with the old profile is filtered and clamped before the next layer uses it. The Q-filter is the block that keeps short wavelengths, where the fixed lead is wrong, out of the learning.

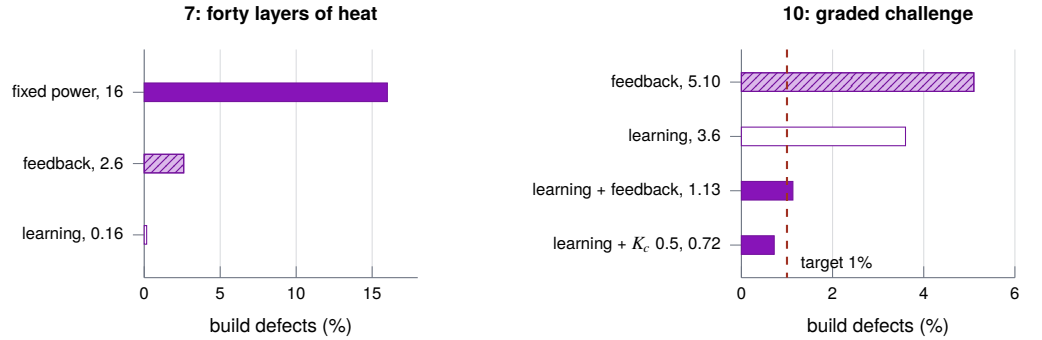


Figure 10.3 Build defects of two cases of Table 10.4, with each value printed beside its label. Over forty layers of heat buildup, feedback removes the slow drift but not the corners, and learning removes both. In the graded challenge, learning alone leaves most of its defects on the first bracket layer and the short-fed layer, and learning with feedback underneath passes the 1 percent target (dashed) only with the feedback gain lowered to 0.5 W/µm.

At short wavelengths a first-order lag contributes much less phase than a delay of the same time constant, so a fixed lead overcompensates there, and the bracket in Equation (10.7) can exceed one [3]. The Q-filter removes those wavelengths from the learning.

🕮 **Pause.** Try Exercise 10.2 before reading on.

10.5 Results

Table 10.3 lists the ten cases, and Table 10.4 collects the outcomes measured at the app's default seed (Section 1.6.1). Each case was run from its loaded settings, with any change made before Run. Figure 10.3 draws the build defects of cases 7 and 10 from that table.

10.5.1 What Feedback Can and Cannot Do

At a fixed 200 W the width reaches 140 µm at every corner on the first layer and 4.7 % of the path is in the keyhole regime. By the fifth layer the part is about 50 K warmer, the corners reach 150 µm and 7.2 % of the path is defective.

Feedback at 1 W/µm takes the corners to about 135 µm and the keyhole share to about 2.5 % on every layer, with an rms error near 6.7 µm, and it removes the slow heat drift completely. It cannot do more, because by the time the photodiode sees the pool widen, the filter, the 300 µs latency and the melt pool lag have used most of the corner. Volume 3, Chapter 6 meets the same limit on a

Table 10.3 Scenario Sequence

#	Scenario	Teaching point
1	Fixed power	corners and heat
2	Feedback on the melt pool	what feedback cannot catch
3	More feedback gain	latency and spatter
4	Learn from the last layer	iterative learning
5	Without the Q-filter	divergence
6	A bracket with an overhang	learning with feedback
7	Forty layers of heat	slow drift
8	The part changes shape	broken repetition
9	Learning with feedback	a layer that does not repeat
10	Challenge	graded build

Table 10.4 Measured Results

Case and condition	Result
1, layer 1 / layer 5	corners 140 / 150 µm, defects 4.7 / 7.2%
2, feedback $K_c = 1$, 10 layers	corners ≈ 135 µm, ≈2.5%, 6.7 µm rms
3, K_c 1 / 2 / 3 / 5	build defects 2.6 / 2.9 / 4.1 / 36.4%
3, spatter 10%, K_c 3 / 1	build defects 8.2 / 2.7%
4, ILC $\gamma = 0.5$	4.7, 1.8, 0% from layer 3, 3.0 µm at layer 10
4, $\gamma = 1.0$ / 1.5	3.4 µm at layer 3 / 1.0–3.1% on layers 6–10
5, no Q, 25 layers	low 5.3 µm at layer 4, then 17 µm, 24.8%
5, ℓ_Q 0.2 / 0.6 mm	2.5–3.1 / 4.7 µm
5, no Q, lead 0.2 mm	converges to 1.6 µm
6, bracket, feedback	6.3% (3.6 keyhole, 2.8 lack of fusion)
6, feedback and ILC	0.7% at layer 2, then ≤0.4%
6, $K_c = 0.5$ with ILC	0% from layer 3, 4.2 µm rms
7, 40 layers, fixed power	+140 K, 4.7 to 23.6% per layer, build 16%
7, feedback / ILC	build 2.6 / 0.16%
8, shape change, keep / reset	13.2 / 23.2% on first new layer
8, keep, with feedback	5.5%
9, short feed, ILC only	12.5% then peak 127 µm, build 1.7%
9, ILC and feedback $K_c = 0.5$	0, 0.4, 0.7%, build 0.40%
10, FB / ILC / ILC+FB / K_c 0.5	build 5.10 / 3.6 / 1.13 / 0.72%

