Least Squares Regression and R^2

Mr. Merrick · September 27, 2025

1) Total Variance in y: squares to the mean (the "null model")

Point	A	B	C	D	E	F	G
x	2.0	3.5	5.0	6.0	7.5	9.0	11.0
y	1.10	1.925	2.25	2.90	3.175	3.90	4.35

The dashed horizontal line marks $\bar{y} = 2.800$. Each dotted arrow is a vertical deviation $(y_i - \bar{y})$. For every point, draw a **square** using that arrow as a side. The total area of all squares is

$$\text{SST} = \sum_{i=1}^{n}(y_i - \bar{y})^2 \quad \text{(total variation in } y\text{)}.$$

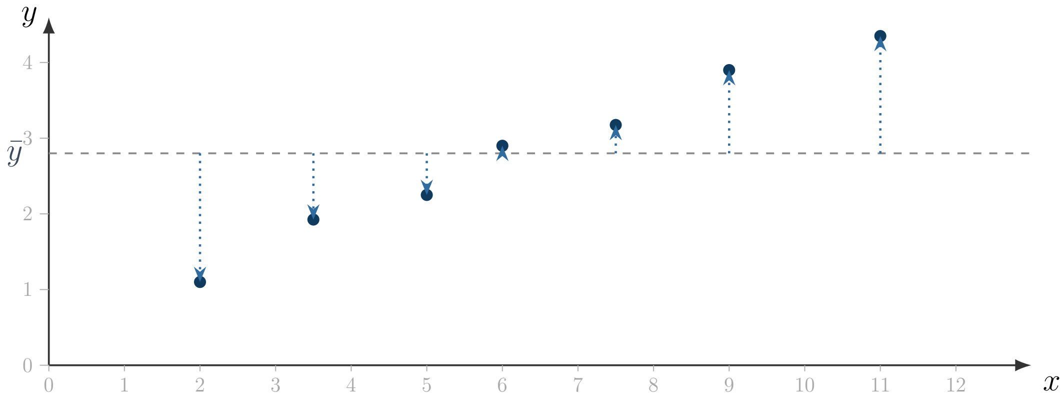

Record your total: $\text{SST} = \sum(y_i - \bar{y})^2 =$

Quick questions

1. The *null model* predicts every value of y with $\bar{y}$. Does it take x into consideration, or use x to explain variation?

2. If we changed the units of y (e.g. cm → m), how would the area of each square change?

3. For this dataset, does it look like there is a relationship between y and x?

2) Least Squares: choose a model to minimize squared residuals

A linear model predicts $\widehat{y} = a + bx$. Each residual is $e_i = y_i - \widehat{y}_i$ (Actual − Predicted — remember AP). We choose $(\hat{a}, \hat{b})$ that *minimizes* the total *sum of squared errors*. Draw squares for each model's residuals.

$$\mathrm{SSE}(a,b) = \sum_{i=1}^{n} \Big(y_i - \hat{y}_i\Big)^2 = \sum_{i=1}^{n} \Big(y_i - (a + bx_i)\Big)^2.$$

Bad model: $\widehat{y} = 4.2 - 0.35x$

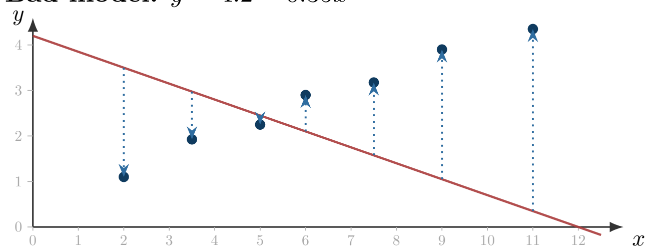

	x_i	y_i	$\widehat{y}_i$	$e_i = y_i - \widehat{y}_i$	e_i^2
A	2.0	1.10			
B	3.5	1.925			
C	5.0	2.25			
D	6.0	2.90			
E	7.5	3.175			
F	9.0	3.90			
G	11.0	4.35			

