INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 1 Process Dynamics, Modeling, and Identification

CHAPTER 11

# Model Structure Selection and Identifiability on a Temperature Control Lab

Pages 258–283

Edited by John D. Hedengren

**How to cite this chapter.** "Model Structure Selection and Identifiability on a Temperature Control Lab," in *Process Dynamics, Modeling, and Identification*, John D. Hedengren, Ed., vol. 1 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 11, pp. 258–283.

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

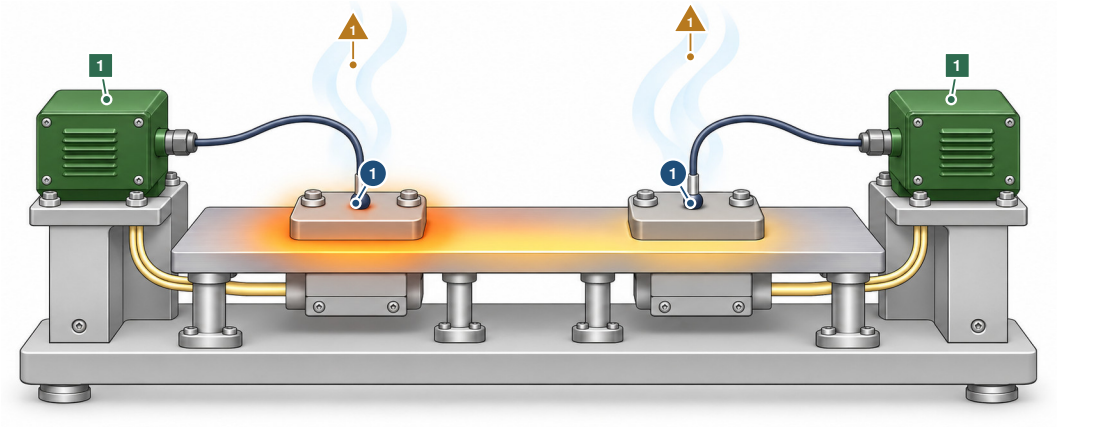
**The app for this chapter**
https://apmonitor.com/apps/tclab-sim/

**The volume online**: contents, other chapters and citation records
https://apmonitor.com/apps/books/v01-dynamics-modeling/#tclab-sim

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

CHAPTER 11

# Model Structure Selection and Identifiability on a Temperature Control Lab




**The controller moves** 1 heater power, set by a stored program, not by feedback

**The controller reads** 1 temperature sensor beside each heater

**What disturbs it** 1 heat lost to the room air

*Notice that the air carries heat away from both tabs while the plate carries it from one tab to the other: the chapter asks which of these paths one recorded temperature can tell apart. The drive boxes stand for the heater switching that sits on the board. Illustration, not to scale.*

### IN PLAIN TERMS Which Model Do the Data Support?

**The main idea.** A small board warms up after a heater turns on. The board, or a computer stand-in for it, records its temperature once a second. Three models of how the board loses heat are fitted to that record, each with more detail than the last.

**The challenge.** A model with an extra term always fits its own record at least as well. It does so even when the heat loss it adds is not there, and only a second test that the fit never saw can catch this.

**Why it matters.** Every controller and computer copy of a device built on a model inherits its mistakes, so the choice of model matters.

### THREE WAYS TO READ THIS CHAPTER

| | | |
|---|---|---|
| **Explore** | 11 min | the first page, the history of the digital twin (Section 11.2), the try-it on the lumped groups and the try-it on Best fit (Section 11.4), the two facts in the same section and the key points |
| **Study** | 55 min | the whole chapter with its three pause points and the twelve exercises at its end |
| **Practice** | 25 min | the first page, the results (Section 11.4), Technology in Practice (Section 11.5), Limitations (Section 11.6) and the open challenge |

**The control question** A heater step test can be explained by more than one energy balance. Which terms of a lumped balance can one record identify, and how does a modeler reject a term that lowers the fitting objective but describes a mechanism the data do not contain?

**Learning objectives**

- Show that the first-order and convective balances are one two-parameter model, and name the parameter groups a step test identifies.
- Compute the steady-state gain and mean lag of the four-state data generator and use them to judge a fitted model.
- Explain why emissivity and the convective coefficient trade along a flat valley over the temperatures a TCLab reaches.
- Test a nested model by hold-out cross-validation on a second step test and read a structured residual as a missing state.
- Explain why an online moving horizon estimate settles away from the whole-record fit when the model form omits the sensor lag.

**Prerequisites** Sections 1.4.2 (energy balance), 1.4.4 (first-order response), 1.4.8 (least squares and the simplex search) and 1.4.9 (validation, identifiability and the information criterion), and Chapter 4 for the board and its four-state model. Section 1.6.3 is needed only for the online estimate.

**The app** https://apmonitor.com/apps/tclab-sim/

**See also** Volume 5, Chapter 11 for a soft sensor of a temperature that no instrument reaches, Volume 5, Chapter 10 for a soft sensor whose bias a delayed laboratory corrects, Volume 5, Chapter 12 for an observer of an unmeasured stored inventory and Volume 5, Chapter 13 for an extended Kalman filter that applies late measurements.



## 11.1 Introduction

The Temperature Control Lab (TCLab) is an Arduino shield that heats two transistors and reads two thermistors, published as a hardware benchmark on the argument that a real device imposes obligations a simulation study never meets, among them a sampling interval that has to be met, traffic on a serial link and a model that is never exactly right [1]. Most of what has been written about the device since then concerns control, from undergraduate tuning exercises to feedforward and cascade structures and the passage from a continuous design to a sampled one, and kits of this kind sit among the hands-on resources that control education had taken up by 2018 [2]. The step response that comes before all of that is usually treated as a preliminary, and the model fitted to it is accepted once the curves look close.

Closeness is the wrong test. Two energy balances that differ in physical content can produce prediction curves that lie within the width of a plotted line over one step test, and the objective function will prefer the one with more parameters whether or not the extra physics is present. The question a modeler actually faces is which structure the data can support, and that question has an answer only when a candidate set is written down, fitted under one objective, and then tested somewhere the fit did not see. Model discrimination in this sense is old work in chemical engineering [3] and is a routine step in estimation practice [4], but it rarely reaches an introductory course because it needs several models, several data sets and a way to compare them quickly.

Two pieces of undergraduate theory carry the whole chapter, and Chapter 1 develops both on a heated tank (Sections 1.4.2 and 1.4.4). A lumped-capacitance energy balance equates the rate of change of stored internal energy to the electrical power dissipated in the transistor minus the heat leaving its surface, and the driving force for that loss is a temperature difference for convection and a difference of fourth powers for radiation [5]. The dynamics that follow are first order, so the response carries exactly two independent pieces of information, a steady-state gain and a time constant, and any parameterization with more than two adjustable symbols is redundant by construction [6]. The principles at work are therefore lumped-capacitance modeling, first-order dynamics and step-test identification, structural identifiability and parameter correlation, the choice of objective function,

input design, cross-validation, and the distinction between fitting a record and choosing a mechanism. The industrial warrant is direct. Every gain error in an identified model is inherited by the controller tuned from it, and a structure chosen for the wrong reason is inherited by everything built on top of it.

This chapter works with a browser application in which a TCLab step test is recorded once and then fitted by hand against three nested energy balances under a single sum-of-absolute-errors objective. It measures what the second and third terms of those balances buy on data whose true structure is known. Lumped thermal modeling, absolute-deviation regression and identifiability analysis are all settled subjects. The chapter writes down the recorder, the three model forms, a re-prediction across the entire horizon on every parameter change, the loss breakdown and a bounded simplex search in enough detail to be rebuilt elsewhere. It then reports objective values, fitted parameters and cross-validation results read off the running application, which locate the boundary between a term the data can identify and a term they cannot. The experiments end with a model that fits better and is rejected.

Section 11.2 places the problem between empirical and physics-based modeling. Section 11.3 carries every constant needed to reproduce the application: the data generator, which is the four-state model that Chapter 4 uses as its twin, the three nested balances (Section 11.3.3), the objective and the search (Section 11.3.5), an online moving horizon estimate of the same parameters (Section 11.3.7) and the hardware path. Section 11.4 measures which parameter groups the data identify, what the radiative term does to the fit and to a second step test (Section 11.4.3), how flat the emissivity valley is, how input design and noise bound the comparison, and how the online estimate compares with the regression (Section 11.4.8). Section 11.5 describes the devices and platforms behind such a record, and Sections 11.6 and 11.7 state what the simulator cannot show and what carries over to plant models.

## 11.2 Background

Process engineering has two traditions of building dynamic models and they answer different questions. The empirical tradition treats the plant as a black box and identifies a transfer function or a difference equation from an input that the engineer designs, which is fast, requires no knowledge of the mechanism, and produces a model that is trustworthy only inside the operating band that the input visited. Four decades of that tradition are surveyed in the system identification literature, along with the recurring difficulty that an identified model is a compromise between what the data contain and what the model set can express [7]. The industrial evidence is that identification, not control design, is where the effort goes, because model development has been reported as the dominant cost of a model predictive control project [8].

> ⧖ HISTORY **From Cybernetics to the Digital Twin, 1948 to 2000**
>
> In 1948 Norbert Wiener defined cybernetics as the scientific study of control and communication in the animal and the machine, a definition that a 2019 review of digital twins links directly to the idea. The term digital twin itself can be traced back only to 2000, when it was mentioned in the context of manufacturing. Several industries had used the idea long before, without the name: online monitoring of process plants, weather forecasting that assimilates measurements into a running model, leak detection in oil and water pipelines, and the remote control and maintenance of satellites and space stations. Data assimilation, the review adds, is in other words the process of fitting models to data, which engineers call filtering and estimation [9].

The physics-based tradition writes conservation equations first and estimates whatever coefficients remain. Its advantage is extrapolation, since a balance that is correct in form stays correct outside the range where its coefficients were measured. Its cost is that the coefficients are often correlated, so the data determine a combination of them rather than each one separately, as Section 1.4.9 shows for the area and level of a tank. Identifiability in this practical sense, meaning whether a particular data set can separate the parameters of a particular model, is normally assessed through the sensitivity

**Table 11.1** Recorded and Derived Signals

| Symbol | Role | Units | Range |
|---|---|---|---|
| $T_1$, $T_2$ | measured output, fitted | °C | unbounded |
| $Q_1$, $Q_2$ | input, programmed or manual | % | 0 to 100 |
| $\hat{T}_1$, $\hat{T}_2$ | model prediction | °C | unbounded |
| $T_a$ | model parameter and generator setting | °C | 10 to 40 |
| $T_{H1}$, $T_{H2}$ | generator state, unmeasured | °C | unbounded |
| $T_{C1}$, $T_{C2}$ | generator state, measured | °C | unbounded |
| $J$ | objective, absolute deviation | °C | $\geq 0$ |
| $T_{\text{cut}}$ | hardware cutoff | °C | 50 to 90 |

matrix and its collinearity, and parameter subsets are chosen so that the remaining directions are distinguishable [10]. The complementary activity is to design the experiment so that the awkward directions become observable, which is the subject of model-based experimental design [11].

The objective function is a modeling decision of the same weight as the model structure. A sum of squared errors, Equation (1.19), is the maximum-likelihood choice for Gaussian noise and gives smooth derivatives, while a sum of absolute errors is far less sensitive to an isolated bad measurement, which is the classical argument for robust estimation [12]. The absolute-deviation form is available as a model option in nonlinear estimation packages and is often preferred for dynamic data reconciliation on plant measurements that contain outliers [13]. The application of this chapter uses the absolute-deviation form throughout, partly for that robustness and partly because a sum of absolute errors in degrees Celsius is a quantity a student can read and reason about without converting units in their head.