slower scale, where a coagulant loop on the settled water acts on water that has already left the plant. Raising the gain makes things worse. The defect share over ten layers is 2.6 % at 1, 2.9 % at 2 and 4.1 % at 3 W/µm, because the extra gain turns spatter into power and pushes the width below the lack-of-fusion threshold. Over eight other noise seeds the share is 2.6 to 2.7 % at 1 and 2.7 to 3.5 % at 2, higher at 2 for seven of the eight, so the first step is small. At 5 the latency makes the loop oscillate, growing layer by layer until 74.0 % of the tenth layer and 36.4 % of the build are defective. With 10 % spatter the gap widens to 8.2 % against 2.7 %.

Pause. Try Exercise 10.1 before reading on.

DID YOU KNOW? More Gain, More Defects

From the simulator: raising the feedback gain from 1 to 2, 3 and 5 W/µm raises the defective share of a ten-layer build from 2.6 to 2.9, 4.1 and 36.4 percent. The extra gain turns spatter into power, and at the highest gain the latency makes the loop oscillate (Table 10.4).

10.5.2 What Learning Does

With feedback off and learning on at $\gamma = 0.5$, the first layer is the fixed-power case at 4.7 %, the second is 1.8 %, and from the third there are no defects, while the rms error continues down to 3.0 µm by layer ten, less than half of what feedback achieved (Table 10.4). Figure 10.4 shows the defects

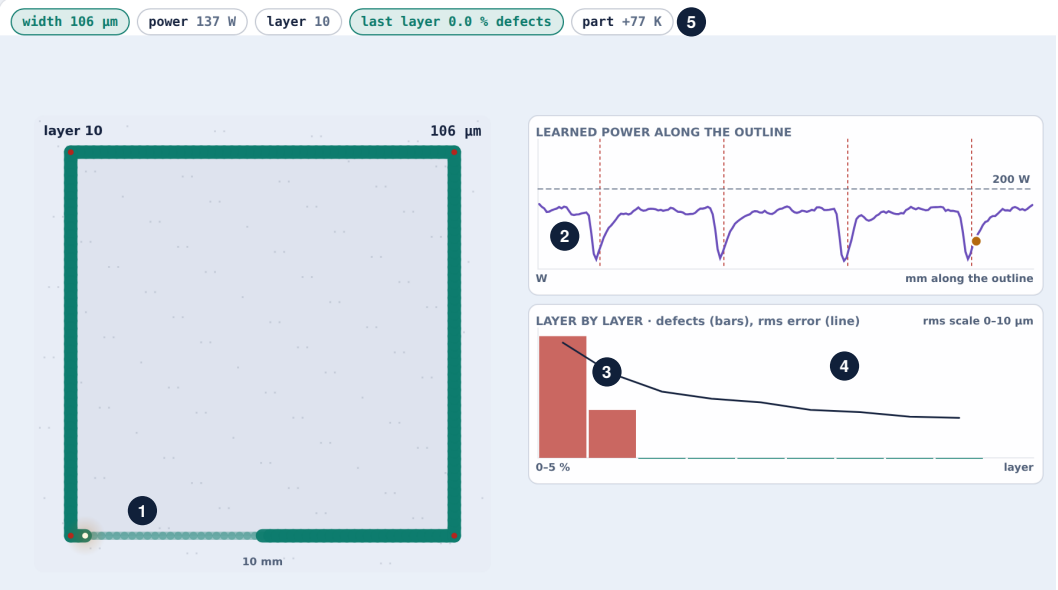


Figure 10.4 The powder bed application running scenario 4, learning without feedback on the square outline, cropped to the plant view. The run is at 0.515 s, 35 mm into the tenth and last layer. (1) The melt pool moves along the outline, and the melted track is colored by width over the faint track of the previous layer, with the four corners marked. (2) The learned power profile dips a fraction of a millimeter before each corner, each marked by a dashed line. (3) The layer-by-layer bars show defects on the first two layers only. (4) The rms error of each layer falls from layer to layer. (5) The status chips give the width, the power, the layer and the part 77 K warmer.

confined to the first two layers and the rms error falling. A learning gain of 1 reaches 3.4 µm by the third layer rather than the eighth, at the cost of 2.1 % lack of fusion on the second. A gain of 1.5 overcorrects, and because each update also copies the noise of the previous layer it leaves 1.0 to 3.1 % defective on layers 6 to 10.

