University of Exeter

Data Visualisation

Week 03-BEM2031

Term2: 2024/25

University of Exeter

Today:

- Assess the clarity of a visualisation (e.g. a graph) examples
- Understand various visualisation tools and what they are best suited for

- Most queries I have had are related to not setting the working directory. R needs to know where to look for a file, such as a csv, if asked to access one.

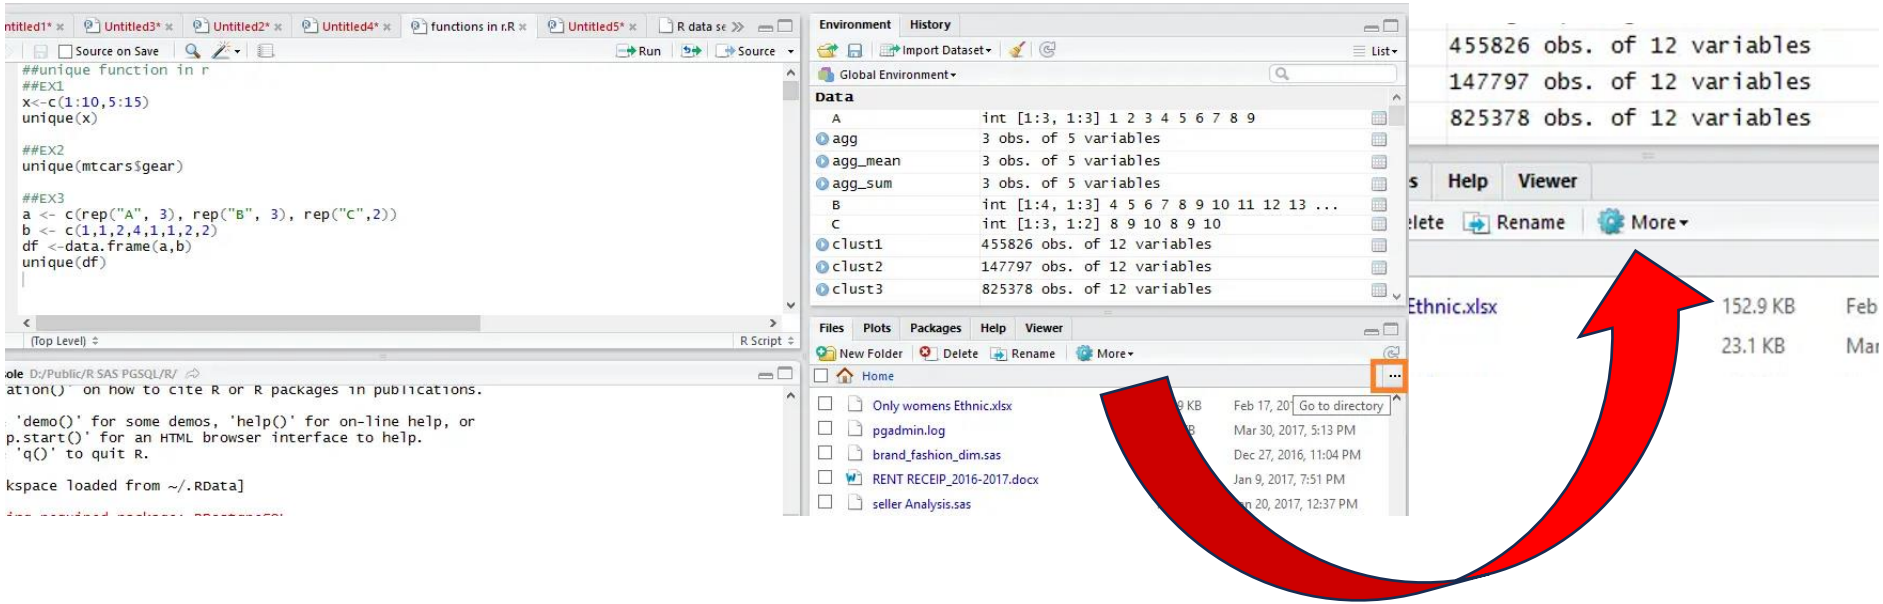

- Make sure you can knit to pdf and/or html from your R markdown files. If you are having problems knitting to pdf you may need to use **install.packages('tinytex')**

- I have had several requests to please QUIETEN DOWN in workshop labs – it is hard for everyone to concentrate with the amount of noise in the room!!