Selecting among candidate structures closes the loop. When two models are nested, the larger one cannot fit worse, so the comparison must be made against a penalty for the extra parameter or against data the fit did not see. Information criteria such as Equation (1.23) supply the first route [14], and cross-validation on a second record, as in Section 1.4.9, supplies the second. The measurements in Section 11.4 take the second route, because a step test at a different heater level is available for free and because the failure it exposes is a physical one rather than a statistical score.

## 11.3 Process Models and Data Generator

The physical board holds a pair of TIP31C power transistors and a pair of thermistors, one of each at either end of a small shield. Electrical power enters each transistor as a duty cycle between 0 and 100%, heat leaves both of them to the room, and the board itself carries heat from one end to the other. No thermistor is bonded to the case of the transistor it watches, so heat reaches the reported temperature only after crossing a thermal resistance. Figure 11.1 draws the arrangement and carries every symbol used in the balances below.

### 11.3.1 Signal Classification

The application records two temperatures and two heater commands once per second. Both temperatures are outputs to be matched, and only the first enters the objective unless the second heater is switched into it. Both heater commands are inputs supplied by a stored program or by the operator, and neither is computed from a measurement, because there is no feedback anywhere in the application. The ambient temperature is a parameter of the fitted model and also a setting of the data generator, and those two roles are separate. Table 11.1 lists the classification.

### 11.3.2 The Data Generator

Measurements come either from a physical board or from a four-state simulator that stands in for it. The simulator is deliberately not one of the three models being fitted, so no parameter set can recover

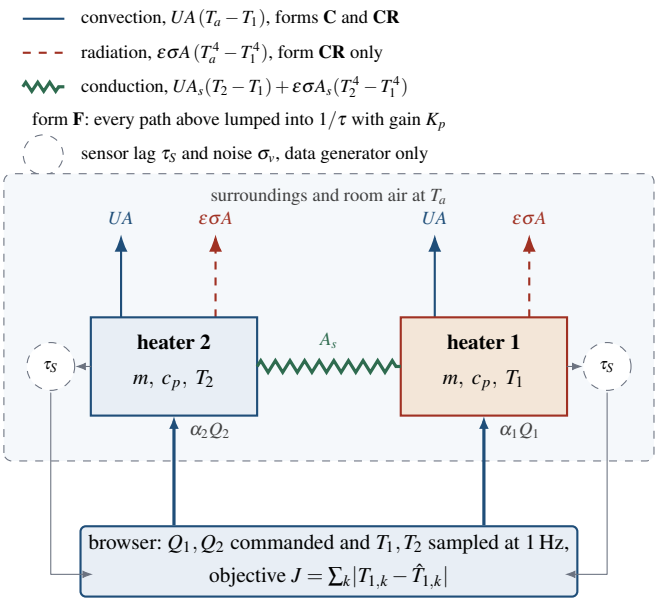



**Figure 11.1** Two heaters on the Temperature Control Lab shield and the three loss paths that the nested energy balances resolve one at a time. The first-order form **F** collapses all of them into a single time constant, the convective form **C** resolves the solid convection path, and the convective plus radiative form **CR** adds the dashed radiation path. Conduction across the board is drawn as a zigzag and the data generator's sensor lags as dashed circles.

the structure that generated the data. It is the twin of Chapter 4 at its default settings. It carries a separate state for each transistor and for each thermistor, as given by Equations (11.1) to (11.4).

$$\frac{dT_{H1}}{dt} = \frac{P_1 Q_1}{C_H} + \frac{T_a - T_{H1}}{\tau_{HA}} - \frac{T_{H1} - T_{H2}}{\tau_{HH}} \tag{11.1}$$

$$\frac{dT_{H2}}{dt} = \frac{P_2 Q_2}{C_H} + \frac{T_a - T_{H2}}{\tau_{HA}} + \frac{T_{H1} - T_{H2}}{\tau_{HH}} \tag{11.2}$$

$$\frac{dT_{C1}}{dt} = \frac{T_{H1} - T_{C1}}{\tau_S}, \qquad \frac{dT_{C2}}{dt} = \frac{T_{H2} - T_{C2}}{\tau_S} \tag{11.3}$$

$$T_i = T_{Ci} + \sigma_v z, \qquad z \sim \mathcal{N}(0,1), \qquad i = 1,2 \tag{11.4}$$

where $T_{H1}$ and $T_{H2}$ hold the two transistor temperatures, $T_{C1}$ and $T_{C2}$ hold the two thermistor temperatures, $P_1$ and $P_2$ scale the available heater power and carry no units, $C_H$ turns that power into a rate of temperature change and is fixed at 5720 % s/°C, $\tau_{HA}$ governs the loss to the room, $\tau_{HH}$ governs conduction from one transistor to the other, $\tau_S$ is the lag of the thermistor behind its transistor and $\sigma_v$ sets the width of the measurement noise. The four states are advanced together with explicit Euler steps of 0.1 s and sampled once per simulated second. A speed control from 1× to 10× takes more substeps per timer tick and leaves the sampled record unchanged, so a ten-minute experiment takes one minute, which no board allows. As in the framework apps of Section 1.6.1, the speed does not change what a run computes. Unlike them, this standalone app draws its noise from the browser's unseeded generator, so two noisy runs differ.

Three properties of this generator matter for everything that follows. Its loss to the room is linear in temperature difference, so radiation is provably absent from the data. Its thermistor lag of 140 s is a state that none of the three fitted models contains, so no parameter set can drive the objective

to zero. Its steady-state gain from $Q_1$ to $T_1$ is 0.599 °C/% at every heater level, obtained by solving Equations (11.1) and (11.2) at steady state with the default parameters, and the mean lag of the same response is 157.6 s, obtained as the sum of the thermistor lag and the 17.6 s first moment of the transistor pair. Those two numbers are the ground truth against which every fit below is judged.

### 11.3.3 Three Nested Energy Balances

The first form is the one an introductory course writes down, a single lumped capacity relaxing toward ambient with a linear gain on heater power, the first-order model of Section 1.4.4, as shown in Equation (11.5).

$$\tau \frac{\mathrm{d}\hat{T}_1}{\mathrm{d}t} = (T_a - \hat{T}_1) + K_p Q_1 \tag{11.5}$$

where $\tau$ is the thermal time constant in seconds, $K_p$ is the steady-state gain in °C/% and $T_a$ is the ambient temperature. Every transport mechanism is inside $\tau$, and nothing in the equation says how the heat leaves.

The second form is an energy balance on a lumped mass whose only loss is convection to the room, as shown in Equation (11.6).

$$mc_p \frac{\mathrm{d}\hat{T}_1}{\mathrm{d}t} = UA\,(T_a - \hat{T}_1) + \alpha_1 Q_1 \tag{11.6}$$

where $m$ is the heated mass in kg, $c_p$ is its specific heat capacity in J/(kg K), $U$ is the overall heat transfer coefficient in W/(m$^2$ K), $A$ is the exposed surface area in m$^2$ and $\alpha_1$ is the electrical power delivered per percent of heater command in W/%. The lumped-capacity assumption behind this form is the standard one, valid when internal conduction is fast compared with the surface resistance [15], and the coefficient $U$ stands for natural convection from a small package in still air [16].

Equation (11.6) is algebraically the same model as Equation (11.5) with $\tau = mc_p/(UA)$ and $K_p = \alpha_1/(UA)$, and the application reports both of those groups as derived quantities so the identity stays visible while the parameters are adjusted. What changes is not the dynamics but the bookkeeping, because the four separate symbols invite a student to give each one a physical value and then discover that only two combinations of them matter.

The third form adds radiation exchange with the surroundings through the Stefan-Boltzmann law, as shown in Equation (11.7).

$$mc_p \frac{\mathrm{d}\hat{T}_1}{\mathrm{d}t} = UA\,(T_a - \hat{T}_1) + \varepsilon\sigma A\,(T_a^4 - \hat{T}_1^4) + \alpha_1 Q_1 \tag{11.7}$$

where $\varepsilon$ is the surface emissivity, $\sigma$ is the Stefan-Boltzmann constant of $5.67 \times 10^{-8}$ W/(m$^2$ K$^4$) and the fourth-power terms are evaluated with absolute temperatures in kelvin. Setting $\varepsilon = 0$ recovers Equation (11.6) exactly, so the three forms are strictly nested, and the reported objective at $\varepsilon = 0$ is a direct check of that nesting. The radiative term is the mechanism usually offered for the observation that the measured steady-state gain of a TCLab falls at elevated temperature [1].

### 11.3.4 The Two-Heater Extension

An option adds the second heater and the heat transferred between the two heaters, and adds $\sum_k |T_{2,k} - \hat{T}_{2,k}|$ to the objective. For the convective plus radiative form the extended pair is given by Equations (11.8) and (11.9).

$$mc_p \frac{\mathrm{d}\hat{T}_1}{\mathrm{d}t} = UA\,(T_a - \hat{T}_1) + \varepsilon\sigma A\,(T_a^4 - \hat{T}_1^4) + UA_s(\hat{T}_2 - \hat{T}_1) + \varepsilon\sigma A_s(\hat{T}_2^4 - \hat{T}_1^4) + \alpha_1 Q_1 \tag{11.8}$$

$$mc_p \frac{\mathrm{d}\hat{T}_2}{\mathrm{d}t} = UA\,(T_a - \hat{T}_2) + \varepsilon\sigma A\,(T_a^4 - \hat{T}_2^4) - UA_s(\hat{T}_2 - \hat{T}_1) - \varepsilon\sigma A_s(\hat{T}_2^4 - \hat{T}_1^4) + \alpha_2 Q_2 \tag{11.9}$$

where $A_s$ is the area that couples the two heaters and $\alpha_2$ is the power factor of the second heater. The coupling terms enter the two balances with opposite signs, so heat moved between the heaters is conserved. The two heaters share $U$, $A$, $m$ and $c_p$, which is a structural constraint rather than a measurement, and Section 11.4 shows what that constraint costs. The convective form drops the two fourth-power terms, and the first-order form couples the two temperatures through a single time constant $\tau_c$ instead, with $\tau(\hat{T}_2 - \hat{T}_1)/\tau_c$ added to the first balance and subtracted from the second, and with a separate gain $K_{p2}$ on the second heater.

### 11.3.5 Objective Function and Numerical Evaluation

The objective is the sum of absolute errors between the measurement and the prediction at every recorded sample, the objective of Equation (1.19) with absolute values in place of squares, as shown in Equation (11.10).

$$J = \sum_{k=1}^{N} |T_{1,k} - \hat{T}_{1,k}| + \delta \sum_{k=1}^{N} |T_{2,k} - \hat{T}_{2,k}| \tag{11.10}$$

where $N$ is the number of recorded samples and $\delta$ is one when the second heater is included in the objective and zero otherwise. The objective is reported in degrees Celsius together with the mean absolute error $J/N$, so a run of 303 samples with an average deviation of 0.47 °C reads as 141.9 °C. No weighting, no normalization and no penalty on the parameters appear anywhere in Equation (11.10).

The prediction is regenerated over the entire recorded horizon whenever any parameter changes, rather than only from the point of the change onward. Each one-second measurement interval is integrated with the classical fourth-order Runge-Kutta method subdivided so that no step exceeds 0.5 s, giving two substeps per interval for the reported prediction and one substep inside the search, where the coarser step halves the cost of an objective evaluation. The heater command is held at the value recorded at the start of the interval, which matches how the firmware and the simulator apply a new setting. The initial condition is the first measured sample rather than the ambient parameter, so a noisy first reading displaces the whole predicted trajectory.

An optional bounded simplex search, the Best fit button, is available as a reference answer. It runs the Nelder-Mead method of Section 1.4.8 with reflection, expansion, contraction and shrink coefficients of 1, 2, 0.5 and 0.5 on parameters rescaled to the unit interval spanned by their adjustable range, which keeps quantities as differently scaled as $A$ at $1.2 \times 10^{-3}$ m$^2$ and $c_p$ at 500 J/(kg K) comparable inside one simplex. The search stops at 2500 objective evaluations, at a simplex smaller than $10^{-6}$, or when the best and worst vertices agree to $10^{-9}$, and it is restarted twice with initial spreads of 0.12, 0.06 and 0.04. It moves two parameters for a single-heater fit, $\{\tau, K_p\}$ for the first-order form and $\{U, \alpha_1\}$ for the other two, and four when the second heater is included. Geometry, material properties, emissivity and ambient temperature are deliberately held, and any fitted parameter that finishes within $10^{-4}$ of the end of its range is reported as a likely sign that the run cannot separate it from another parameter.