> **TRY IT IN THE APP** **4 · Learn From the Last Layer**
>
> Load the scenario and press Run. Feedback is off and learning is on for ten layers. Watch the layer-by-layer plot: 4.7 percent of the first layer is defective, 1.8 percent of the second, and none from the third on, while the rms error falls to 3.0 µm by layer ten (Table 10.4).

10.5.3 Why the Q-Filter Matters

Without the Q-filter the learned profile improves for a few layers, with the rms error lowest at 5.3 µm on layer four, and then short-wavelength ripple grows layer after layer until 24.8 % of the path is defective at layer 25 with an rms error of 17 µm. With a 0.2 mm filter the error settles at 2.5 to 3.1 µm with no defects after layer two, and with a 0.6 mm filter it converges more slowly, with defects until layer eight, and settles near 4.7 µm because the filter also blurs the corners. A lead of 0.2 mm, one bin and closer to the true lag, converges to 1.6 µm even without a filter, while a lead of 0.4 mm already drifts back up to 4.4 µm by layer 25. The filter is insurance against the model error in the lead, which a real machine never knows exactly, and choosing it is most of the design effort in a practical learning controller [2, 3].

TRY IT IN THE APP **5 · Without the Q-Filter**

Load the scenario, press Run and let all 25 layers run. The rms error is lowest at 5.3 µm on layer four, then climbs to 17 µm, with 24.8 percent of layer 25 defective. Press Reset, set Q-filter length to 0.2 mm and run again: the error settles at 2.5 to 3.1 µm with no defects after layer two.

10.5.4 Overhangs, Heat and Broken Repetition

The bracket adds an overhang edge that runs hot and an inside corner that runs cold. Feedback alone leaves about 6.3 % of every layer defective, 3.6 % keyhole and 2.8 % lack of fusion. With learning added to feedback the defects fall to 0.7 % on the second layer and stay at or below 0.4 %, and with the feedback gain lowered to 0.5 W/µm they are zero from the third layer at 4.2 µm rms, because feedback no longer has the repeating part to chase and only adds noise. Over forty layers at fixed power the part warms by about 140 K, the defect share rises from 4.7 to 23.6 % per layer and 16 % of the build is defective. Feedback removes the drift and leaves 2.6 %, all at the corners, and learning alone leaves 0.16 %, all on the first two layers.

When the square becomes a bracket at layer nine, the learned profile lowers power where a convex corner has become a cold concave one and has nothing for the two new corners and the overhang, so 13.2 % of that layer is in the keyhole regime, and one layer later the new shape has been learned. Forgetting the profile is worse, 23.2 %, because three corners and the heat buildup were still right. With feedback running underneath, the first bracket layer is 5.5 % defective. A layer fed short of powder is the opposite failure. Learning alone leaves that layer 12.5 % lack of fusion (12 to 22 % over eight other noise seeds) and then applies the extra power it learned to the next layer, where the width peaks at 127 µm, for 1.7 % of the build. With feedback at 0.5 W/µm underneath, the short layer has no defects, the next is 0.4 % defective, the one after 0.7 % with lack of fusion, and the build 0.40 %.

Pause. Try Exercise 10.3 before reading on.

10.5.5 The Graded Challenge

The challenge builds twenty layers, six of the square and then the bracket, with a short powder feed on layer fifteen and the part heating throughout, starting with feedback only at 1 W/µm. It is met with less than 1 % of the scanned path outside the width window and the width within 30 µm of setpoint over the last third. As loaded, feedback leaves 5.1 % of the build defective, about 6 % of every bracket layer, and the width 48.6 µm from setpoint. Learning alone leaves 3.6 %, most of it on the first bracket layer and on the short-fed layer, where 50 % of the path is lack of fusion. Learning with the default feedback gain leaves 1.13 %, just over the target, with the width 34.5 µm from setpoint, and learning with the feedback gain lowered to 0.5 W/µm leaves 0.72 % and 26.0 µm, with the first bracket layer the worst at 8.7 % and the short-fed layer at 0.6 %. That run passes with 0.28 percentage points and 4.0 µm to spare, and the same settings carried through Reset give the same result. The margins are narrow against the noise. Over eight other seeds the default gain passes once and the lower gain fails once, on the width.