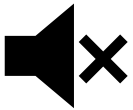

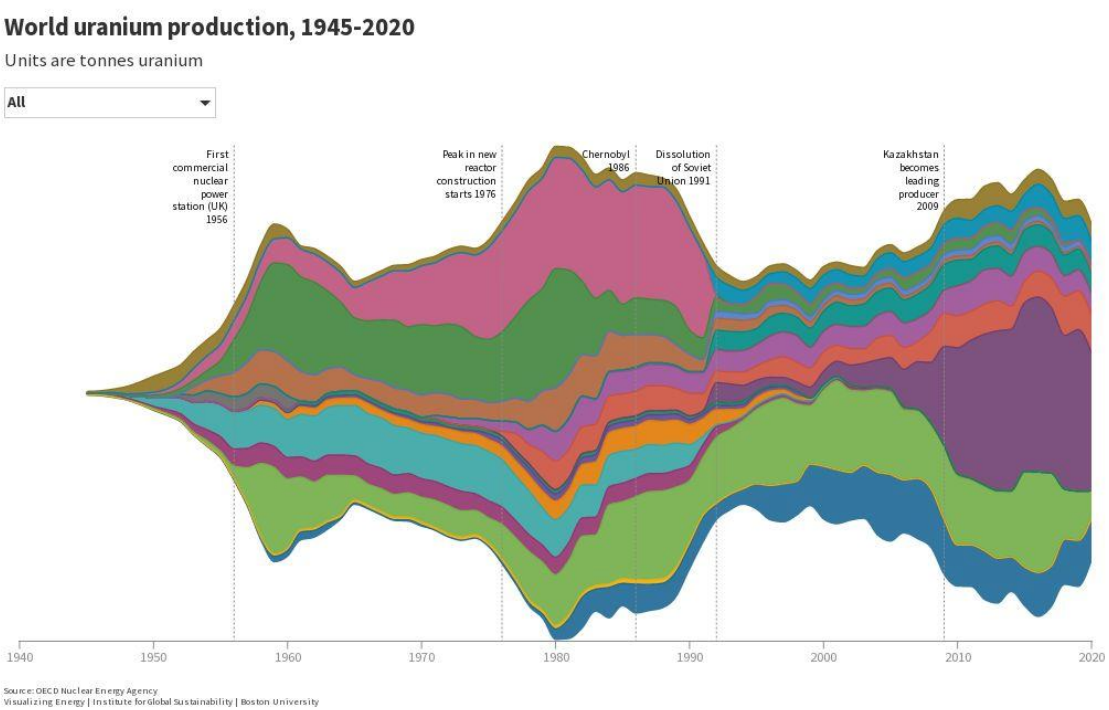

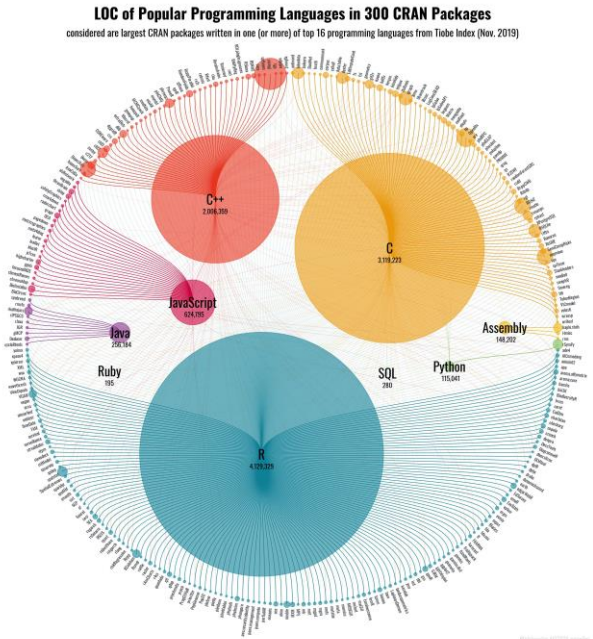

Why visualize?