**Table 11.2** Parameters of the Fitted Models and of the Data Generator

| Symbol | Meaning | Default | Range | Status |
|---|---|---|---|---|
| *fitted model, first-order form* ***F*** | | | | |
| $\tau$ | time constant, s | 120 | 10 to 600 | fitted |
| $K_p$ | gain, °C/% | 0.8 | 0 to 3 | fitted |
| $K_{p2}$ | gain 2, °C/% | 0.6 | 0 to 3 | $T_2$ only |
| $\tau_c$ | coupling time, s | 1000 | 60 to 4000 | $T_2$ only |
| *fitted model, forms* ***C*** *and* ***CR*** | | | | |
| $U$ | coefficient, W/m$^2$K | 10, 5 | 1 to 40 | fitted |
| $A$ | area, m$^2$ | 0.0012 | 0.0002 to 0.006 | held |
| $A_s$ | coupled area, m$^2$ | 0.0002 | 0 to 0.002 | $T_2$ only |
| $m$ | mass, kg | 0.004 | 0.0005 to 0.02 | held |
| $c_p$ | heat capacity, J/kg K | 500 | 100 to 1500 | held |
| $\alpha_1$ | power factor, W/% | 0.01 | 0 to 0.05 | fitted |
| $\alpha_2$ | power factor 2, W/% | 0.0075 | 0 to 0.05 | $T_2$ only |
| $\varepsilon$ | emissivity | 0.9 | 0 to 1 | held |
| $T_a$ | ambient, °C | 23 | 10 to 40 | seeded |
| $\sigma$ | S.-B. constant, W/m$^2$K$^4$ | $5.67 \times 10^{-8}$ | — | hidden |
| *data generator* | | | | |
| $T_a$ | ambient, °C | 21 | 0 to 45 | setting |
| $P_1$ | heater 1 capacity | 200 | 50 to 255 | setting |
| $P_2$ | heater 2 capacity | 100 | 25 to 255 | setting |
| $\tau_{HA}$ | loss to the room, s | 20 | 5 to 100 | setting |
| $\tau_{HH}$ | heater to heater, s | 100 | 20 to 300 | setting |
| $\tau_S$ | thermistor lag, s | 140 | 20 to 300 | setting |
| $\sigma_v$ | sensor noise, °C | 0.2 | 0 to 1 | setting |
| $C_H$ | capacity constant, % s/°C | 5720 | — | hidden |
| $h$ | integration step, s | 0.1 | — | hidden |

### 11.3.6 Parameters, Ranges and Visibility

Table 11.2 lists every parameter of the fitted models and of the data generator, with the value each takes when the application opens. All thirteen model parameters are operator adjustable, and only those that appear in the selected balance are exposed. Each form keeps its own copy of the whole parameter record, so switching forms never discards a fit in progress.

The default $U$ is 10 W/(m$^2$ K) for the first-order and convective forms and 5 W/(m$^2$ K) for the radiative form, which are the published starting values of the two course assignments the forms follow and already encode the expectation that radiation will carry about half the loss. Ambient temperature is not fitted. At the start of every simulator run the first reading is rounded to 0.1 °C and written into $T_a$ for all three forms, because the board begins at room temperature and a direct reading of a measurable quantity is preferable to an estimate that can trade against the heater gain.

### 11.3.7 Online Moving Horizon Estimate

The simplex search answers once, after the record is complete. A second view, Online estimate, runs the same fit while the record is still arriving, as a moving horizon estimate of exactly the parameters the search moves [17], and draws each of them against time beside the search's answer. Section 1.6.3 describes the estimator options of the framework apps in general. This one is specific to the fit. It never writes a parameter, so the regression remains the reference. Each estimated parameter enters the state in the unit coordinate of its adjustable range, as shown in Equation (11.11).

$$x_k = \left[\hat{T}_{1,k},\ (\hat{T}_{2,k}),\ \tilde{\theta}_k^{\mathsf{T}}\right]^{\mathsf{T}}, \qquad \tilde{\theta}_i = \frac{\theta_i - \theta_i^{\min}}{\theta_i^{\max} - \theta_i^{\min}} \tag{11.11}$$

where $\theta$ holds $\{\tau, K_p\}$ for the first-order form or $\{U, \alpha_1\}$ for the other two, with $\{K_{p2}, \tau_c\}$ or $\{\alpha_2, A_s\}$ added when $T_2$ is included, and $\theta^{\min}$ and $\theta^{\max}$ are the ranges of Table 11.2. Those ranges become the hard bounds $0 \le \tilde{\theta} \le 1$, the box the search itself uses, while the temperatures are unbounded and $A$, $m$, $c_p$, $\varepsilon$ and $T_a$ stay at their set values. Each parameter is a random walk. The temperature rows of the transition $f$ integrate one measured interval of the selected balance with the two Runge-Kutta substeps of the displayed prediction and the heater command held, and $f$ is linearized once per step at the previous estimate $\hat{x}_{k-1}$ by central differences of $10^{-5}$, which gives $A_k = \partial f/\partial x$ and $c_k = f(\hat{x}_{k-1}) - A_k \hat{x}_{k-1}$. At the linearization point the affine step reproduces the integrator to $1.4 \times 10^{-14}$ °C. At every sample the estimator solves Equation (11.12) over the last $L = 30$ intervals, $s = \max(0, k - L)$.

$$\begin{aligned} \min\ & \|x_s - \bar{x}_s\|^2_{\bar{P}_s^{-1}} + \sum_{j=s}^{k-1} \|w_j\|^2_{W^{-1}} + \sum_{j=s}^{k} \frac{\|y_j - C x_j\|^2}{\sigma_y^2} \\ \text{s.t. } & x_{j+1} = A_{j+1} x_j + c_{j+1} + G w_j, \qquad 0 \le \tilde{\theta}_j \le 1 \end{aligned} \tag{11.12}$$

In it $y_j$ holds the recorded $T_1$ (and $T_2$), $C$ selects the modeled temperatures and $G$ drives the parameters alone, so the estimate is an output-error fit like the search. The weights are $\sigma_y = 0.2$ °C and $W = (1/6)^2/3600$ per normalized parameter per step. The arrival term anchors the window at $\bar{x}_s$, the one-step prediction from the estimate made at $s-1$, weighted by a covariance $\bar{P}_s$ that a Kalman recursion propagates through each stored linearization. Both start at the ambient temperature with a spread of 2 °C and at the published parameter values with a sixth of each range. These weights were fixed before any estimate was computed. The value of $\sigma_y$ is the generator's default noise held as a constant of the estimator, never read from the noise setting, and the walk adds the prior's variance again over 3600 steps, the longest run the application records. The window of 30 was the longest of 10, 30 and 60 whose 99th percentile cost per update on the costliest form stayed under 10 ms, a tenth of the sample period at 10×, which was also a rule set beforehand. Equation (11.12) is a dense quadratic program of 61 decisions with one heater and 122 with two, solved by a dual active-set method. It departs from the search in two ways, because it sums squared rather than absolute errors and because its parameters may drift along the record while the search holds one set over the whole of it.

While no bound is active, Equation (11.12) with each step's model fixed on entry has the same solution as the extended Kalman filter on the same linearization, for every window length [17]. A separately written filter matched the unbounded window to $9.6 \times 10^{-9}$ of the range over every update examined, so the unbounded window serves as the Kalman filter in Section 11.4.8, and the window length changes the answer only while a bound is active. A window the solver rejects is replaced by a Kalman update clipped to the bounds and counted, and no measured run rejected one. The estimate consumes the recorded samples in order and nothing else, so it is causal and identical at 1× and 10×, and a record finished before the view was opened is replayed from its first sample. It reads the recorded temperatures and heater commands, the held parameters, the published values and the ranges, and seven altered copies of the source that route the generator's settings or state into it each fail the application's tests.

### 11.3.8 Absence of Feedback and the Overtemperature Comparator

Nothing in the application computes an actuator value from a measurement. There is no setpoint, no automatic mode, no proportional-integral-derivative algorithm and no tuning. Heater commands come from one of four stored programs or from two operator-set inputs, and the stored programs are piecewise constant schedules evaluated as functions of elapsed time alone.

One automatic action does exist and applies only to a physical board. An overtemperature limit, 70 °C by default and settable anywhere between 50 °C and 90 °C while nothing is plugged in, stops the run the moment either board reading reaches it, zeroes both heater commands, writes `X` to the

firmware and latches the alarm. That is a comparator with a latch rather than a controller, and it has no dynamics of its own.

### 11.3.9 Assumptions

The three fitted forms rest on six assumptions, each of which fails somewhere.

The heated object is isothermal. One temperature stands for the silicon die, the package and the board copper attached to it, which is a poor description in the first seconds of a step while the die is still running well ahead of everything bonded to it.

The measured temperature equals the modeled temperature. None of the three forms contains a sensor state, although the data generator does and a real thermistor does, so the whole of that lag is pushed into the fitted parameters. This single assumption accounts for most of the residual reported in Section 11.4.

The heat transfer coefficient is constant. Natural convection from a small package has a coefficient that grows with the temperature difference, so a single $U$ fitted over a 26 °C rise is an average over that rise [16].

The surroundings radiate as a black body at the air temperature. The room air temperature and the temperature of the surfaces the transistor exchanges radiation with are treated as the same number, and the view factor is folded into $\varepsilon$.

Heater power is proportional to duty cycle. The electrical power a pulse-width-modulated transistor actually dissipates depends on its own junction temperature, so the true relation bends slightly as the device heats up.

The measurement noise is white, Gaussian and additive. A ten-bit converter reading a thermistor delivers quantized noise instead, although at $\sigma_v = 0.2\,°\mathrm{C}$ the difference is small.

### 11.3.10 Hardware Interface

The hardware path uses Web Serial directly from the page. The transport requests a port, opens it at 115 200 baud and falls back to 9600 baud, waits 1500 ms for the board to finish its reset, writes `VER` and accepts the connection only if a reply line matches the TCLab firmware pattern. Three commands follow, `Q1 0`, `Q2 0` and `X`, which guarantee that the run begins with both heaters off. Each second of a run it reads `T1` and `T2`, then writes the heater command that holds until the next reading, with a 1800 ms timeout on every reply and all operations serialized through a promise chain. Cable removal, a closed stream, a command timeout, the page being hidden and the page being unloaded each trigger a best-effort `Q1 0`, `Q2 0` and `X`, and the watchdog compiled into the firmware backstops all of them. Hardware always runs at 1×, because a board cannot be accelerated. The measurements reported in this chapter were taken from the simulator, since the hardware path could not be exercised in the environment used for the experiments.

## 11.4 Results

Five recorded runs support the measurements below and are listed in Table 11.3. Every run was taken from the built-in simulator at 10× speed with the default generator parameters, and the sensor noise was set to zero except where stated, so that each objective value is reproducible to the digit displayed. No thermistor allows zero noise, but a noise-free run shows the structural residual on its own and gives the reproducible objective values that a comparison between models needs. Objective values and fitted parameters were read from the application's own readouts, and the trajectories plotted in Figure 11.3 were regenerated from Equations (11.6) and (11.7) using the exported parameter header, which reproduces the displayed objective to 0.01 °C.