10.6 Technology in Practice

What Measures It

The loop of this chapter reads one signal, the melt pool width seen by a coaxial photodiode and filtered with a 0.2 ms time constant, and the learning controller averages the same signal in 0.2 mm bins along the outline. Figure 10.5 draws both loops on machine hardware. A research testbed carries more than one view of the process: a coaxial high-speed camera behind a band-pass filter at 850 nm, which records up to 10 000 frames a second, and a layer camera above the bed that images every layer before and after the powder is spread [14]. The chapter treats its photodiode as a calibrated width sensor, which Section 10.7 lists as a limitation.

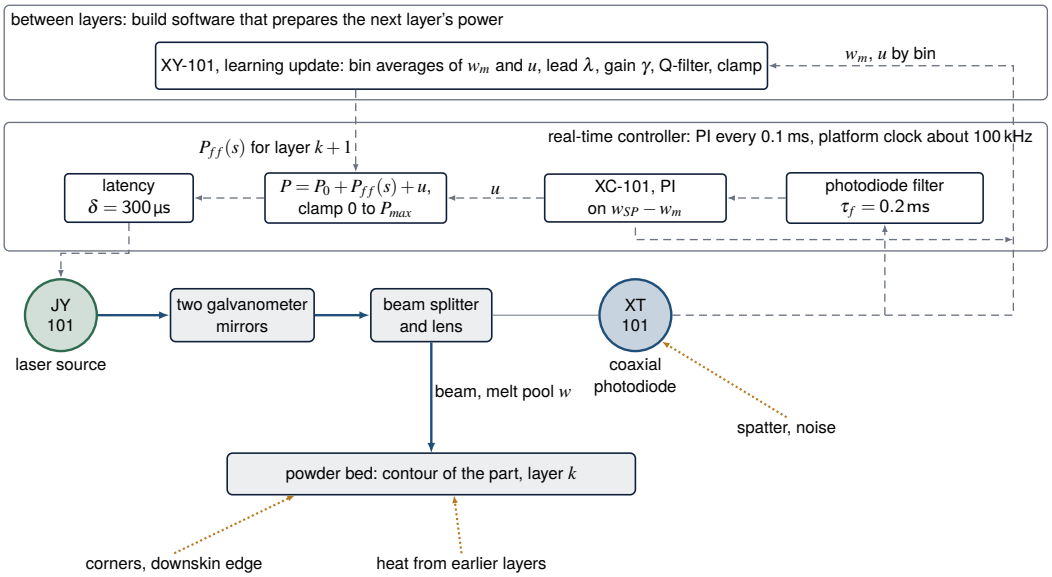


Figure 10.5 The chapter's two loops on machine hardware. Inside a pass, the PI loop on the filtered photodiode width trims the power that the learned profile asks for, and the command reaches the melt pool after the latency. Between passes, the learning update turns the binned width error and feedback effort of layer k into the power profile of layer $k+1$. Dotted arrows mark disturbances. Tags are the book's; the chapter names none.

TECHNOLOGY **Sensor: A Coaxial Photodiode**

A coaxial camera looks down the laser's optical path, so the melt pool stays still in its view wherever the mirrors point [14]. A photodiode placed the same way gives one signal that follows the pool's area [11]. It integrates the thermal emission continuously and has no spatial resolution, but its high sampling rate and its sensitivity to thermal transients complement a camera [15]. One research machine sampled its coaxial photodiode at 50 kHz in a band of 700 to 1000 nm [16]. On a trend by layer, sharp rises in intensity marked hot spots, a surface raised above the build plane and swelling [15]. The model's photodiode reads the width itself, with 3 % spatter and 2 µm of noise (Table 10.1).

What Moves It

Laser power is the only manipulated variable of the chapter, and the path and the speed along it are fixed by the outline. On a real machine the scan path and the power are generated from the sliced part before the build and interpolated into digital position and power commands. A power adjustment module can sit between those commands and the controller and change the power before it is sent [13]. That is where a learned profile such as $P_{ff}(s)$ would enter, as a power written into the next layer's commands.

TECHNOLOGY Final Element: Laser and Galvanometer Mirrors

Two mirrors turned by a galvanometer system steer the beam, a lens focuses it to a spot 0.05 to 0.5 mm wide, and scan speeds reach 7 m/s [11]. The controller sends two digital streams: positions to the galvanometer drive, which closes its own PID loop, and a power command to the laser, usually as an analog voltage [13]. Command and result differ. On one testbed the true power was slightly less than commanded, and a galvanometer position error built the parts slightly contracted [14]. In the model the power reaches the pool 300 µs after its command, and the spot slows to half speed near every corner (Table 10.2).

