

Goals

This week:

- model a linear relationship by fitting a least-squares line to the data
- use a residual plot to assess the appropriateness of fitting a linear model to the data
- interpret the intercept and slope of the fitted line
- use the equation of a fitted line to make predictions
- distinguish between interpolation and extrapolation when using the fitted line to make predictions, recognising the potential dangers of extrapolation
- write up the results of the above analysis in a systematic and concise matter

Theoretical Components

Resources:

For this week the theory work is in the *PDF file*:
Week 4 Notes & Exercises

Line of best fit

<https://www.youtube.com/watch?v=DmGLQkUm-4g>

Interpolation and extrapolation

<https://www.youtube.com/watch?v=bEANDIjkqcU>

Knowledge Checklist

- Fitting straight lines to bivariate data
- Balancing points
- Equation of a straight line
- Least squares regression
- Dependent and independent variable
- Interpolation
- Extrapolation

Practical Components

There are questions to be answered in the booklet *Week 4 Notes & Exercises*

Investigation

See the end of the brief 😊

Quiz

No mathspace.co for this week.

However, please remember to complete your reflection for this week 😊

MATHEMATICAL APPLICATIONS 3

WEEK 4 NOTES & EXERCISES


Fitting straight lines to bivariate data

Fitting a line of 'best fit' to bivariate data (scatter plots) enables us to analyse data and possibly make predictions. Regression analysis is concerned with finding these straight lines.

In our previous work, when we displayed bivariate data as a scatterplot, the independent variable was placed on the horizontal axis and the dependent variable was placed on the vertical axis. When the relationship between two variables (x and y) is described in equation form, such as $y = mx + c$, the subject, y , is the dependent variable and x is the independent variable.

Fitting a straight line by eye


Consider the set of bivariate data points below. In this case, the x -values could be heights of married women, why y -values could be the heights of their partners. We wish to determine a linear relationship between these two random variables.



Of course, there is no single straight line which would go through all the points, so we can only estimate such a line. Furthermore, the more closely the points appear to be on or near a straight line, the more confident we are that such a linear relationship may exist and the more accurate our fitted line should be.

Consider the estimate, drawn 'by eye' in the figure below. It is clear most of the points are on or very close to this straight line. This line was easily drawn since the points are very much part of an apparent linear relationship. However, note that some points are below the line and some are above it. Furthermore, if x is the height of married women and y is the height of their partners, it seems that their partners are generally taller than them.

Regression analysis is concerned with finding these straight lines using various methods so the number of points above and below the lines are 'balanced'.




Method of fitting lines by eye

There should be an equal number of points above and below the line. For example, if there are 12 points in the data set, 6 should be above the line and 6 below it. This may appear logical or even obvious but fitting by eye involves a considerable margin of error.

Example


Fit a straight line to the data in the figure using the equal-number-of-points method.



Solution


The number of points (n) is 8. We need to fit a line where 4 points are above and below the line. Using a clear plastic ruler, try to fit the best line.

Attempt 1:




The first attempt has 3 points below the line where there should be 4. We need to make refinements.

Attempt 2:






The second attempt is an improvement, but the line is too close to the points above the line. We can improve on the position of the line until there is a better 'balance' between the upper and lower points is achieved.
Final attempt:




Exercise 1

Fit a straight line to the data in the scatter plots using the equal-number-of-points method.



Fitting a straight line – least-squares regression

This is used when the data shows a linear relationship and there are no obvious outliers (or these are discarded). To understand the underlying theory behind least-squares, consider the regression line shown below.



We want to minimise the total of the vertical lines, or 'errors' in some way, such as balancing the errors above and below the line. This is reasonable, but for mathematical reasons it is preferable to minimise the sum of the squares of each of these errors. This is the essential mathematics of least-squares regression.

The formula for a straight line when using the x and y axis is:

$$y = mx + c$$

where:

- m = gradient of the line $\left(\frac{\text{rise}}{\text{run}}\right)$
- c = y intercept (i.e. the value of y when $x = 0$)

Fortunately, we can use an online calculator to find the equation of the line. This calculator is found at

<https://www.socscistatistics.com/tests/regression/default.aspx>

Example


Use the online calculator, find the line of best fit.

x	y
4	10
6	8
7	13
9	15
10	14
12	18
15	19
17	23

Solution

The calculator gives the equation $y = 1.04x + 4.57$ to be the line of best fit.