The application organizes its experiments as thirteen scenarios, numbered and named as on its Instructions page and its progress badge: 1, Record and fit the first-order form; 2, Prove the lumped groups; 3, Add radiation; 4, Check the $\varepsilon = 0$ collapse; 5, Run Best fit, then distrust it; 6, Find the identifiability valley; 7, Fix it with a better experiment; 8, Read the residual's shape; 9, Two heaters; 10, Use the 100% step; 11, Fit real hardware; 12, Export and re-fit elsewhere; and 13, Watch the fit

**Table 11.3** Recorded Runs

| Label | Heater program | Length | Samples |
|---|---|---|---|
| E1 | $Q_1$ 0 to 50% at 2 s | 302 s | 303 |
| E2 | $Q_1$ 0 to 100% at 2 s | 302 s | 303 |
| E3 | manual, $Q_1$ 80% at 2 s, 30% at 300 s | 602 s | 603 |
| E4 | doublet, $Q_1$ 80, 30, 60% with $Q_2$ 40% | 600 s | 601 |
| E5 | $Q_1$ 100% at 10 s, $Q_2$ 100% at 300 s | 600 s | 601 |

converge online. Scenarios 11 and 13 are optional, so eleven are required. A scenario is credited for a program recorded to its end, the stated model form or parameter change, and a confirmation where the reading cannot be inferred from the record. Selecting a scenario alone earns nothing. Two credits are weaker than their names. Scenario 3 records that the predicted loss split was read, which cannot establish radiation in a simulator whose balance contains none. Scenario 4 records a confirmation without checking that the other parameters match the convective form, and the collapse to Equation (11.6) is exact only when they do. The records show that an activity was done, not what was learned.

### 11.4.1 Structural Identifiability of the Convective Form

The convective form carries six symbols and the data can move only two combinations of them. Doubling $A$ from 0.0012 $\mathrm{m}^2$ to 0.0024 $\mathrm{m}^2$ while halving $U$ from 8.46 to 4.23 $\mathrm{W/(m^2\,K)}$ leaves the objective on E1 at 141.9 °C, unchanged in every displayed digit, and doubling $m$ from 0.004 to 0.008 kg while halving $c_p$ from 500 to 250 J/(kg K) does the same. Only the products $UA$ and $mc_p$ are identifiable, exactly as the derived readouts $\tau = mc_p/(UA)$ and $K_p = \alpha_1/(UA)$ predict. Doubling the fitted $\alpha_1$ alone instead raises the objective from 141.9 to 5039.9 °C, which shows that the power factor is not part of either group.

> **TRY IT IN THE APP** **2 · Prove the Lumped Groups**
>
> Set the sensor noise to 0 in Settings and the speed to $10\times$, then record the Q1 = 50% · 5 min test to its end. Choose the Convective form, double $U$ and halve $A$ from their published values, and watch the objective and the Implied $\tau$ and Implied gain readouts. None of them moves, because only the products $UA$ and $mc_p$ enter the prediction.

The same identity appears between the two smallest model forms. Fitting the first-order form to E1 returns $\tau = 197$ s and $K_p = 0.676\,^\circ\mathrm{C}/\%$ with $J = 141.9\,^\circ\mathrm{C}$, and fitting the convective form to the same run returns $U = 8.46\,\mathrm{W/(m^2\,K)}$ and $\alpha_1 = 0.006\,86\,\mathrm{W}/\%$, whose derived groups are 197 s and 0.676 °C/% with the identical objective. Two independent simplex searches over different symbols reach the same model, which is the cleanest demonstration in the application that renaming parameters is not the same as adding information.

**Pause.** Try Exercise 11.1 before reading on.

### 11.4.2 Effect of the Radiative Term on the Fitted Run

Adding the Stefan-Boltzmann term reduces the objective on E1 from 141.9 to 134.2 °C, an improvement of 5.4%, and the maximum single-sample error falls from 1.67 to 1.63 °C. The radiative form introduces a nonlinear loss mechanism, but its automatic search still moves only two coefficients because emissivity is held fixed. The smaller residual therefore does not establish that radiation generated the data, which contain only linear loss. The parameters that deliver the improvement carry the argument. The heat transfer coefficient falls from 8.46 to 2.46 $\mathrm{W/(m^2\,K)}$, the derived time constant grows from 197 to 679 s, and the loss breakdown reports that 70.1% of the energy leaving heater 1 over the run went out by radiation. The truth is 0%.

**Table 11.4** Objective After Fitting, by Model Form and Recorded Run

| Run | Published **C** | Fitted **F** | Fitted **C** | Fitted **CR** |
|---|---|---|---|---|
| E1 | 1876.4 | 141.9 | 141.9 | 134.2 |
| E2 | 3752.8 | 283.8 | 283.8 | 250.2 |
| E3 | – | 567.1 | 567.1 | 658.3 |
| E4 | – | 594.0 | 594.0 | 653.9 |
| E5 | – | 963.2 | 963.2 | 1211.2 |

All entries are $J$ in °C. E5 includes $T_2$ in the objective, and the other four do not.

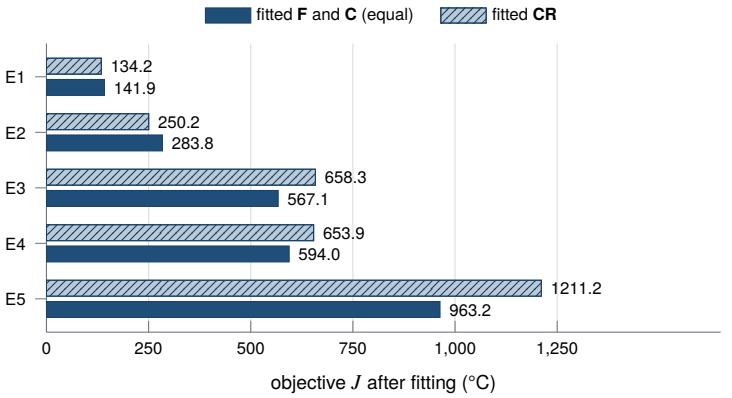



**Figure 11.2** The objective after fitting, by model form and recorded run, from Table 11.4. Solid bars are the first-order and convective forms, which reach the same value on every run, and hatched bars the convective plus radiative form. Notice that the radiative term wins only on the two single steps E1 and E2 and loses on the two-level input E3, the doublet E4 and the two-heater run E5.

Results for all five runs are collected in Table 11.4. The radiative form improves both step-test fits, by 5.4% on E1 and 11.8% on E2. It loses by 16.1% on the two-level input E3, by 10.1% on the doublet E4, and by 25.7% on the two-heater run E5, where the search drives $U$ onto its lower bound of 1 W/(m$^2$ K) and the application flags the parameter as unidentifiable in its status line. Figure 11.2 draws the same results as bars.

> **TRY IT IN THE APP** **5 · Run Best Fit, Then Distrust It**
>
> With the sensor noise at 0, record the Two-heater · 10 min test, tick Include T2 and choose the Convective + radiative form. Press Best fit and read the status line. The objective ends at 1211.2 °C, worse than the convective 963.2 °C, and $U$ finishes on its lower bound of 1 W/(m$^2$ K), which the status line reports as a parameter the run cannot separate.

### 11.4.3 Cross-Validation Between the Two Step Tests

The clearest evidence sits in the pair of step tests at 50% and 100% heater power. Both models were fitted to E1 and then applied without change to E2, and results are given in Table 11.5. The convective parameters transfer perfectly, in the sense that a fresh search on E2 cannot improve on them and reports so. The objective doubles from 141.9 to 283.8 °C only because the temperature rise doubles, from 25.95 °C to 51.90 °C at the end of the run, and the measured ratio of 2.000 is itself the signature of a linear loss. The radiative parameters do not transfer. Applied to E2 they give 429.0 °C against 283.8 °C for the convective set, which is 51% worse, and the residual in Figure 11.3 shows why. The fitted radiative model runs too slowly and stays 1.49 °C below the measurement for the whole second

**Table 11.5** Cross-Validation Between the Two Step Tests

| Parameter set | $J$ on E1 | $J$ on E2 | $\hat{T}_1$ at $Q_1 = 100\%$ |
|---|---|---|---|
| **C** fitted on E1 | 141.9 | 283.8 | 88.56 |
| **CR** fitted on E1 | 134.2 | 429.0 | 80.48 |
| **CR** fitted on E2 | 209.7 | 250.2 | 84.47 |
| data generator | – | – | 80.94 |

$J$ in °C, steady state in °C. The convective set fitted on E2 is the set fitted on E1, because the search on E2 could not improve on it.

half of the run, while the convective residual stays inside 0.92 °C over the same window. Fitting on one step and evaluating on an untouched second step is hold-out cross-validation, and it is what makes the degradation from 283.8 to 429.0 °C a measurement rather than an anecdote.

🕮 **Pause.** Try Exercise 11.2 before reading on.

The steady-state column of Table 11.5 carries the opposite verdict. Extrapolated to a full-power steady state, the convective fit predicts 88.56 °C against the generator's true 80.94 °C, an error of 7.62 °C, while the radiative fit predicts 80.48 °C, an error of 0.46 °C. The mechanism is visible in the chord gain of the fitted radiative model, which falls from 0.655 to 0.595 °C/% between $Q_1 = 50\%$ and $Q_1 = 100\%$ while the generator holds 0.599 °C/% at both. A falling curve crosses a constant one somewhere, and E2 happens to be near the crossing. The radiative term therefore improves an extrapolation that the data never constrained and worsens a transient that they did, which is the practical shape of overfitting on a physical model.

✦ **DID YOU KNOW? Better on Its Own Record, Worse on the Next**

From the simulator: the radiative form fitted to the 50% step lowers the objective from 141.9 to 134.2 °C. Carried unchanged to the 100% step, it scores 429.0 °C against 283.8 °C for the convective set, 51% worse (Table 11.5).

### 11.4.4 Emissivity Scan

Emissivity is excluded from the search on purpose, and a scan explains the decision. Setting $\varepsilon$ by hand and refitting $U$ and $\alpha_1$ at each value gives Table 11.6. Across the whole admissible range of $\varepsilon$ from 0 to 1 the objective moves from 141.9 to 133.4 °C, a total of 6.0%, while the fitted heat transfer coefficient sweeps a factor of 4.7 from 8.46 down to 1.79 W/(m$^2$ K) and the derived time constant sweeps the same factor from 197 to 931 s. A student who reads only the objective sees a monotone and almost featureless slope. The reported loss breakdown at the same time swings from 100% convective to 78.1% radiative, so the application reports a completely different account of the physics for a 6% change in fit quality. The entry at $\varepsilon = 0$ reproducing the convective objective of 141.9 °C exactly is the numerical check that the two forms are nested.

The physical reason the valley is flat is the temperature range. Over a rise from 21 to 47 °C the radiative driving force $T_a^4 - T_1^4$ is close enough to linear in $T_a - T_1$ that the two loss terms are nearly proportional, so a reduction in $U$ can be traded against an increase in $\varepsilon$ almost exactly. Doubling the step to E2 barely helps, because at the published values the instantaneous radiative share at the end of the run moves only from 55.1% to 58.6% when the final temperature goes from 53.7 to 83.1 °C. Separating the two mechanisms needs a temperature range that a TCLab is not permitted to reach, which is a statement about the experiment rather than about the physics.