Where the Controller Runs

A typical controller of the process works on a clock of 100 kHz [13], and one testbed issues each position and power command on a 10 µs step [14]. The PI loop of this chapter executes every 0.1 ms, ten of those steps, and its learning update runs once per layer, after the pass. A control loop may act discretely between the completion of build layers rather than continuously [11], and commercial systems already adjust the energy input at every layer from layer-wise measurements [16]. Read as machine hardware, the PI loop belongs in the real-time controller beside the mirror and power streams, and the learning update with its lead, Q-filter and clamp belongs in the software that prepares the next layer between passes.

IN PRACTICE A Qualified Process Is a Locked Process

The space agency standard NASA-STD-6030, *Additive Manufacturing Requirements for Spaceflight Systems*, released in 2021, sets qualification and certification strategies for additive parts. A qualified and locked process defines the flaw state a part is expected to carry, and each part must be shown to represent that process. In-situ monitoring today looks for the conditions that lead to flaws rather than for the flaws themselves, and a monitoring baseline of one component can serve as a fingerprint for its later builds [17]. The defect share of this chapter is likewise a proxy for porosity, not a measurement of it (Section 10.3).

What Is Modernizing It

Three developments touch this loop. The first is feedforward from geometry. A geometric conductance factor computed from the share of solid and powder near the melt pool lowers the power near edges and overhangs, and the parts built with the strongest adjustment showed the lowest variation of melt pool intensity [13]. The second adds a model. A vector-level feedforward controller that couples a fast thermal model to a reduced melt pool model reduced geometric inaccuracy in key dimensions by 62 percent, porosity by 16.5 percent and photodiode variation by 6.8 percent on average [16]. Both act before the pool widens, as the learned profile of this chapter does, but from a model rather than from the last layer. The third is richer monitoring: photodiode traces mapped layer by layer beside camera images to classify anomalies, as groundwork for machine learning [15], and monitoring baselines used as a fingerprint of a qualified build [17].

10.7 Limitations

Only the contour is scanned. The hatch that fills the interior dominates the energy input and has its own features, such as vector turnarounds, that a learning controller would also address.

The melt pool model is phenomenological. A real pool changes mode abruptly at the keyhole threshold, and its width is only one of the quantities that determine porosity [4].

Geometry repeats exactly on every layer of the same outline. Real cross-sections change continuously with height, which blurs the distinction between the repeating error and the new one [7].

OPEN CHALLENGE **Which Loop, on Which Time Scale**

Most reported control of powder bed fusion uses melt pool temperature or size to set the laser power or speed. Loops could also act once between layers rather than continuously, but it is not yet clear which signatures are best measured or modeled, and which inputs are best controlled on which time scale. A review of the field called for a control architecture that names every possible loop and shows which one suits which link between parameter, signature and part quality [11]. This chapter compares two such loops on one contour. Progress would be that comparison made on whole parts, with the hatch and evolving cross-sections.

The photodiode is treated as a calibrated width sensor. Relating a real photodiode signal to a width is itself a modeling problem [10].

The part cools by a fixed fraction per layer, whereas real interlayer temperature depends on part geometry, scan strategy and recoat time, which is the subject of layer-to-layer temperature control [8].

OPEN CHALLENGE **Learning Under a Locked Recipe**

Feedback on the melt pool stays reactive and cannot act ahead of a known overhang, while changing process parameters in real time, layer by layer or faster, complicates certifying parts for critical applications [16]. Qualification rests on a locked process [17], and a learned profile changes the power on every layer until it converges. Progress would be a way to qualify the learning controller itself, with bounds on what it may change and each learned profile kept as part of the build record.

10.8 Conclusion

A melt pool that overheats at the same corners on every layer is a disturbance feedback cannot catch and learning can. In the simulator described here, fixed power leaves 4.7 % of the first layer in the keyhole regime, feedback with 300 µs of latency trims it to about 2.5 %, and more gain raises the defect share to 4.1 % at 3 W/µm. Learning from the previous layer removes the defects from the third layer on and more than halves the rms error to 3.0 µm. Without its Q-filter the learning diverges to 24.8 % defects by layer 25, because a fixed lead overcompensates the melt pool lag at short wavelengths. When the repetition breaks, keeping the learned profile (13.2 % on the first new layer) beats discarding it (23.2 %), and feedback underneath the learning limits the damage of a layer that does not repeat, taking the build from 1.7 % to 0.40 % defective. The model scans one contour, uses a phenomenological melt pool and counts a width proxy rather than porosity. Further work would add the hatch, a keyhole transition, geometry that evolves with height and a sensor model. The app is a way to practice control of a process that repeats itself.