1. Exploratory Data Analysis
2. Communicating results

University of Exeter

Histograms

Statistics

(summary statistics)

Descriptive Statistics

Inferential Statistics

Statistics

(summary statistics)

Descriptive Statistics

Inferential Statistics

Central tendency	• Mean • Median • Mode
Dispersion / variability	• Range • Interquartile range • Variance • Standard deviation
Skewness	• Symmetric • Left • Right

Statistics

(summary statistics)

Descriptive Statistics		Inferential Statistics
Central tendency	• Mean • Median • Mode	Estimation parameters
Dispersion / variability	• Range • Interquartile range • Variance • Standard deviation	Hypotheses testing
Skewness	• Symmetric • Left • Right	Tests of difference and similarity

Statistics

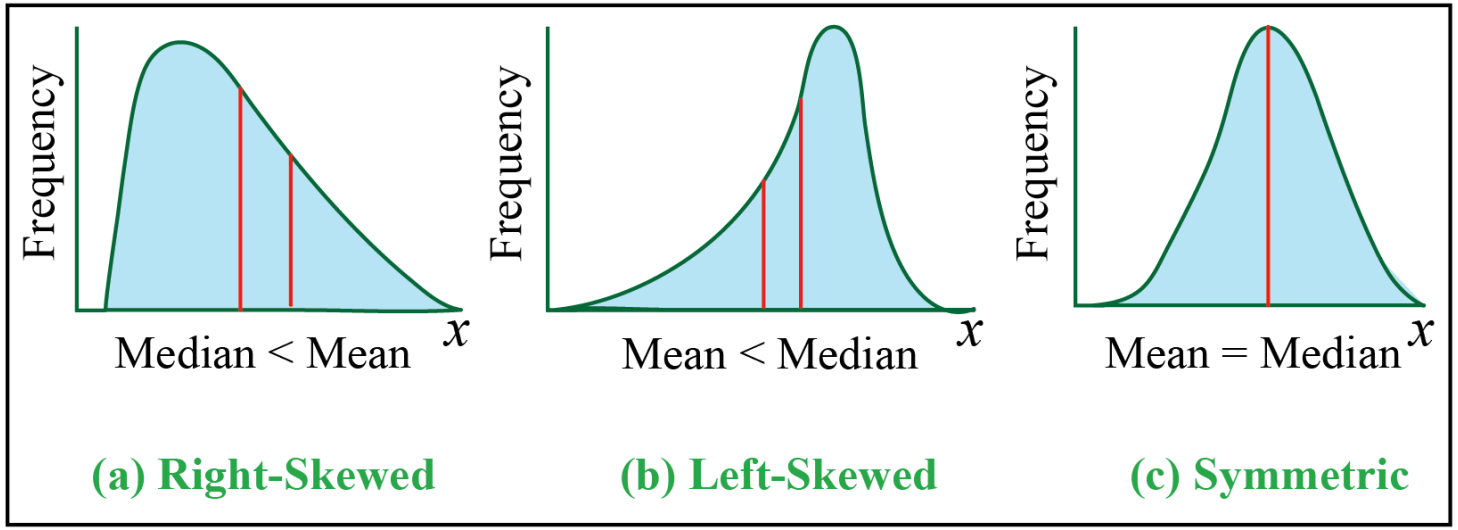
Frequency
Median < Mean
x
(a) Right-Skewed
Frequency
Mean < Median
x
(b) Left-Skewed
Frequency
Mean = Median
x
(c) Symmetric