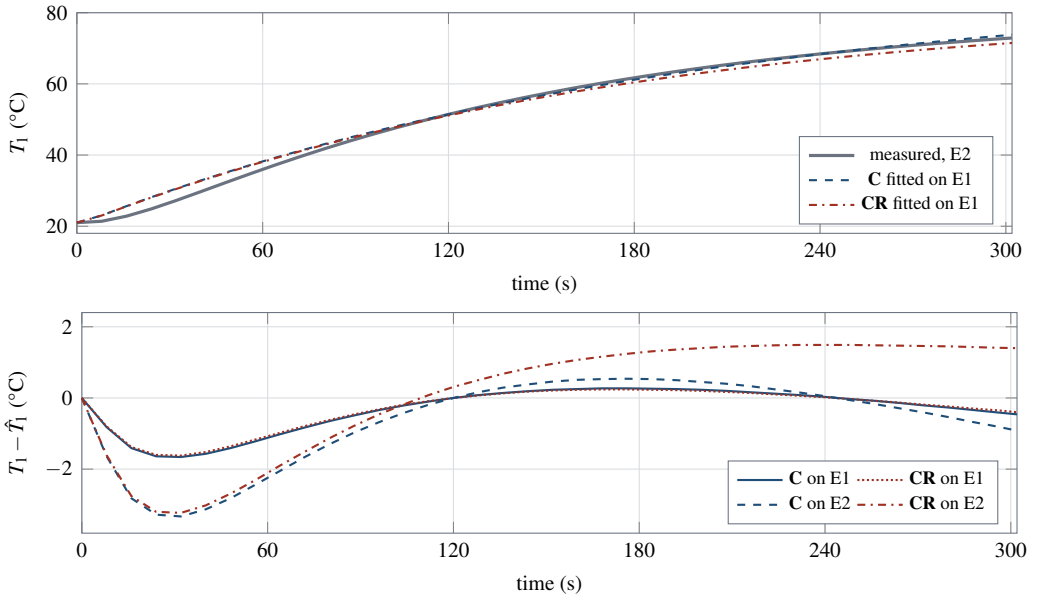



**Figure 11.3** A fit made on the 50% step test carried across to the 100% step test. The upper axis plots the measured E2 response (thick gray) against the two predictions, which were fitted to E1 and then left alone: convective dashed, radiative dash-dotted. The lower axis plots the signed residual of both models on both runs: on E1 convective solid and radiative dotted, on E2 convective dashed and radiative dash-dotted. The persistent 1.49 °C bias of the radiative form on E2 is the excess that Table 11.5 counts.

**Table 11.6** Emissivity Scan on Run E1, Refitting $U$ and $\alpha_1$

| $\varepsilon$ | $U$ (W/(m$^2$ K)) | $mc_p/(UA)$ (s) | radiative share (%) | $J$ (°C) |
|---|---|---|---|---|
| 0.0 | 8.46 | 197 | 0.0 | 141.9 |
| 0.2 | 7.12 | 234 | 15.2 | 140.2 |
| 0.4 | 5.78 | 288 | 30.6 | 138.5 |
| 0.6 | 4.45 | 374 | 46.2 | 136.8 |
| 0.8 | 3.12 | 534 | 62.1 | 135.1 |
| 0.9 | 2.46 | 679 | 70.1 | 134.2 |
| 1.0 | 1.79 | 931 | 78.1 | 133.4 |

> ✦ **DID YOU KNOW?** **A Flat Valley With Two Different Stories**
>
> From the simulator: sweeping the emissivity from 0 to 1 on E1 moves the objective by only 6.0%, while the fitted heat transfer coefficient falls by a factor of 4.7 and the reported radiative share climbs from 0.0 to 78.1% (Table 11.6).

### 11.4.5 Ambient Temperature and Input Design

Ambient temperature is the other parameter held out of the search, and the reason is the same trade in a different pair. Scanning $T_a$ by hand and refitting gives Table 11.7. On the single step E1 a 6 °C error in $T_a$ costs 2.3% of objective while moving the fitted power factor by 18%, so a search allowed to move $T_a$ travels a long way for very little reduction in the objective. On the two-level input E3 the same 6 °C window costs 63% of objective, which is the sensitivity that makes the parameter

**Table 11.7** Ambient Temperature Held at a Wrong Value, Convective Form Refitted

| $T_a$ | run E1, one heater level | | run E3, two heater levels | |
|---|---|---|---|---|
| (°C) | $\alpha_1$ (W/%) | $J$ (°C) | $\alpha_1$ (W/%) | $J$ (°C) |
| 19 | 0.0073 | 140.9 | 0.0077 | 476.7 |
| 21 | 0.0069 | 141.9 | 0.0075 | 567.1 |
| 25 | 0.0060 | 144.1 | 0.0072 | 778.2 |

identifiable. An input that visits two heater levels supplies two independent relations between $T_a$ and $\alpha_1$, while an input that holds one level supplies only the combination $UA\,T_a + \alpha_1 Q_1$.

Sensitivity is not the same as accuracy. On E3 the objective keeps falling as $T_a$ is reduced, reaching 354.6 °C at 15 °C, which is below any plausible room temperature. The minimum of the objective over $T_a$ sits outside the physical range because the fitted forms have no sensor state, and the search buys a better match to the lagged measurement by pretending the room is cold. Seeding $T_a$ from the first reading of the run avoids the whole question, and it costs nothing because the board always starts at room temperature.

The two-level input also improves the fitted gain. Fitting the convective form to E1 returns a gain of 0.676 °C/% against the generator's true 0.599 °C/%, an error of 12.8%, because a five-minute step reaches only 86.6% of the asymptotic rise and the search has to guess the rest. Fitting the same form to E3 returns 0.617 °C/%, an error of 3.0%, and the doublet E4 returns 0.619 °C/%. The identified time constant improves in the same way, from 197 s on E1 to 164 s on E3 against the generator's 157.6 s mean lag.

### 11.4.6 The Noise Floor

Every number above was taken with the sensor noise disabled. Four further recordings use the default 0.2 °C noise and independent, reproducible pseudorandom streams. Their convective objectives are 134.6, 159.8, 161.1 and 137.4 °C, with a mean of 148.2 °C and a sample standard deviation of 14.2 °C. The automatic searches use the ambient reading taken before recording and initialize the prediction from the first stored sample. Those are different noisy readings in the simulator.

The paired radiative fits reduce the objectives by 6.7 to 7.0 °C on these same four records. This repeatable reduction is smaller than the variation between recordings, but that comparison alone does not make the paired improvement a noise artifact. The noise-free result and the hold-out test establish the relevant limitation: a smaller training objective does not verify the added mechanism. Across the four convective fits, $U = 8.39 \pm 0.28\,\mathrm{W/(m^2\,K)}$ and $\tau = 199 \pm 7\,\mathrm{s}$, where the spreads are sample standard deviations rather than confidence intervals.

A low initial reading shifts the entire predicted trajectory and can cancel part of the structural residual. The first record starts 0.33 °C below the noise-free initial temperature and fits to 134.6 °C, below the noise-free 141.9 °C. Ambient seeding provides a second source of variation. The deployed application uses unseeded browser randomness, so the exact four noisy records are a reproducibility experiment rather than a fixed sequence that every learner will obtain.

### 11.4.7 Two Heaters and One Shared Coefficient

Including $T_2$ in the objective adds two parameters and two constraints at once. On run E5 the convective form reaches 514.6 °C on $T_1$ and 448.6 °C on $T_2$ for a total of 963.2 °C, with $\alpha_2 = 0.0031\,\mathrm{W/\%}$ against $\alpha_1 = 0.0076\,\mathrm{W/\%}$ and $A_s = 0.000\,43\,\mathrm{m^2}$. The ratio $\alpha_2/\alpha_1$ of 0.41 underestimates the generator's true capacity ratio $P_2/P_1$ of 0.50, and the derived time constant rises to 213 s, both of which are the price of forcing one $U$, one $A$ and one $mc_p$ to serve two heaters that the generator gives different capacities. The radiative form on the same run reaches 1211.2 °C, 25.7% worse, and finishes with $U$ against its lower bound, which the application reports in words rather than leaving the student to notice.

**Table 11.8** Final Error of the Online Estimate Against Its Reference

| Run | Form | KF | $L = 10$ | $L = 30$ | $L = 60$ | noisy |
|---|---|---|---|---|---|---|
| S1 | **C**, $T_2$ | 0.05 | 0.05 | 0.05 | 0.05 | 0.93 |
| S2 | **CR** | 0.48 | 0.48 | 0.48 | 0.48 | 1.12 |
| E1 | **F** | 14.43 | 14.43 | 14.43 | 14.43 | 9.35 |
| E1 | **CR** | 9.91 | 1.42 | 0.95 | 1.20 | 3.39 |
| E4 | **C** | 10.13 | 9.90 | 9.93 | 9.90 | 10.59 |
| E5 | **CR**, $T_2$ | 71.53 | 6.39 | 6.29 | 6.20 | 6.52 |
| E5 | **C**, $T_2$ | 65.70 | 35.55 | 27.41 | 27.45 | 25.89 |
| $E1_{20}$ | **F** | 0.57 | 0.57 | 0.57 | 0.57 | 6.42 |

Largest final error over the estimated parameters, in percent of each range, against the truth (S1, S2) or the search's answer (the rest). KF is the unbounded window. The last column is $L = 30$, the mean of three seeds at 0.2 °C, and the others are noise-free. $E1_{20}$ is E1 with $\tau_S = 20$ s.

**Table 11.9** Final Estimate When One Candidate Cause Is Changed

| | whole record | | walking $\theta$ | | constant $\theta$ | |
|---|---|---|---|---|---|---|
| Run, value | $\lvert e\rvert$ | $e^2$ | window | exact | window | exact |
| E1 **F**, $\tau$ | 197.1 | 221.1 | 111.9 | 10.0 | 146.1 | 174.4 |
| E1 **CR**, $U$ | 2.456 | 1.551 | 2.826 | 1.513 | 3.588 | 3.522 |
| E4 **C**, $U$ | 10.06 | 9.739 | 6.188 | 6.082 | 10.43 | 10.57 |
| E5 **C**$T_2$, $U$ | 7.83 | 7.386 | 1.758 | 1.777 | 5.026 | 8.144 |

Noise-free. Whole record: the search on absolute or on squared errors. Window: Equation (11.12) in one pass with $L = 30$. Exact: the same problem over the whole record, relinearized to convergence. Constant: $W = 0$. $\tau$ in s, $U$ in W/(m$^2$ K).

### 11.4.8 Online Estimate Beside the Regression

The moving horizon estimate of Section 11.3.7 was run on six records, with the answer of the search, or the truth where it is known, as the reference in every comparison. Three are the simulator runs E1, E4 and E5, which carry the 140 s thermistor lag that no fitted form contains, and a fourth repeats E1 with that lag at 20 s. Two are synthetic records generated by the fitted balances themselves, the doublet input with the convective form and $T_2$ at $U = 9$, $\alpha_1 = 0.008$, $\alpha_2 = 0.004$ and $A_s = 0.0004$ (S1), and the full-power step with the radiative form at $U = 6$ and $\alpha_1 = 0.009$ (S2). Each record was taken noise-free and with 0.2 °C of noise on three seeds, the truth was read by the measurement script alone, and the estimator ran unchanged on all of them. Table 11.8 lists the final errors.

On the synthetic records, where the selected form is the generating one, every estimator ends within 0.05% (S1) and 0.48% (S2) of each range of the truth without noise and within 0.93% and 1.12% on average with it, while the search ends within 0.13% and 0.11% with noise. The second power factor shows when a parameter is not yet excited. On S1, $\alpha_2$ stays at its published 0.0075 W/% until heater 2 switches on at 450 s and is within 2% of its range of the truth by 452 s.

On the simulator's records the estimate settles away from the search. The first-order time constant ends at 111.9 s against 197.1 s on E1, the convective $U$ at 6.19 against 10.06 W/(m$^2$ K) on E4, and the convective $U$ on E5 with $T_2$ at 1.76 against 7.83. The last estimate is not a whole-record model either. Held constant over E1 it gives $J = 866.1$ °C against the search's 141.9 °C, and over E4 it gives 1918.8 °C against 594.0 °C.