KEY POINTS

- A melt pool that overheats at the same corners on every layer is a disturbance feedback cannot catch and learning can.
- Feedback with 300 µs of latency trims the keyhole share to about 2.5 %, and more gain makes the build worse.
- Learning from the previous layer removes the defects from the third layer on and more than halves the rms error to 3.0 µm.
- Without its Q-filter the learning diverges, because a fixed lead overcompensates the melt pool lag at short wavelengths.
- When the outline changes, keeping the learned profile beats discarding it, and feedback underneath the learning limits the damage of a layer that does not repeat.

Where else this applies. Chapter 7 makes the same choice between reacting and remembering with one number per cycle (Section 1.4.5), and Section 1.4.6 sets out the learning update on the batch kettle.

References

[1] S. Arimoto, S. Kawamura, and F. Miyazaki, "Bettering operation of Robots by learning," *Journal of Robotic Systems*, vol. 1, no. 2, pp. 123–140, 1984.

[2] R. W. Longman, "Iterative learning control and repetitive control for engineering practice," *International Journal of Control*, vol. 73, no. 10, pp. 930–954, 2000.

[3] M. Norrlöf and S. Gunnarsson, "Time and frequency domain convergence properties in iterative learning control," *International Journal of Control*, vol. 75, no. 14, pp. 1114–1126, 2002.

[4] W. E. King, H. D. Barth, V. M. Castillo, G. F. Gallegos, J. W. Gibbs, D. E. Hahn, C. Kamath, and A. M. Rubenchik, "Observation of keyhole-mode laser melting in laser powder-bed fusion additive manufacturing," *Journal of Materials Processing Technology*, vol. 214, no. 12, pp. 2915–2925, 2014.

[5] T. DebRoy, H. Wei, J. Zuback, T. Mukherjee, J. Elmer, J. Milewski, A. Beese, A. Wilson-Heid, A. De, and W. Zhang, "Additive manufacturing of metallic components – Process, structure and properties," *Progress in Materials Science*, vol. 92, pp. 112–224, 2018.

[6] T. Craeghs, F. Bechmann, S. Berumen, and J.-P. Kruth, "Feedback control of Layerwise Laser Melting using optical sensors," *Physics Procedia*, vol. 5, pp. 505–514, 2010.

[7] D. J. Hoelzle and K. L. Barton, "A new spatial Iterative Learning Control approach for improved micro-Additive Manufacturing," in *2014 American Control Conference*, 2014, pp. 1805–1810.

[8] B. Kavas, E. C. Balta, M. Tucker, A. Rupenyan, J. Lygeros, and M. Bambach, "Layer-to-layer closed-loop feedback control application for inter-layer temperature stabilization in laser powder bed fusion," *Additive Manufacturing*, vol. 78, p. 103847, 2023.

[9] S. Berumen, F. Bechmann, S. Lindner, J.-P. Kruth, and T. Craeghs, "Quality control of laser- and powder bed-based Additive Manufacturing (AM) technologies," *Physics Procedia*, vol. 5, pp. 617–622, 2010.

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

[11] M. Mani, B. Lane, A. Donmez, S. Feng, S. Moylan, and R. Fesperman, "Measurement science needs for real-time control of additive manufacturing powder bed fusion processes," National Institute of Standards and Technology, Gaithersburg, MD, Tech. Rep. NISTIR 8036, 2015.

[12] M. Seifi, A. Salem, J. Beuth, O. Harrysson, and J. J. Lewandowski, "Overview of Materials Qualification Needs for Metal Additive Manufacturing," *JOM*, vol. 68, no. 3, pp. 747–764, 2016.

[13] H. Yeung, B. Lane, and J. Fox, "Part geometry and conduction-based laser power control for powder bed fusion additive manufacturing," *Additive Manufacturing*, vol. 30, p. 100844, 2019.

[14] B. Lane and H. Yeung, "Process monitoring dataset from the additive manufacturing metrology testbed (AMMT): Overhang part X4," *Journal of Research of the National Institute of Standards and Technology*, vol. 125, p. 125027, 2020.

[15] H. Hyer, L. Scime, Z. Snow, C. Joslin, T. McDonald, C. Leach, J. Kendall, B. Schreiber, and C. Petrie, "Demonstrating advanced sensors for in situ monitoring toward qualification of nuclear-relevant components," Oak Ridge National Laboratory, Oak Ridge, TN, Tech. Rep. ORNL/TM-2025/4056, 2025.