Statistics

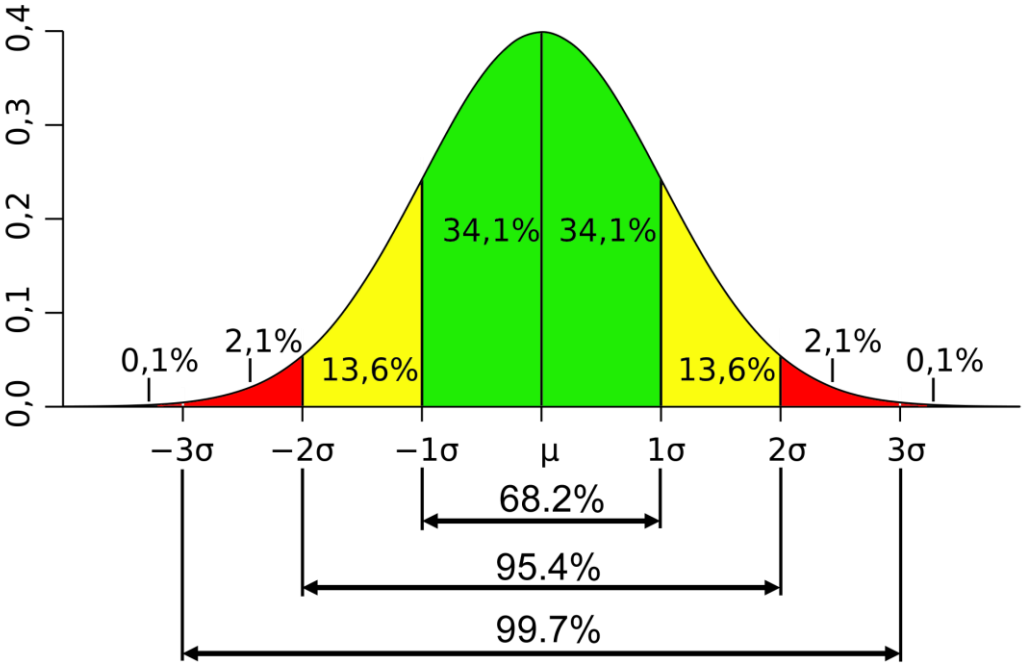

Variance: the average squared differences from the mean

Standard deviation: the square root of the variance

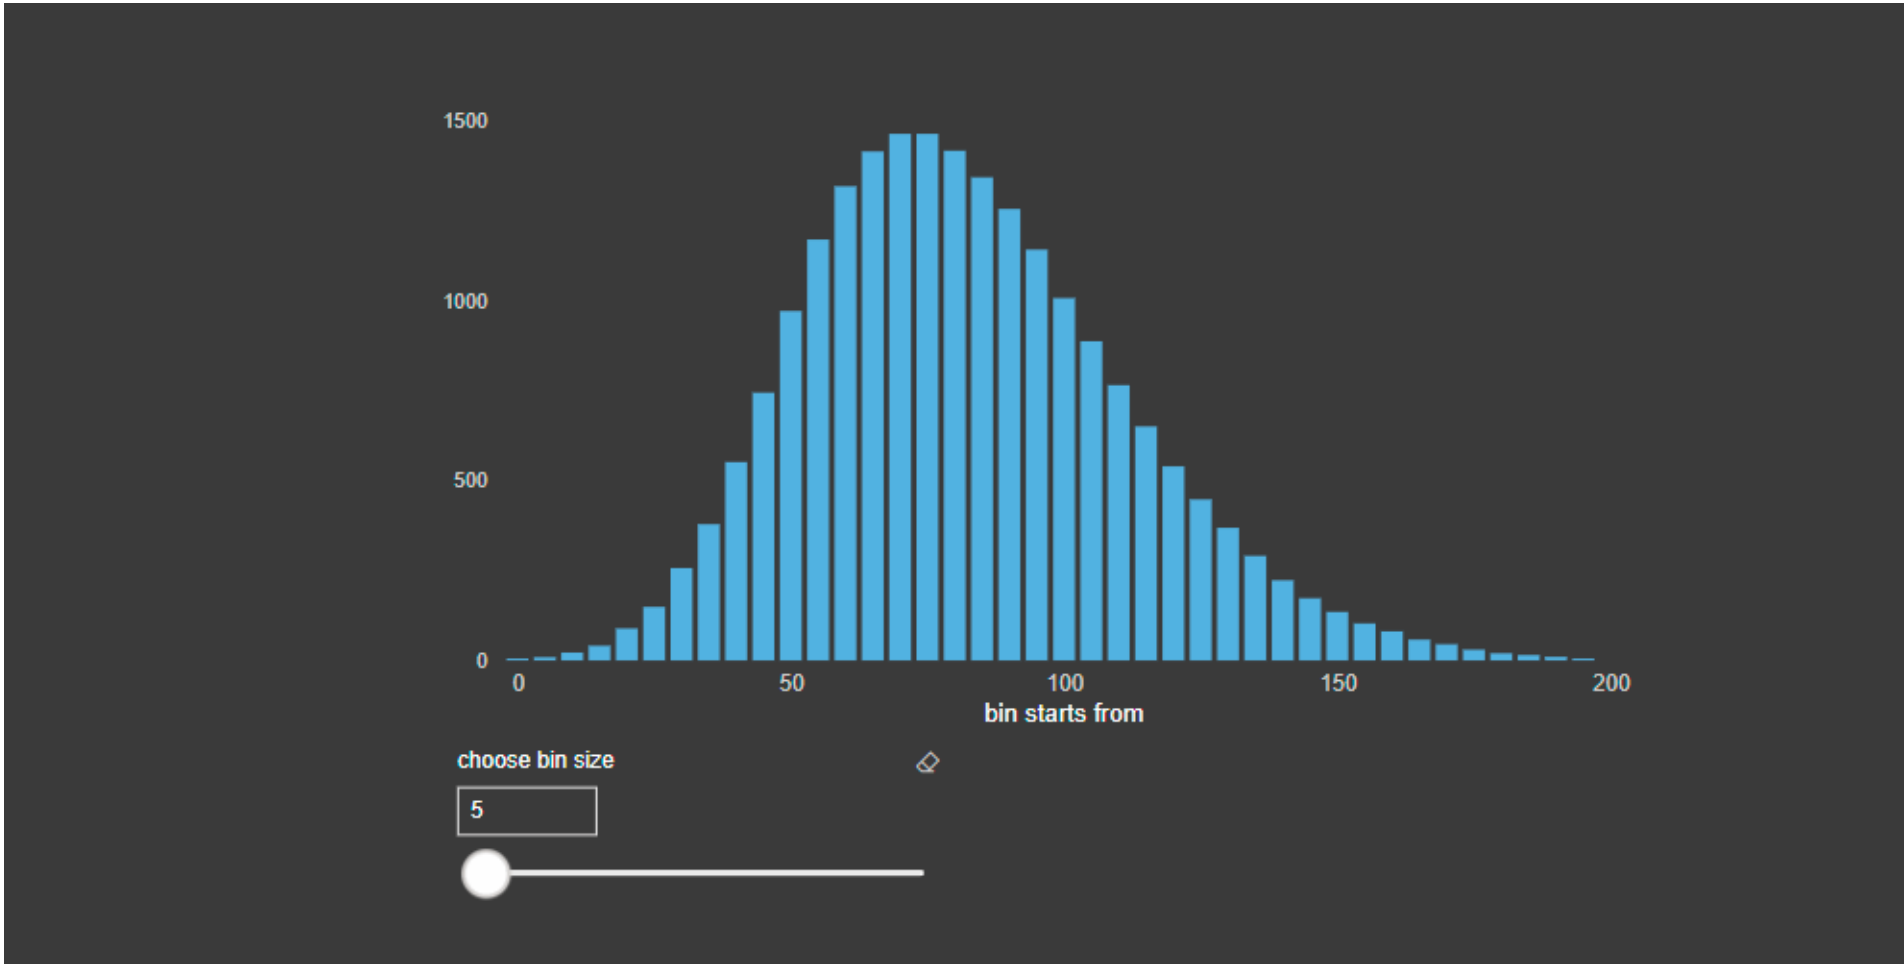
1500
1000
500
0
0
50
100
150
200
bin starts from
choose bin size
5