Changing one candidate cause at a time locates the difference, as listed in Table 11.9. The missing structure is the precondition. With the generator's thermistor lag set to 5, 20, 50, 80 and 140 s, the first-order $\tau$ of the window ends 0.2, 0.6, 3.2, 5.6 and 14.4% of its range from the search's, and with the lag at its 20 s minimum the two agree within 0.57% noise-free, although not under noise (6.42%). The random walk carries the difference. Solving the estimator's own problem exactly over the whole of E4 gives $U = 6.08$ against the window's 6.19, so the one-step linearization costs 0.3%

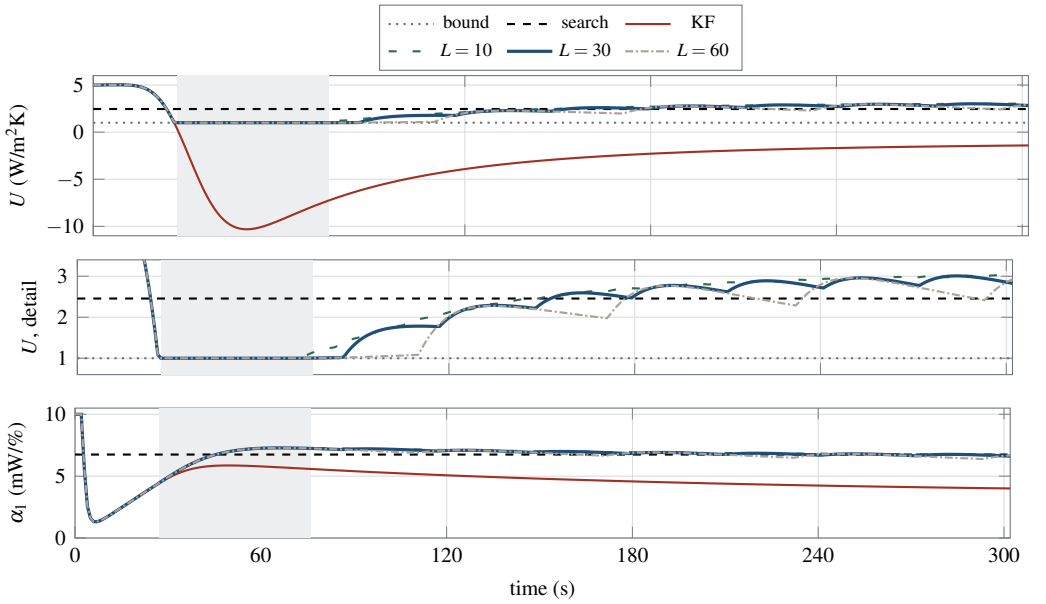



**Figure 11.4** Online estimates of $U$ and $\alpha_1$ on the noise-free 50% step E1 with the radiative form. The unbounded Kalman filter (KF) takes $U$ to $-10.3$ while the three windows hold it on its lower bound, shaded from 27 to 76 s for $L = 30$. The middle axis enlarges the windows. KF is thick solid, $L = 10$ loosely dashed, $L = 30$ very thick solid and $L = 60$ dash-dotted. The straight dashed lines are the search's answer and the dotted line is the bound.

of the range, while holding the parameters constant brings $U$ to 10.43 in one pass and 10.57 solved exactly, within 1.3% of the search. With the walk off, the doublet's $U$ stays within 0.4 to 1.2% of the search's at lags from 50 to 140 s, against 9.9 to 24.9% with it. Inside the window estimate $\alpha_1$ falls from 0.0052 to 0.0016 W/% during the 30% stretch of the doublet between 300 and 450 s and then recovers, so the parameters drift to absorb the lag the form leaves out, and the last value describes the end of the record rather than the whole of it. The objective is a small part, because refitting the whole record by squares moves the answer by 0.8 (E4) to 4.1% (E1, first order) of the range. On the first-order form the walking problem is degenerate. Solved exactly it holds $\tau$ on its 10 s bound for 68% of the record and lets $K_p$ grow from 0.013 to 0.52 °C/%, tracing the lagged rise as a growing gain, so relinearization moves the estimate further from the search rather than closer. Halving the estimator's substep or the generator's Euler step moved no value here by more than 0.02% of a range.

🕮 **Pause.** Try Exercise 11.3 before reading on.

The bound matters on E1 with the radiative form, as shown in Figure 11.4. The unbounded filter drives $U$ below its lower bound of 1 W/(m$^2$ K) in 91% of its updates, reaching $-10.3$ at 50 s, 29% of the range below the bound, and ends at $-1.41$, a negative heat transfer coefficient whose pair gives $J = 1687$ °C over the record. An unbounded filter leaving the physical region in this way is a documented weakness of the extended Kalman filter that bounded moving horizon estimation avoids [18]. The window with $L = 30$ holds $U$ on the bound from 27 to 76 s, 16.5% of its updates, then leaves it, and the three windows end at 3.01, 2.83 and 2.93 for $L = 10$, 30 and 60 against the search's 2.456, with $J = 246.0$ °C for $L = 30$ against 134.2 °C. Under noise the filter ends between $-0.93$ and $-0.54$ and the window between 3.14 and 4.16 on the three seeds. On E5 with the radiative form and $T_2$ the filter ends at $U = -6.65$ and $\alpha_2 = -0.000\,37$ W/%, both outside their ranges, while the window stays inside every range and ends within 6.3% of the search. Where no bound is active the window length does not matter, and the filter and all three windows agree to $10^{-6}$ of the range on

E1 with the first-order form, on S1 and S2, and with the 20 s lag. Where one is active, $L = 10$ and $L = 60$ differ by up to 8.1% (E5, convective), and $L = 30$ and $L = 60$ stay within 1.2% of each other in every case of Table 11.8.

The estimator's covariance exposes the trade that the search reports only through a parameter on its bound. Its Riccati covariance puts the correlation of $U$ and $\alpha_1$ at 0.89 to 0.98 on every record of the physical forms, against 0.54 to 0.59 for $\tau$ and $K_p$ on E1, because each record fixes the steady gain $\alpha_1/(UA)$ far better than $U$ and $\alpha_1$ separately. The application does not display this correlation. On E5 with the radiative form, where the search ends with $U$ on its lower bound, the window rides that bound for 10.5% of its updates early in the run and then leaves it, so it does not repeat the warning of the search.

An update with $L = 30$ cost a median of 1.81 ms and a 99th percentile of 3.55 ms in Node.js on the four-parameter radiative form, and 6.21 ms at the 99th percentile on the convective form with $T_2$, where $L = 60$ reached 43.6 ms, on a computer carrying other work (one-minute load average 1.8). In Chrome on the built application at 10× with the same four parameters, the median was 2.7 ms and the 99th percentile 6.5 ms against the 10 ms budget, with 600 of 600 samples taken, no long task and a load average of 2.7, a margin of only a factor of 1.5. At 1× the 99th percentile was the cold first update, 42.2 ms of a 100 ms budget. The estimate runs only while its view is shown, so the default view costs nothing.

### 11.4.9 Fitted Result in the Application

Figure 11.5 shows the current application after noise-free run E2 with the radiative form fitted. The heat transfer coefficient is 1.56 W/(m$^2$ K) at a held emissivity of 0.9, and the objective is 250.2 °C over 303 points. The integrated loss split is 19.7% convective and 80.3% radiative. A mean error of 0.83 °C therefore accompanies a convective time constant of 1072 s, far above the generator's 157.6 s mean lag, and a radiative loss absent from the data generator. The run and its prediction are computed in the browser and export as a CSV file for independent analysis, under comment lines naming the source, the step test, the model form and each parameter value.

## 11.5 Technology in Practice

This chapter has no loop to close. Its hardware is a measurement path that ends in a record and a fit, and Figure 11.6 draws that path with the devices of the board.

### What Measures It

The data generator reads each thermistor through a first-order lag and white noise, Equations (11.3) and (11.4), while the three fitted forms carry no sensor at all, the assumption that Section 11.4 finds behind most of the residual. A record used to choose between models must first be trusted, and calibration compares a sensor's response with a precisely known input [19]. A zero error would shift every reading and with them the ambient value seeded from the first one. No range or accuracy for the board's sensors is given here, because none was found in an open source that names no maker.

> **TECHNOLOGY** **Sensor: A Temperature Reading Checked Against a Standard**
>
> To calibrate an instrument is to expose it to a precisely known input and adjust its response so the output matches across the whole range. A zero shift moves every reading by the same amount, while a span shift changes the slope, so its error is unequal across the range. Technicians record each calibration as found and as left, because the difference over time measures drift, and excessive drift often signals an impending failure. Water supplies two convenient reference points, its freezing and boiling points, and other substances supply the fixed points of the international temperature scale [19]. For a fitted model, a span error of the sensor would scale the measured rise and with it the identified gain.

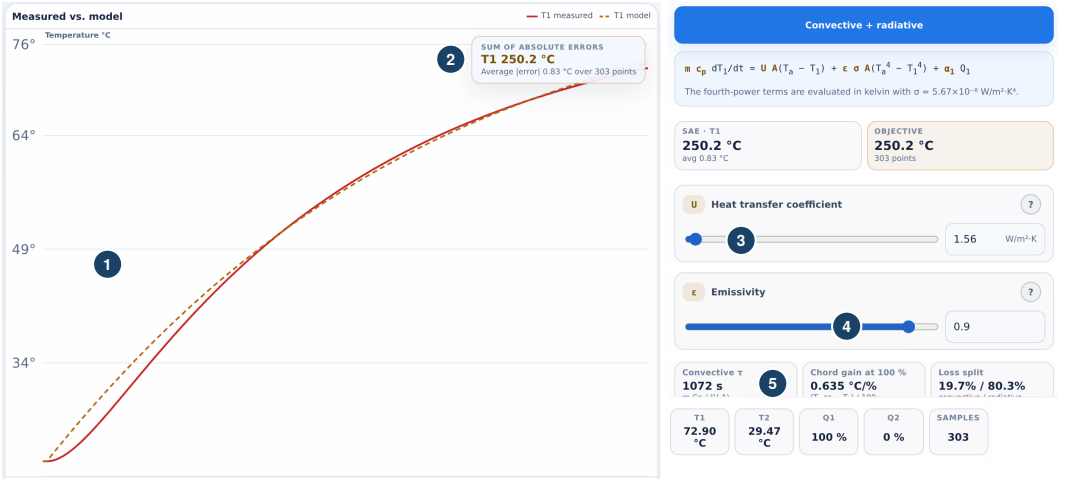



**Figure 11.5** The current application after the noise-free full-power step E2, shown as a composite of six parts of the screenshot: the measured and modeled temperature on the left, and on the right the selected form, the objective tiles, two sliders, the readouts and the measurement tiles of the run. The heater trend and the time axis are cut. Heater 1 steps to 100% at 2 s (Table 11.3), and the final temperature is 72.90 °C at 302 s. The radiative best fit has an objective of 250.2 °C, $U = 1.56\,\mathrm{W/(m^2\,K)}$, and an integrated radiative loss fraction of 80.3%. Callouts: (1) the measured temperature (solid) and the model (dashed), which rises first; (2) the sum of absolute errors and its count of points; (3) $U$ near the low end of its range; (4) the emissivity, held; (5) the readouts of the convective time constant, the chord gain and the loss split.

### What Moves It

Each heater is a transistor whose duty cycle sets its share of full power, the final element that Chapter 4 describes. Here no measurement sets it: a stored program or the operator does, so the design of the input matters more than the device. Industrial test signals must excite both the steady state at low frequency and the dynamics at medium and high frequency, and step-like binary signals are the ones most used [8]. Table 11.7 shows the same need on the board: one heater level ties $T_a$ to $\alpha_1$, and two levels separate them.

> **IN PRACTICE** **Record Where the Model Was Validated**
>
> A government standard for models and simulations, NASA-STD-7009B, was first issued in 2006 and drawn up partly in answer to the findings on the Columbia accident. It requires that a model be validated and that a record of its domain of validation be kept: the region of inputs over which its responses compare favorably with a referent, the real system or an accepted stand-in. Calibration, also called tuning, adjusts parameters to improve agreement with a referent and is a separate act. A model may extrapolate outside the validated domain only within its stated limits, and the aspects not validated are to be recorded [20].