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

[17] W. Tilson, D. Wells, E. Lanigan, B. West, R. Russell, A. Park, B. McEnerney, and A. Glendening, "NASA's certification and qualification challenges for additively manufactured hardware: What is next?" in *Proceedings of the 17th European Conference on Spacecraft Structures, Materials and Environmental Testing (ECSSMET)*, Toulouse, France, 2023, nASA Technical Reports Server 20230003172.

10.9 Exercises

10.1 ○ Concept check · 1 min

In scenario 2 a PI loop on the photodiode, running every 0.1 ms at $K_c = 1\,\mathrm{W}/\mu\mathrm{m}$, still lets every corner reach about 135 µm. Which reason does the chapter give?

(a) The filter, the latency and the melt pool lag use up most of a corner that lasts about a millisecond

(b) The power limit of 370 W stops the loop from moving the power as far as a corner needs

(c) The integral time of 1 ms is too long for the integral to build up before each corner

(d) The 3 % spatter on the photodiode hides the widening at a corner from the loop

10.2 Concept check · 1 min

Without its Q-filter the learning controller of scenario 5 improves for four layers and then gets worse until a quarter of the path is defective. Why?

(a) The learning gain of 0.5 is too high for the nominal plant gain, so every wavelength overshoots

(b) The fixed 0.6 mm lead overcompensates the lag at short wavelengths, so those components grow

(c) The part heats from layer to layer, so the profile learned on one layer no longer fits the next

(d) The clamp on the learned power cuts the corner corrections and leaves an error that accumulates

10.3 Concept check · 2 min

In scenario 8 the square outline becomes a bracket at layer 9. Keeping the learned profile leaves 13.2 % of that layer defective and forgetting it leaves 23.2 %. Why is keeping it better?

(a) The kept profile already holds a correction for the overhang, learned from the bracket on earlier layers

(b) Choosing Forget also switches the feedback loop off until the new outline has been scanned once

(c) Three corners and the heat buildup are unchanged, so most of what was learned is still right

(d) The kept profile is filtered again at the change, which removes the corner it no longer needs

10.4 Theory and calculation · 6 min

The scan speed falls to $\rho = 0.5$ of $v_0 = 800\,\mathrm{mm/s}$ within $d_c = 0.5\,\mathrm{mm}$ of every corner, by Equation (10.3). At the apex of a convex corner $g = 1.15$. Use $n = 0.7$, $m = 0.35$, $w_0 = 100\,\mu\mathrm{m}$ at $P_0 = 200\,\mathrm{W}$ and a cold part, and $\int_{-\pi}^{\pi} \mathrm{d}x/(A - B\cos x) = 2\pi/\sqrt{A^2 - B^2}$.

(a) Find the time the spot spends in the slowed zone of one corner, and the time of one pass around the 40 mm square with its four corners. *Checkpoint:* 1.77 ms and 52.1 ms.

(b) Find the width the melt pool would settle to at the apex at 200 W, and say why the first layer of scenario 1 peaks lower, at 140 µm. *Checkpoint:* 146.6 µm.

(c) Find the power that holds 100 µm at the apex, and the share of the time in the slowed zone that the photodiode filter, the latency and the melt pool time constant take together. *Checkpoint:* 115.8 W; 0.8 ms, about 45 %.

10.5 Theory and calculation · 6 min

At a fixed 200 W a pass of the square takes 52.06 ms. The part gains $c_E = 1.36\,\mathrm{K/J}$ per joule delivered and keeps $\kappa = 0.9$ of its rise through each recoat, by Equation (10.4), and the width rises by $\alpha = 1.5 \times 10^{-3}\,\mathrm{K^{-1}}$ per kelvin.

(a) Find the energy of one pass and the temperature rise the part settles to after many layers. *Checkpoint:* 10.41 J and 141.6 K.

(b) Find the rise during layer 5, after four passes, and the factor by which it widens the pool. *Checkpoint:* 48.7 K and 1.073.

(c) The corners of layer 1 peak at 139.6 µm. Predict the peak on layer 5, and the width on a straight edge, 100 µm on a cold part, once the part has settled. *Checkpoint:* 149.8 µm and 121 µm.

10.6 Theory and calculation · 7 min

The Q-filter of Section 10.4 runs $y_i = a\,y_{i-1} + (1-a)\,x_i$ over the 0.2 mm bins once forward and once backward, with $a = \exp(-0.2\,\mathrm{mm}/\ell_Q)$. A component that turns by θ radians per bin is multiplied by $(1-a)/|1-a\,\mathrm{e}^{-\mathrm{j}\theta}|$ on each pass.