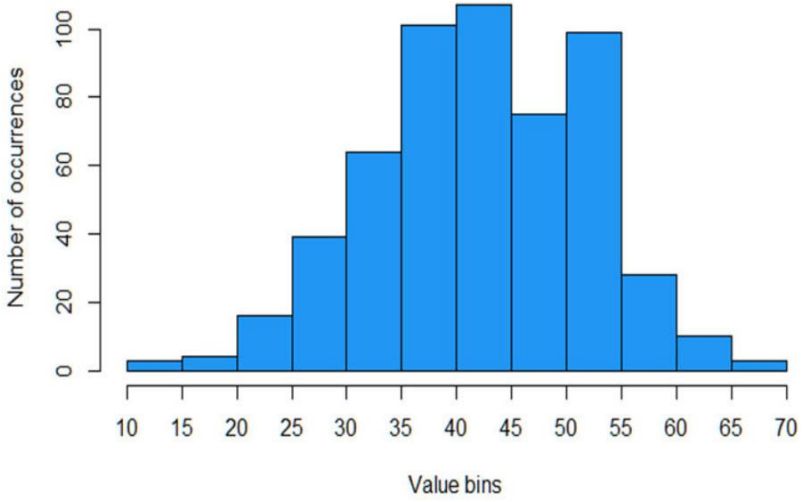

Symmetric (normal) vs skewed and uniform distriutions

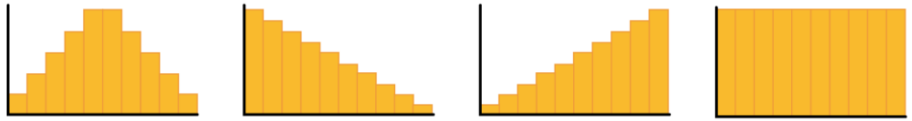

Normal distribution
(unimodal, symmetric, the "bell curve")

Right-skewed distribution
(Positively-skewed)

Left-skewed distribution
(Negatively-skewed)