### Where the Controller Runs

The platform is the browser page. Each second it reads both temperatures over the serial link and writes the held heater command (Section 11.3.10), the cycle that a standards glossary defines as reading sensor signals into memory, executing the algorithm and sending the actuator signals [21]. Plant loops run with cycle times from milliseconds to minutes [21], and the 1 s sample sits inside that range, short against the 197 s time constant fitted to E1. The simplex search and the moving horizon

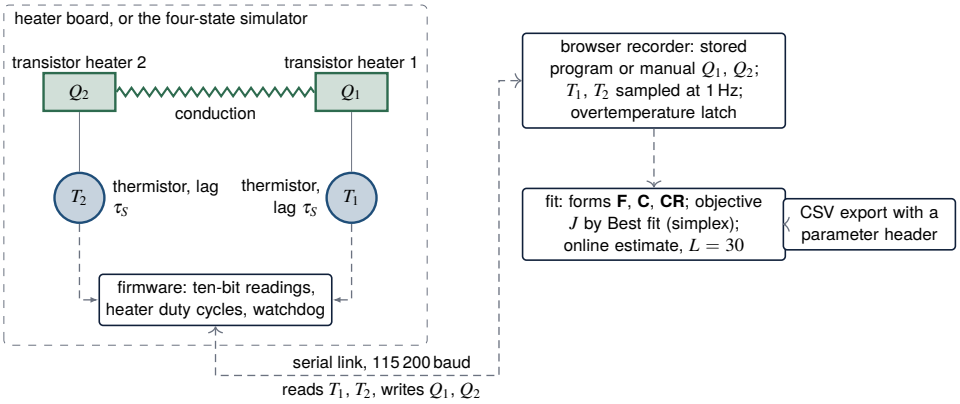



**Figure 11.6** The measurement path of the chapter on the board. Each thermistor sits beside its transistor heater, the firmware converts the readings and applies the duty cycles, and the browser records both temperatures once a second and fits the three forms to the record. Notice that no arrow runs from a temperature back to a heater: the heater commands come from a program, and the only automatic action is the overtemperature latch.

estimate run in the same page. In a plant the record would come from an archive of the control system rather than from the loop itself.

> **TECHNOLOGY** **Platform: A Data Historian**
>
> A data historian is a computer that polls process data from the control system at regular intervals, stores them on a high-capacity drive, and serves them to other computers on request, typically over years. Historians existed in distributed control systems long before network security zones, because the controllers have too little memory to archive process data and plant staff wanted long-term records for running the process and diagnosing problems. In a secured network the historian sits in a buffer zone between the plant and office networks and supplies its records without any direct reading from the controllers [19]. The recorder of this chapter plays the same part at one sample per second, and its exported file names the source, the test, the model form and every parameter.

### What Is Modernizing It

Three developments from the literature on digital twins reach this fit. The first is hybrid modeling: a physics-based model explains part of the data, the interpretable residual it leaves is modeled by an explainable data-driven method, and only what remains goes to a black-box model before the last residual is discarded as noise [9]. The structured residual of Figure 11.3 is such a residual, the trace of the missing sensor state. The second is the automatic recovery of governing equations by sparse regression over many candidate terms, which is fast but performs poorly on data corrupted with noise [9]. The third is data assimilation, routine in weather centers, which combines a model with observations by variational methods or Kalman filters [9]. The moving horizon estimate of Section 11.3.7 applies that idea to two parameters of one heater.

## 11.6 Limitations

The data generator is linear in its loss term, so the negative verdict on the radiative form in Table 11.4 is a verdict about this data set and not about a physical TCLab. Measurements on real boards show a steady-state gain that falls at elevated temperature, and radiation is the usual explanation [1], so the same experiment run on hardware may well rank the three forms differently. The method of this chapter does not depend on which way the ranking falls.

The application supplies no confidence intervals, no sensitivity matrix and no information criterion, so the comparison between models is made on a raw objective value and on cross-validation alone. A collinearity index computed from the sensitivity matrix would quantify the trade that Table 11.6 only exhibits [10], and an information criterion such as Equation (1.23) would price the third parameter directly instead of leaving the noise floor of Section 11.4 to do it [14]. Chapter 12 adds both kinds of evidence for linear models, with a moving-block bootstrap for the uncertainty of each fit and a validation interval that the fit never sees.

The derived time-constant readout is $mc_p/(UA)$ for both physical forms, and the current application labels the radiative value as the convective time constant. It ignores the fourth-power term, so the 1072 s in Figure 11.5 describes the convective path alone. The gain readout is the chord gain from the full balance, 0.635 °C/% for this fit. Using $\alpha_1/(UA)$ instead would give 3.545 °C/%, overstating that gain by a factor of 5.58.

The online estimate is a tracking estimate, not a second regression. Its parameters are random walks, so on a record that no form describes its last value summarizes the end of the record and scores $J = 866.1\,°\mathrm{C}$ over the whole of E1 against the search's 141.9 °C. Its model omits the thermistor lag exactly as the three forms do, linearizes each step once at the estimate of that moment, and carries fixed weights rather than weights fitted to the run. It reports no confidence interval, and the parameter correlation it computes is not displayed.

The three fitted forms share one heat transfer coefficient, one area and one thermal capacity between the two heaters, so a board whose two channels differ cannot be matched exactly at any parameter setting.

The simplex search moves two parameters, or four with the second heater included, and holds the rest. It is a reference answer rather than an estimator, and it carries no global-search strategy beyond two restarts from progressively smaller simplexes, so a badly placed starting point can leave it in a local minimum.

The hardware path was documented from the source code and was not exercised, because no physical board was attached during this work. Reading a board needs Web Serial, which desktop Chromium-family browsers expose and other browsers do not, and which a page reached over plain HTTP cannot use at all.

## 11.7 Discussion

A parameter set is a property of one data record, while a model structure is a claim about mechanism that every later use inherits. The distinction becomes expensive downstream. A controller tuned from an identified model inherits the model's gain error directly, and a predictive controller inherits its extrapolation as well, which is why model development rather than control design dominates the cost of an industrial application [8]. A digital twin makes the same inheritance permanent, because the twin is consulted precisely in the conditions where running the plant would be too expensive, and those are the conditions in which no data were ever taken [22]. Chapter 4 runs such a twin beside the board it models.

The experiment of this chapter is small enough that both outcomes can be seen in one afternoon. A model with an extra physical term fits the data it was given by 5.4% better and predicts a different experiment 51% worse, and the same model extrapolates a steady state to within 0.46 °C where the simpler one is off by 7.62 °C. Neither result is a general law about radiation on transistors. Both are what happens when a term is added to a balance and the data are not asked whether they can support it, which is the habit this chapter sets out to break.

The same choices reappear when the fit runs online. Declaring which parameters an estimator may move is the subset-selection decision that Table 11.6 forces by hand, and an estimator left free to move $\varepsilon$ and $U$ together would wander along the flat valley for the same reason the batch search would. Adding the sensor state that all three forms omit turns the same formulation into a soft sensor

for the transistor temperature, which no thermistor on the board reports. Volume 5, Chapter 11 builds a soft sensor of that kind for a temperature no instrument reaches.

> **OPEN CHALLENGE** **Twins That Extrapolate and Say How Far to Trust Them**
>
> A review of digital twins counts predictive capacity, transparency and generalization among the main technical challenges. Physics-based models generalize well to new problems with similar physics and allow their errors to be bounded, yet they capture only part of the known physics because of the assumptions made along the way. Data-driven models take in long histories but generalize poorly to unseen problems, and their errors cannot be bounded. The review proposes hybrid models that join the two [9]. Progress would be a twin whose predictions beyond its training record carry an error bound that holds when the device is tested there.

## 11.8 Conclusion

This chapter studied an application that records a Temperature Control Lab step test and fits three nested energy balances to it under one sum of absolute errors, with no controller anywhere. The first-order and convective forms are shown to be the same two-parameter model in different symbols, returning identical objectives of 141.9 °C on the 50% step and 963.2 °C on the two-heater run from independent searches, and the products $UA$ and $mc_p$ are shown to be the only identifiable groups by direct substitution. Adding the Stefan-Boltzmann term improves the 50% step fit by 5.4%, to 134.2 °C, while degrading the untouched 100% step from 283.8 to 429.0 °C and reporting a radiative share of 70.1% in data generated with no radiation at all.

An emissivity scan moves the objective by 6.0% across the entire range from 0 to 1 while the fitted heat transfer coefficient moves by a factor of 4.7, and four repeats of the same step test with 0.2 °C of sensor noise give objectives of $148.2 \pm 14.2$ °C. That spread is larger than the improvement the extra term buys, although the paired reduction of 6.7 to 7.0 °C repeats on every record, so it is the hold-out test and not the noise that rejects the term.

A moving horizon estimate of the same parameters agrees with the search within 0.48% of each range on records its own form generated, settles 14.4% of the range away on the simulator's 50% step because its drifting parameters absorb the thermistor lag, and holds $U$ inside its range where an unbounded Kalman filter ends at $-1.41$ W/(m$^2$ K).

A fuller study would add a confidence interval on each fitted parameter, a collinearity index over the selected subset, an information criterion beside the raw objective, a sensor state as a fourth candidate structure, and the same experiments on physical boards. The question the application poses, which model a measurement can support, is answered by a second record and not by the objective of the first.

> **KEY POINTS**
>
> - Which model a measurement can support is decided by a second record, not by the objective of the first.
> - The first-order and convective forms are the same two-parameter model in different symbols, and only the products $UA$ and $mc_p$ are identifiable.
> - Adding radiation improves the 50% step fit by 5.4% while degrading the untouched 100% step from 283.8 to 429.0 °C, and reports a radiative share of 70.1% in data generated with none.
> - A moving horizon estimate settles 14.4% of the range away on the 50% step because its drifting parameters absorb the thermistor lag.
>
> **Where else this applies.** Chapter 12 adds bootstrap uncertainty and a validation interval for linear models, and Volume 5, Chapter 11 builds a soft sensor for a temperature that no instrument reaches.

## References

[1] J. Park, R. A. Martin, J. D. Kelly, and J. D. Hedengren, "Benchmark temperature microcontroller for process dynamics and control," *Computers & Chemical Engineering*, vol. 135, p. 106736, 2020.

[2] J. A. Rossiter, B. Pasik-Duncan, S. Dormido, L. Vlacic, B. Jones, and R. Murray, "A survey of good practice in control education," *European Journal of Engineering Education*, vol. 43, no. 6, pp. 801–823, 2018.

[3] G. E. P. Box and W. J. Hill, "Discrimination among mechanistic models," *Technometrics*, vol. 9, no. 1, pp. 57–71, 1967.

[4] C. Kravaris, J. Hahn, and Y. Chu, "Advances and selected recent developments in state and parameter estimation," *Computers & Chemical Engineering*, vol. 51, pp. 111–123, 2013.

[5] D. E. Seborg, T. F. Edgar, D. A. Mellichamp, and F. J. Doyle, III, *Process Dynamics and Control*, 4th ed. Hoboken, NJ: John Wiley & Sons, 2016.

[6] B. W. Bequette, *Process Control: Modeling, Design, and Simulation*. Upper Saddle River, NJ: Prentice Hall, 2003.

[7] L. Ljung, "Perspectives on system identification," *Annual Reviews in Control*, vol. 34, no. 1, pp. 1–12, 2010.

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

[9] A. Rasheed, O. San, and T. Kvamsdal, "Digital twin: Values, challenges and enablers," arXiv preprint 1910.01719 (version 1); published in revised form as IEEE Access 8, 21980–22012, 2020, 2019.

[10] R. Brun, P. Reichert, and H. R. Künsch, "Practical identifiability analysis of large environmental simulation models," *Water Resources Research*, vol. 37, no. 4, pp. 1015–1030, 2001.

[11] G. Franceschini and S. Macchietto, "Model-based design of experiments for parameter precision: State of the art," *Chemical Engineering Science*, vol. 63, no. 19, pp. 4846–4872, 2008.

