PART 1 / LINEAR REGRESSION 20 minutes · Week of 23 November 2026

Predict electricity use with one line

At 07:00 the building manager knows the morning forecast in degrees Celsius. The daily meter total will be available at midnight. Use 25 independent simulated days to fit a straight line for forecasts between 17 and 35°C.

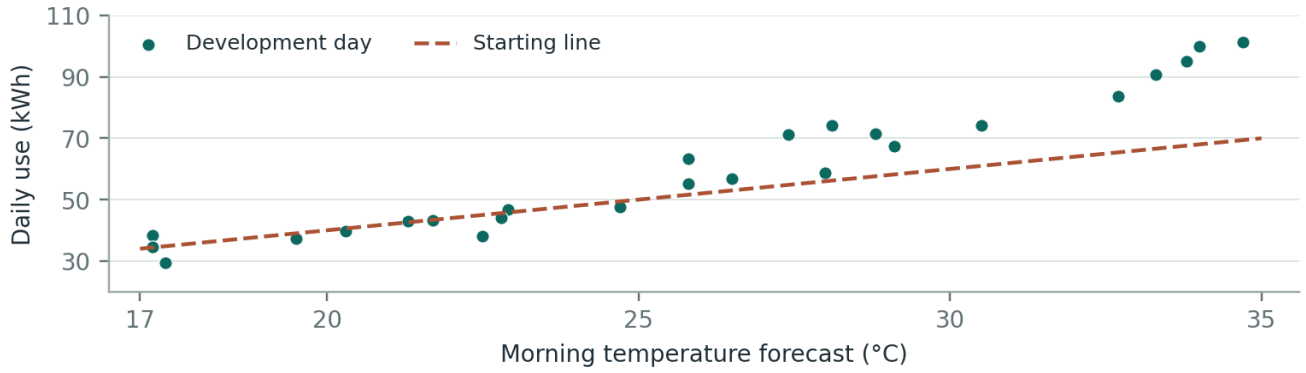

Each dot is one pair of measurements. The dashed line is the starting model, before gradient descent.

$$x = (T - 25)/5, \qquad \hat{y} = b + wx$$

T is the forecast in °C and ŷ is predicted electricity use in kWh. When T = 25°C, x = 0 and the predicted use is b. Raising T by 5°C raises x by one, so w is the change in predicted use per 5°C.

01 Start with b = 50 and w = 10. At 25°C the forecast is 50 kWh. Calculate its value at 30°C, then at 25.8°C.

02 Open **CS203_Lab11_Starter.ipynb** in Jupyter or Colab. D01 is (25.8°C, 63.4 kWh) and D02 is (25.8°C, 55.3 kWh). Explain why the line gives them the same prediction despite the different meter readings.

03 Complete **predict** and **mse**. Mean squared error (MSE) is the average of the squared differences between a prediction and its observed value. Use it to measure how well the line fits its training pairs.

Two-hour session

Linear regression 20 min · Gradient descent 35 min · Break 10 min · Five-fold cross-validation 40 min · Final forecast 15 min. The notebook contains every measurement and runs on a CPU with NumPy and Matplotlib.

My predictions at 30°C and 25.8°C

PART 2 / GRADIENT DESCENT 35 minutes

Update both parameters by hand

Use two practice pairs before training on the building data. Start at b = 50, w = 10 and learning rate α = 0.1. Write the prediction error as predicted minus observed.

T (°C)	x	Observed y	Predicted ŷ	Error e	Product ex
20	−1	38	_____	_____	_____
30	+1	66	_____	_____	_____

$$e_i = \hat{y}_i - y_i, \qquad L = \frac{1}{n}\sum_i e_i^2$$

$$g_b = \frac{2}{n}\sum_i e_i, \qquad g_w = \frac{2}{n}\sum_i e_i x_i$$

$$b_{\text{new}} = b - \alpha g_b, \qquad w_{\text{new}} = w - \alpha g_w$$

01 Calculate the starting MSE and both gradients using the same two old predictions. Update b and w together. Use the new line to predict both days again. Did the average squared error fall?

02 Implement **gradient_step** and complete the repeated update in **fit_line**. Run exactly two updates on the practice pairs. The loss history should contain three entries, including the starting loss.

03 Compare two learning rates on the same 20 training days using the notebook's plot. Start each line from b = 50, w = 10. Read both training losses at step 10 and explain which rate moved faster.

What one pass through the loop does

Use the current b and w for all predictions, calculate both derivatives and take one step. The initial values are set once when a fresh fit starts.

My two gradients, updated line and new MSE

PART 3 / FIVE-FOLD CROSS-VALIDATION 40 minutes

Take turns holding out five days

The 25 development days have been shuffled into five fixed groups. Each row below defines one round. Start a fresh line in every round at b = 50 and w = 10. Use α = 0.1 for 200 updates.

Hold out	Five validation day IDs	Train on groups	MAE (kWh)
A	D25 D13 D07 D06 D14	B, C, D, E	_____
B	D18 D08 D11 D22 D20	A, C, D, E	_____
C	D02 D21 D05 D04 D23	A, B, D, E	_____
D	D01 D09 D24 D10 D15	A, B, C, E	_____
E	D03 D12 D17 D19 D16	A, B, C, D	_____

$$\mathrm{MAE}_{\text{fold}} = \frac{1}{5} \sum_{i \text{ in fold}} |y_i - \hat{y}_i|$$

01 Complete **cross_validate**. Select the 20 training rows outside the held-out group, fit the line to those 20 and predict the five temperatures in the group. Save each prediction under its original day ID.

02 For round A, compare the five printed predictions with their meter readings. Check the five absolute errors by hand and average them. Enter all five fold scores in the table above.

03 Average the five MAEs. Check the same number by taking the mean absolute error of all 25 held-out predictions. Explain why these averages agree when all five groups are the same size.

04 The notebook adds 100 kWh to group A's meter readings as a check. Explain why the model trained on B–E gives exactly the same A predictions after this change. Which part of the round does change?

Average five-fold MAE, highest-error fold and reason for fitting five lines

PART 4 / APPLYING THE FITTED LINE 15 minutes

Forecast before reading the meter

After estimating the error by five-fold cross-validation, fit one fresh line on all 25 development days. Use the same starting values, learning rate and 200 updates. This final line supplies every new forecast.

Make the predictions first

01 Record the fitted intercept and slope. The slope is expressed per 5°C, so divide it by five to get the change per 1°C. Use the line to forecast the daily electricity use at 30°C.

02 Run the cell listing predictions for 12 fresh days and enter the first three below. Next, reveal their observed meter totals in the following cell. Calculate each absolute error.

Fresh day	Forecast °C	Predicted kWh	Observed kWh	Absolute error (kWh)
T01	32.8	_____	_____	_____
T02	26.7	_____	_____	_____
T03	22.5	_____	_____	_____

03 Give the manager the forecast at 30°C, mean five-fold MAE and MAE on the 12 fresh days. Add one sentence explaining why the last two numbers can differ.

What to keep

Save your notebook with the line equation, worked gradient update, loss curves, five fold scores and fresh-day predictions. Each error you give the manager should be in kWh.

My note to the building manager