Uniform distribution
(equal spread, no peaks)

Unimodal vs bimodal distributions

Normal distribution
(unimodal, symmetric, the "bell curve")

Symmetric bimodal distribution
(two modes)

Non-symmetric bimodal distribution
(two modes)

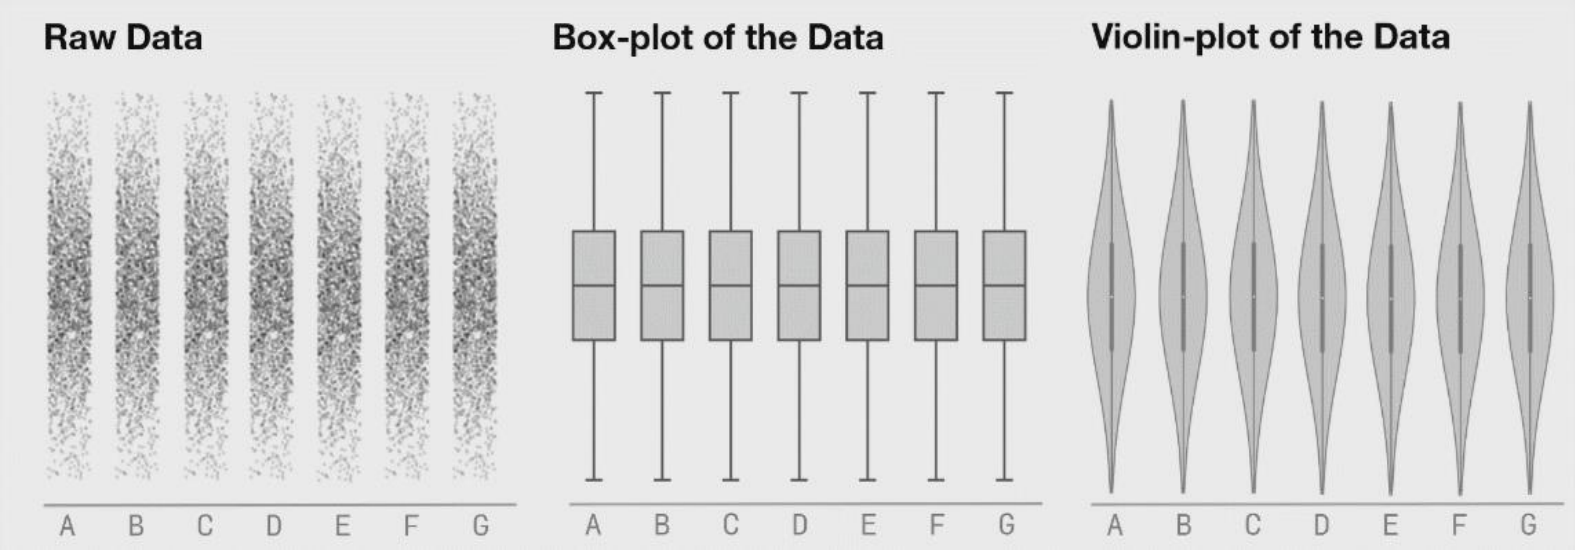
Raw Data
Box-plot of the Data
Violin-plot of the Data
A B C D E F G
A B C D E F G
A B C D E F G