$\mathrm{SSE}_{\mathrm{bad}} = \sum e_i^2 =$

Best (least-squares) model: $\widehat{y} = 0.5440 + 0.3589x$

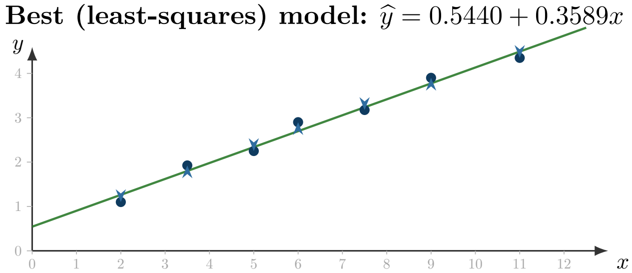

	x_i	y_i	$\widehat{y}_i$	$e_i = y_i - \widehat{y}_i$	e_i^2
A	2.0	1.10			
B	3.5	1.925			
C	5.0	2.25			
D	6.0	2.90			
E	7.5	3.175			
F	9.0	3.90			
G	11.0	4.35			

$\mathrm{SSE}_{\mathrm{best}} = \sum e_i^2 =$

Null model (ignore x): $\widehat{y} = \bar{y}$

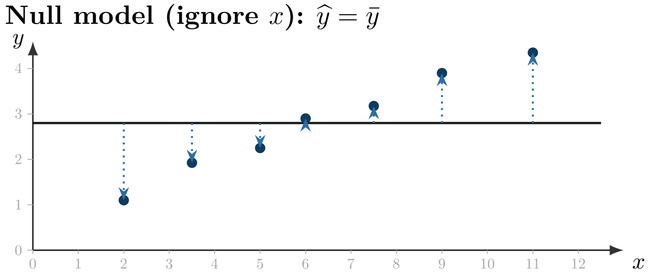

	x_i	y_i	$\widehat{y}_i$	$e_i = y_i - \widehat{y}_i$	e_i^2
A	2.0	1.10			
B	3.5	1.925			
C	5.0	2.25			
D	6.0	2.90			
E	7.5	3.175			
F	9.0	3.90			
G	11.0	4.35			

$\mathrm{SSE}_{\mathrm{null}} = \sum e_i^2 =$

Quick questions

1. Which model has the *smallest* total square area?

2. The bottom row ("ignore x") gives a baseline amount of square area. How can we tell if another model is an *improvement* compared to this baseline?

3. If a model's square area is only a little smaller than the baseline, what does that suggest about x? What if the model's square area is much smaller?

3) Decomposing squares and R^2

Any response y_i can be decomposed into

$$y_i - \bar{y} \;=\; (y_i - \widehat{y}_i) \;+\; (\widehat{y}_i - \bar{y}).$$

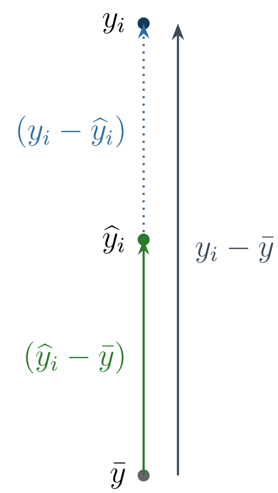

Squaring and summing over points leads to the **sum-of-squares identity**

$$\underbrace{\text{SST}}_{\sum(y_i-\bar{y})^2} \;=\; \underbrace{\text{SSR}}_{\sum(\widehat{y}_i-\bar{y})^2} \;+\; \underbrace{\text{SSE}}_{\sum(y_i-\widehat{y}_i)^2}\;.$$

The coefficient of determination is

$$R^2 \;=\; \frac{\text{SSR}}{\text{SST}} \;=\; 1 - \frac{\text{SSE}}{\text{SST}},$$

the *proportion of total square area explained* by using x.

Shade/identify the squares:

- On the *null model* panel, your squares show SST.
- On the *best model* panel, your squares show SSE.
- The *explained* squares correspond to SSR = SST − SSE.

	SST	SSE (best)	SSR = SST − SSE	$R^2 = \dfrac{\text{SSR}}{\text{SST}}$
Values				

Practice

1. Explain why the explained squares (SSR) must be *nonnegative*.

2. If a different line (not least squares) is used, which quantity necessarily increases, SSE or SST? Why?

3. In this dataset, R^2 is very close to 1. What does that tell you about the usefulness of x for predicting y?

4. What is the lowest possible value of R^2 and what does it mean in context? What is the largest value of R^2 and what does it mean in context?

5. $\star$ Prove SST = SSR + SSE.