[12] P. J. Huber, "Robust estimation of a location parameter," *The Annals of Mathematical Statistics*, vol. 35, no. 1, pp. 73–101, 1964.

[13] J. D. Hedengren, R. Asgharzadeh Shishavan, K. M. Powell, and T. F. Edgar, "Nonlinear modeling, estimation and predictive control in APMonitor," *Computers & Chemical Engineering*, vol. 70, pp. 133–148, 2014.

[14] H. Akaike, "A new look at the statistical model identification," *IEEE Transactions on Automatic Control*, vol. 19, no. 6, pp. 716–723, 1974.

[15] D. W. Hahn and M. N. Özişik, *Heat Conduction*. Hoboken, NJ: Wiley, 2012.

[16] A. Bejan, *Convection Heat Transfer*. Hoboken, NJ: Wiley, 2013.

[17] C. V. Rao, J. B. Rawlings, and D. Q. Mayne, "Constrained state estimation for nonlinear discrete-time systems: stability and moving horizon approximations," *IEEE Transactions on Automatic Control*, vol. 48, no. 2, pp. 246–258, 2003.

[18] E. L. Haseltine and J. B. Rawlings, "Critical evaluation of extended Kalman filtering and moving-horizon estimation," *Industrial & Engineering Chemistry Research*, vol. 44, no. 8, pp. 2451–2460, 2005.

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

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

[21] K. Stouffer, M. Pease, C. Tang, T. Zimmerman, V. Pillitteri, S. Lightman, A. Hahn, S. Saravia, A. Sherule, and M. Thompson, "Guide to operational technology (OT) security," National Institute of Standards and Technology, Gaithersburg, MD, Special Publication 800-82r3, 2023.

[22] A. Rasheed, O. San, and T. Kvamsdal, "Digital twin: Values, challenges and enablers from a modeling perspective," *IEEE Access*, vol. 8, pp. 21 980–22 012, 2020.

## 11.9 Exercises

**11.1** ⌕ Concept check · 2 min

On run E1 the fitted convective form has $U = 8.46\,\mathrm{W/(m^2\,K)}$ and $A = 0.0012\,\mathrm{m^2}$. Doubling $A$ while halving $U$ leaves $J$ at 141.9 °C in every digit, while doubling $\alpha_1$ alone raises it to 5039.9 °C. What does this show?

(a) The objective is insensitive to geometry, because the heater is small beside the board

(b) The record fixes the products $UA$ and $mc_p$, not their factors, and $\alpha_1$ lies outside both

(c) The search stopped in a local minimum, so the area should be added to the fitted set

(d) The area enters only through radiation, which the convective form leaves out

**11.2** Concept check · 2 min

Fitting the convective plus radiative form to E1 lowers $J$ from 141.9 to 134.2 °C and reports that 70.1 % of the loss went out by radiation. What does the lower objective establish on its own?

(a) That radiation carries most of the heat that leaves the transistor in these data

(b) That the convective form is mis-specified and should be dropped for this board

(c) Nothing about the mechanism, since a larger nested model cannot fit its own record worse

(d) That the improvement is noise, since it is smaller than the spread of the noisy repeats

**11.3** Concept check · 2 min

With the noise off, the online estimate of the first-order form on E1 ends at $\tau = 111.9\,\mathrm{s}$ while the search returns 197.1 s. Which explanation does the chapter's evidence support?

(a) A window of 30 samples is too short to see a time constant near 200 s

(b) The estimator sums squared errors, and squared errors favor a short time constant

(c) The lower bound of $\tau$ is active through the run and holds the estimate down

(d) The random walk lets $\tau$ and $K_p$ drift to absorb the thermistor lag the form omits

**11.4** Theory and calculation · 7 min

Take the data generator at its defaults: $P_1 = 200$, $C_H = 5720\,\%\,\mathrm{s/°C}$, $\tau_{HA} = 20\,\mathrm{s}$, $\tau_{HH} = 100\,\mathrm{s}$ and $\tau_S = 140\,\mathrm{s}$, with $Q_2 = 0$. Write $a = 1/\tau_{HA}$ and $b = 1/\tau_{HH}$.

(a) Set the derivatives of Equations (11.1) to (11.3) to zero and find the steady-state gain from $Q_1$ to $T_1$. *Checkpoint:* 0.599 °C/%.

(b) The transfer function from $Q_1$ to $T_{H1}$ is $G(s) = (P_1/C_H)(s+a+b)/\big((s+a)(s+a+2b)\big)$. Its mean lag $-G'(0)/G(0)$ is the sum of the time constants of its poles less that of its zero. Find it, and add the thermistor lag. *Checkpoint:* 17.6 s and 157.6 s.

(c) Run E1 holds $Q_1 = 50\,\%$ from 2 s to 302 s and ends 25.95 °C above its start. What fraction of the final rise is that? *Checkpoint:* 86.6 %.

Say what (c) means for a model fitted to E1 alone.

**11.5** Theory and calculation · 6 min

The convective fit to E1 is $U = 8.46\,\mathrm{W/(m^2\,K)}$ and $\alpha_1 = 0.006\,86\,\mathrm{W/\%}$, with the held values $A = 0.0012\,\mathrm{m^2}$, $m = 0.004\,\mathrm{kg}$ and $c_p = 500\,\mathrm{J/(kg\,K)}$.

(a) Find the time constant and gain of the equivalent first-order form, and the gain error against the generator's 0.599 °C/%. *Checkpoint:* 197 s, 0.676 °C/%, about 13 % high.

(b) At steady state Equation (11.6) fixes only $UA\,T_a + \alpha_1 Q_1$ on a record with one heater level. With $UA$ held, how much must $\alpha_1$ change at $Q_1 = 50\,\%$ to absorb an error of 6 °C in $T_a$, and what fraction of the fitted value is that? *Checkpoint:* About 0.0012 W/%, about 18 %.

Compare (b) with the E1 columns of Table 11.7, and say why the doublet breaks the trade.

**11.6** Theory and calculation · 7 min

The radiative term of Equation (11.7) uses kelvin.

(a) Show that $\varepsilon\sigma A(T_1^4 - T_a^4) = h_r A(T_1 - T_a)$ with $h_r = \varepsilon\sigma(T_1 + T_a)(T_1^2 + T_a^2)$, and evaluate $h_r$ for $\varepsilon = 0.9$ and $T_a = 21$ °C at $T_1 = 21$ and at 47 °C. *Checkpoint:* 5.20 and 5.93 W/(m² K).

(b) The radiative fit to E1 has $U = 2.46\,\mathrm{W/(m^2\,K)}$. Add $h_r$ at both ends of the rise and compare with the convective fit's 8.46. *Checkpoint:* 7.66 and 8.39 W/(m² K).

(c) At the published values $U = 5\,\mathrm{W/(m^2\,K)}$ and $\varepsilon = 0.9$, find the instantaneous radiative share $h_r/(U + h_r)$ at 53.7 and at 83.1 °C. *Checkpoint:* 55.1 % and 58.6 %.

Explain from (a) to (c) why the emissivity valley of Table 11.6 is flat.

**11.7** App challenge · 8 min

Reload the app so every form holds its published values. In Settings set the Sensor noise to 0 and press Apply. With the Simulator source, choose the step test Q1 = 50% · 5 min, set Sim speed to 10× and press Run. The run ends at 5:02, about 30 s of real time, and sets $T_a$ to 21.0 °C in every form. This is scenario 1, Record and fit the first-order form, and scenario 2, Prove the lumped groups.

(a) Select First-order and read the Objective. Then press Best fit and read $\tau$, $K_p$ and the Objective. *Checkpoint:* 2767.9 °C before the fit. Then $\tau = 197$ s, $K_p = 0.676$ °C/% and 141.9 °C, each within one unit of the last digit.

(b) Select Convective and read the Objective, Implied $\tau$ and Implied gain at the published values. Then type 20 for $U$ and 0.0006 for $A$ and read the three again. *Checkpoint:* 1876.4 °C, 167 s and 0.833 °C/%, both times.

(c) Leave $A$ at 0.0006, press Best fit and read $U$, $\alpha_1$, Implied $\tau$, Implied gain and the Objective. *Checkpoint:* $U = 16.91\,\mathrm{W/(m^2\,K)}$ and $\alpha_1 = 0.0069\,\mathrm{W/\%}$, then 197 s, 0.676 °C/% and 141.9 °C.

Say what (b) and (c) show together.

**11.8** App challenge · 8 min

Continue from the previous exercise, or repeat its setup and Best fit on the Convective form. This is scenario 10, Use the 100% step.

(a) Select Convective + radiative and press Best fit on the E1 record. Read $U$, the Objective and the Loss split. *Checkpoint:* $U = 2.46\,\mathrm{W/(m^2\,K)}$, 134.2 °C, 29.9% convective and 70.1% radiative.

(b) Choose Q1 = 100% · 5 min and press Run without touching the parameters. When the run ends, read the Objective of the radiative form, then select Convective and read its Objective. Press Best fit there and see whether it improves. *Checkpoint:* 429.0 °C and 283.8 °C, and Best fit stays at 283.8 °C.

(c) Select Convective + radiative and press Best fit on E2. Read $U$, the Objective and the Loss split. *Checkpoint:* $U = 1.56\,\mathrm{W/(m^2\,K)}$, 250.2 °C, 19.7% and 80.3%.

Which model would you carry to a third experiment, and why?

**11.9** App challenge · 9 min

This is scenario 13, Watch the fit converge online. Reload the app, set the Sensor noise to 0 in Settings and press Apply, choose Q1 = 50% · 5 min and the First-order form, switch the chart view to Online estimate, set Sim speed to 10× and press Run.

(a) When the run ends, read the window estimates of $\tau$ and $K_p$, and check whether either is marked on its bound. Then press Best fit and read its $\tau$ in the same view. Express the gap in percent of the $\tau$ range of 10 to 600 s. *Checkpoint:* $\tau = 111.9$ s and $K_p$ about 0.582, neither on its bound. Best fit 197.1 s. Gap 14.4 %.

(b) In Settings set the Sensor response time to 20 s, press Apply and Run the same test again. Read the window $\tau$ and, after Best fit, the search's $\tau$, and express the gap the same way. *Checkpoint:* About 31.9 s against 35.3 s, a gap of 0.57 % of the range.

What does the change from (a) to (b) say about where the disagreement comes from?

**11.10** Application area · 3 min

An engineer plans to repeat the full-power step E2 on a physical board with the overtemperature limit at its default.

(a) The noise-free simulated E2 passes 70 °C about 259 s after the run starts and ends at 72.90 °C. If the board followed the simulator, what would happen? *Checkpoint:* The run stops near 259 s with both heaters off and the alarm latched.

(b) The limit can be raised to 90 °C before the board is connected. Would a longer, hotter step separate radiation from convection? *Checkpoint:* Hardly: the radiative share moves only from 55.1 to 58.6 %.

**11.11** Application area · 3 min

A model for tuning is fitted to a five-minute step that reached about 87 % of its final rise, and the fitted gain is 0.676 °C/% where the true gain is 0.599 °C/%.

(a) A PI controller is tuned for a chosen loop gain $K_c K_p$ using the fitted gain. By what factor is the actual loop gain off? *Checkpoint:* $0.599/0.676 = 0.886$, about 11 % low.

(b) If the search had also been allowed to move $T_a$ on the two-level record E3, where would it have gone, and what practice prevents it? *Checkpoint:* Down to about 15 °C, where $J$ is 354.6 °C. Measure $T_a$ instead.

**11.12** Review the agent · 5 min

An AI assistant reviews runs E1 and E2 and writes: “The convective plus radiative form is the better model of this board. It lowers the objective on both step tests, by 5.4 % on E1 and 11.8 % on E2, its loss breakdown shows that 70 % of the heat leaves by radiation, and its full-power steady state of 80.48 °C is within 0.46 °C of the true 80.94 °C, where the convective model is 7.62 °C off. Use it as the model for a predictive controller.” Find the flaw and say how you would check it.