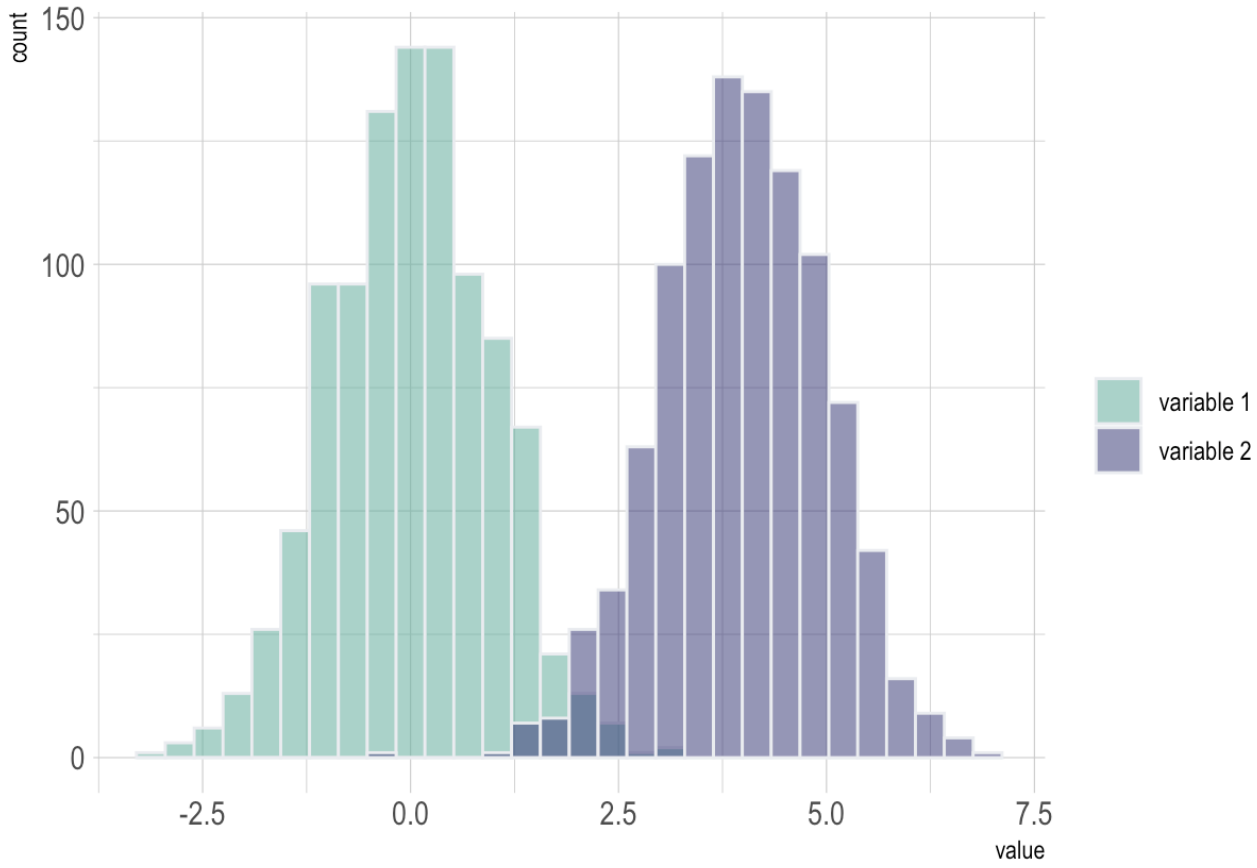
count
150
100
50
0
-2.5
0.0
2.5
5.0
7.5
value
variable 1
variable 2