The calculator also shows the scatter plot and line of best fit. Ideally, the graph would start at $(0,0)$, then we could see the y -intercept is 4.57.



Exercise 2

Find the equation of the linear regression line for the following data set using the least squares method utilising the online calculator:

1.

x	4	6	7	9	10	12	15	17
y	10	8	13	15	14	18	19	23

2.

x	1	2	3	4	5	6	7	8	9
y	35	28	22	16	19	14	9	7	2

Interpretation, interpolation and extrapolation

Once you have a linear regression line, the slope and intercept can give important information about the data set. The slope (m) indicates the rate at which the data are increasing or decreasing. The y-intercept indicates the approximate value of the data when $x = 0$.

The line can be used to make predictions such as finding a value for y for a given value of x – where x is not in the given data.

The line is always of the form: $y = (\text{gradient}) \times x + (\text{y - intercept})$
 $y = mx + c$

This line can be used to 'predict' data values for a given value of x . Of course, these are only approximations since the regression line itself is only an estimate of the 'true' relationship between the bivariate data. However, they can still be used, in some cases, to provide additional information about the data set such as making predictions.

Interpolation

Interpolation is the use of the regression line to predict values in between two values already in the data set. If the data set is highly linear (r is near +1 or -1) then we can be confident that our interpolated value is quite accurate. If the data are not highly linear (r is near 0) then our confidence is duly reduced.

Extrapolation

Extrapolation is the use of the regression line to predict values smaller than the smallest value in the data set or larger than the largest value.

Two problems may arise in attempting to extrapolate from a data set. Firstly, it may not be reasonable to extrapolate too far away from the given data values. For example, suppose there is a weather data set for 5 days. Even if it is highly linear (r is near +1 or -1), a regression line used to predict the same data 15 days in the future is highly risky. Weather has a habit of randomly fluctuating and patterns rarely stay stable for very long.

Secondly, the data may be highly linear in a narrow band of the given data set. For example, there may be data on stopping distances for a train at speeds of between 30 and 60 km/h. Even if they are highly linear in this range, it is unlikely that things are similar at very low speeds (0-15 km/h) or high speeds (over 100 km/h).

Generally, we should feel more confident about the accuracy of a prediction derived from interpolation than on derived from extrapolation. Of course, it still depends upon the correlation coefficient (r). The closer to linearity the data are, the more confident our predictions in all cases.

Example

Using the data below, find the height of the child at age 8 and 15.

Age (years)	1	3	5	7	9	11
Height (cm)	60	76	115	126	145	148

Solution

The calculator gives the equation $y = 9.4x + 55.3$ to be the line of best fit. The gradient is 9.4 thus the average growth is 9.4 cm per year. $c = 55.3$ thus at birth the height of the baby is 55.3 cm.

At age 8, $x = 8$, interpolation:

$$y = mx + c$$

$$y = 9.4(8) + 55.3$$

$$y = 130.5 \text{ cm}$$

At age 15, $x = 15$, extrapolation:

$$y = mx + c$$

$$y = 9.4(15) + 55.3$$

$$y = 196.3 \text{ cm}$$

This prediction is clearly unreliable as it assumes a constant rate of growth throughout life. Thus, even though $r = 0.96$, in this case it is only useful in the range the data provides.

Exercise 3

1. A drug company wishes to test the effectiveness of a drug to increase red blood cell counts in people who have a low count. The following data was collected.

Day of experiment	4	5	6	7	8	9
Red blood cell count	210	240	230	260	260	290

- a. What are the independent and dependent variables in this case?
- b. Find the relationship in the form of $y = mx + c$ using the online calculator
- c. What is the rate at which the red blood cell count is changing?
- d. What was the blood cell count at the beginning of the experiment (i.e. on day 0)?

2. A wildlife exhibition is held over 6 weekends and features still and live displays. The number of live animals that are being exhibited varies each weekend. The number of animals participating, together with the number of visitors to the exhibition each weekend, is shown below.