(a) Find a for $\ell_Q = 0.2$ and 0.6 mm, and the combined gain of the two passes as a function of θ. *Checkpoint:* 0.368 and 0.717; $(1-a)^2/(1-2a\cos\theta+a^2)$.

(b) Evaluate that gain for a constant, for a pattern that alternates from bin to bin (a 0.4 mm wavelength) and for a 2 mm wavelength, for both filters. *Checkpoint:* 1, 0.214, 0.740 at 0.2 mm; 1, 0.027, 0.227 at 0.6 mm.

(c) The filter multiplies the bracket of Equation (10.7). Find the largest bracket magnitude the 0.2 mm filter still turns into convergence at the two wavelengths, and say what the 0.6 mm filter costs at a corner. *Checkpoint:* 4.7 and 1.35.

10.7 App challenge · 8 min

Load scenario 3, More feedback gain, which opens with Feedback gain at 3 W per µm. Use the default seed and a speed of 0.1, which plays 0.1 s of scanning each second, so a ten-layer build takes about 5 s. The run stops itself at the end of the build. Make every change before pressing Run.

(a) Run as loaded and read Build defects. *Checkpoint:* 4.10 %.

(b) Press the reload button, set Feedback gain to 1 and run. Repeat at 2. *Checkpoint:* 2.57 % and 2.93 %.

(c) Press the reload button, set Spatter on the photodiode (1σ) to 10 %, and run at a gain of 3. Then press Reset, set the gain to 1 and run again. Explain the gap between the two. *Checkpoint:* 8.25 % and 2.66 %.

Accept Build defects within 0.02 percentage points.

10.8 App challenge · 9 min

Load scenario 4, Learn from the last layer, which runs ten layers with Feedback off and Learning (ILC) on at a learning gain of 0.5. Use the default seed and a speed of 0.1. After each run export the CSV file: a layer's scores sit in the columns `layer_rms_um`, `layer_defect_pct` and `layer_lof_pct` of the first row whose `layer` column shows the next layer, and those of layer 10 in the last row.

(a) Run as loaded. Read the defects of layers 1, 2 and 3 and the rms error of layer 10. *Checkpoint:* 4.69, 1.84 and 0.00 %; 2.98 µm.

(b) Press Reset, set Learning gain to 1 before Run and run. Read the lack of fusion on layer 2 and the rms error of layer 3. Which layer of the first run first had a lower rms error than that? *Checkpoint:* 2.10 %, 3.44 µm; layer 8.

(c) Press Reset, set Learning gain to 1.5 and run. Read the defects of layers 6 to 10. Then set Spatter on the photodiode and Photodiode noise to zero, press Reset and run again. *Checkpoint:* 0.97 to 3.10 %; no defects from layer 3.

Accept rms errors within 0.02 µm and shares within 0.02 percentage points.

10.9 App challenge · 8 min

Load scenario 5, Without the Q-filter, which runs 25 layers of learning alone with Q-filter length at 0. Use the default seed and a speed of 0.25, so a build takes about 5 s. Read each layer's scores from the exported CSV file, as in the previous exercise.

(a) Run as loaded. Find the layer with the lowest rms error and its value, and the rms error and defect share of layer 25. *Checkpoint:* Layer 4, 5.34 µm; 17.03 µm and 24.79 %.

(b) Press Reset, set Q-filter length to 0.2 mm before Run and run. Repeat at 0.6 mm. Read the rms error of layer 25 and the layers that still have defects. *Checkpoint:* 2.83 µm, defects on layers 1 and 2 only; 4.70 µm, defects through layer 8 and on layers 19 and 24.

(c) Press Reset, set Q-filter length back to 0 and Learning lead (under Advanced) to 0.2 mm, and run. Read the rms error of layer 25 and explain why the shorter lead converges without a filter. *Checkpoint:* 1.65 µm.

Accept rms errors within 0.02 µm and shares within 0.02 percentage points.

10.10 Application area · 3 min

A quality manager proposes recording the defect share the controller reports for each layer as the porosity of the part, for its qualification file. Give two reasons from the chapter why that number cannot serve as porosity, and one thing that would have to be added.

10.11 Application area · 3 min

A machine builder plans to ship iterative learning on its production machines. From scenarios 6, 8 and 9 and the limitations, state three rules for the design and the reason for each.

10.12 Review the agent · 5 min

An AI assistant is asked how to pass the challenge, A build that changes, and writes: “Switch Learning (ILC) on and leave the rest as loaded. Learning removes the corners, the overhang and the heat on every repeated layer, and the default feedback gain of 1 catches the short powder feed. The build will come in under 1 %.” Find the flaw and say how you would check it.