Number of animals	6	4	8	5	7	6
Number of visitors	311	220	413	280	379	334

- a. Find the rate of increase of visitors as the number of live animals is increased.
- b. Find the predicted number of visitors if there are no live animals.
3. An electrical goods warehouse produces the following data showing the selling price of electrical goods to retailers and the volume of those sales.

Selling price (\$)	60	80	100	120	140	160	200	220	240	260
Sales volume (x1000)	400	300	275	250	210	190	150	100	50	0

Perform a least-squares regression analysis and discuss the meaning of the gradient and y-intercept.

4. The following table represents the costs for shipping a consignment of shoes from Melbourne factories. The cost is given in terms of distance from Melbourne. There are two factories that can be used. The data is summarised below.

Distance from Melbourne (km)	10	20	30	40	50	60	70	80
Factory 1 cost (\$)	70	70	90	100	110	120	150	180
Factory 2 cost (\$)	70	75	80	100	100	115	125	135

- a. Find the least-squares regression equation for each factory.

Factory 1

Factory 2

- b. Which factory is likely to have the lowest shipping cost to a shop in Melbourne? (use numbers)

- c. Which factory is likely to have the lowest shipping cost to Mytown, 115 kilometres from Melbourne? (use maths)

5. A factory produces calculators. The least-squares regression line for cost of production (C) as a function of numbers of calculators (n) produced is given by:

$$C = 7.76n + 660$$

Furthermore, this function is only deemed accurate when producing between 100 and 1000 calculators.

- Find the cost to produce 200 calculators
- How many calculators can be produced for \$2000?
- Find the cost to produce 10 000 calculators.
- What are the 'fixed' costs for this production?
- Which of part a. to c. is an interpolation?

Complete the table below.

The first two columns give the values for age (x , in years) and systolic blood pressure (y , in mmHg) for 15 women.

X	Y	X^2	Y^2	XY
42	130	1764	16900	5460
46	115			
42	148			
71	100			
80	156			
74	162			
70	151			
80	156			
85	162			
72	158			
64	155			
81	160			
41	125			
61	150			
75	165			
Total of X $\Sigma = 984$	Total of Y $\Sigma = 2193$	Total of X^2	Total of Y^2	Total of XY

Now use the values from the table and the formulae given to estimate the parameters of the linear regression for the x and y . i.e use the sums and the formulae given to work out 'a', 'b', r and r^2 . Show all working.

Check your answers on

<https://www.socscistatistics.com/tests/regression/default.aspx>

Write out the linear regression rule for the age (x , in years) and systolic blood pressure (y , in mmHg) for 15 women.

Use this rule to estimate the systolic blood pressure for a 59 year old woman.

Formulae:

$$a = \frac{n \sum XY - \sum X \sum Y}{n \sum X^2 - (\sum X)^2}$$

$$b = \bar{Y} - a\bar{X}$$

$$r = \frac{\sum XY - \frac{\sum X \sum Y}{n}}{\sqrt{\left[\sum X^2 - \frac{(\sum X)^2}{n} \right] \left[\sum Y^2 - \frac{(\sum Y)^2}{n} \right]}}$$

Marking Rubric

Week 4

Name:

CRITERIA	EXPECTATIONS	POSS	MULT	GIVEN	TOTAL
Practical	Student completes practical work, including exercises and Mathspace task, of the brief to an acceptable standard set by the teacher.	2	3		/6
Investigation Task	Student completes the investigation task of the week to an acceptable standard set by the teacher.	2	2		/4
Reasoning and Communications	Student responses are accurate and appropriate in presentation of mathematical ideas, with clear and logical working out shown.	4	-		/4
Concepts and Techniques	Student submitted work selects and applies appropriate mathematical techniques to solve practical problems and demonstrates proficiency in the use of mathematical facts, techniques, and formulae.	4	-		/4
	Submission Guidelines				
Timeliness	Student submits the exercises, Mathspace/online task, and investigation by the set deadline. See scoring guidelines for specific details.	2	-		/2
				FINAL	/20

Student Reflection: How did you go with this week's work? What did you learn? What did you find easy?
 What do you need to work on?