University of Exeter

FINISH TIME DISTRIBUTIONS

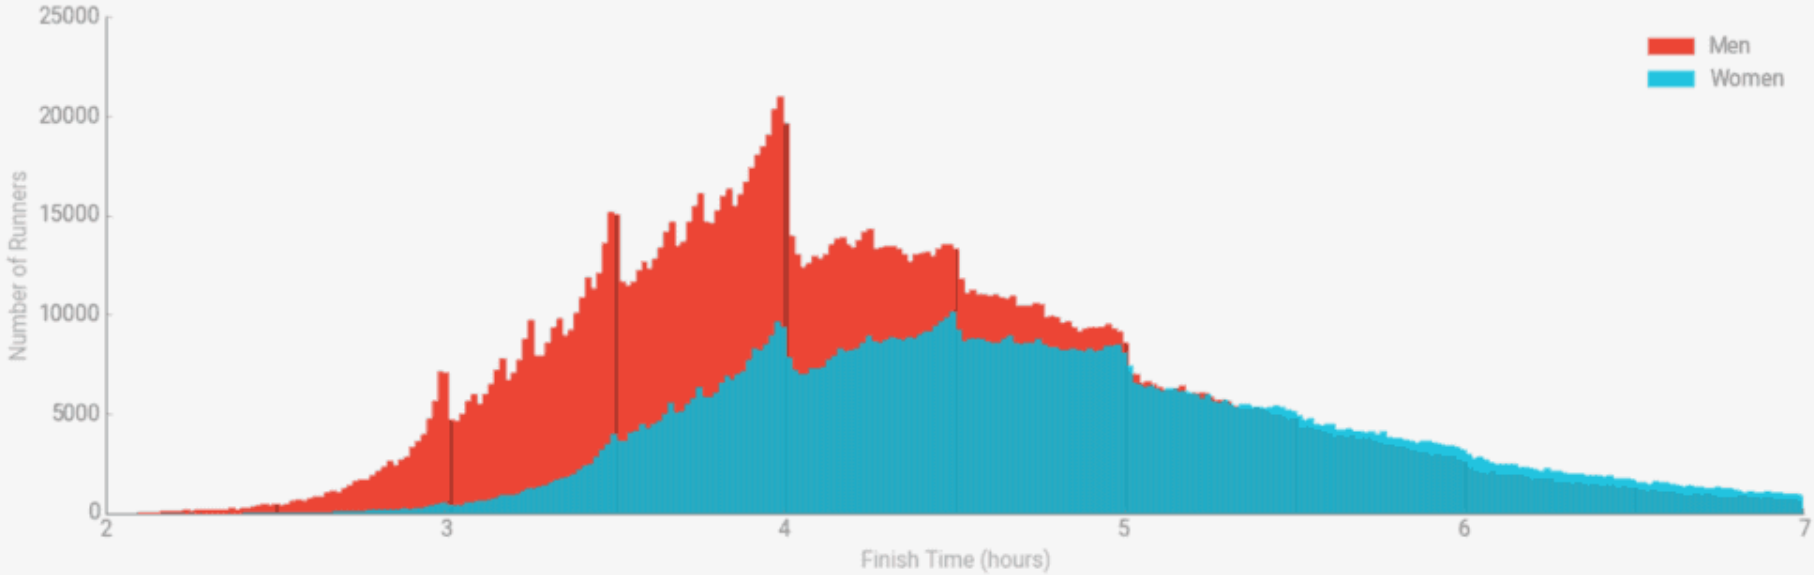

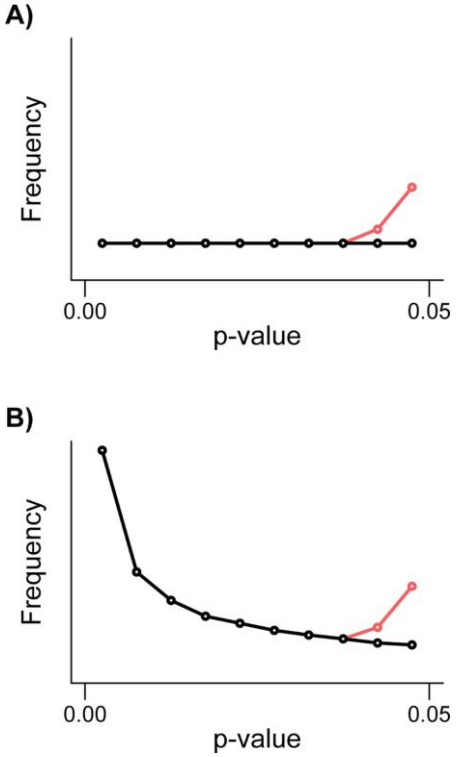

P-hacking is a bias in the scientific literature that occurs when researchers manipulate data or statistical analyses to obtain significant results.

P-hacking can involve looking at many relationships, collecting or selecting data, or choosing different methods until non-significant results become significant.

P-hacking can lead to false or exaggerated findings and undermine the validity of scientific research.

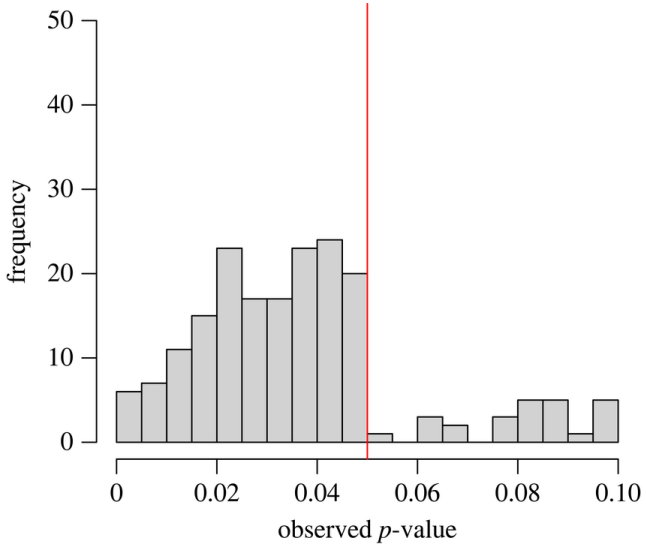

Mesquida, C., Murphy, J., Lakens, D., & Warne, J. (2022). Replication concerns in sports and exercise science: a narrative review of selected methodological issues in the field. *Royal Society Open Science*, *9*(12), 220946.

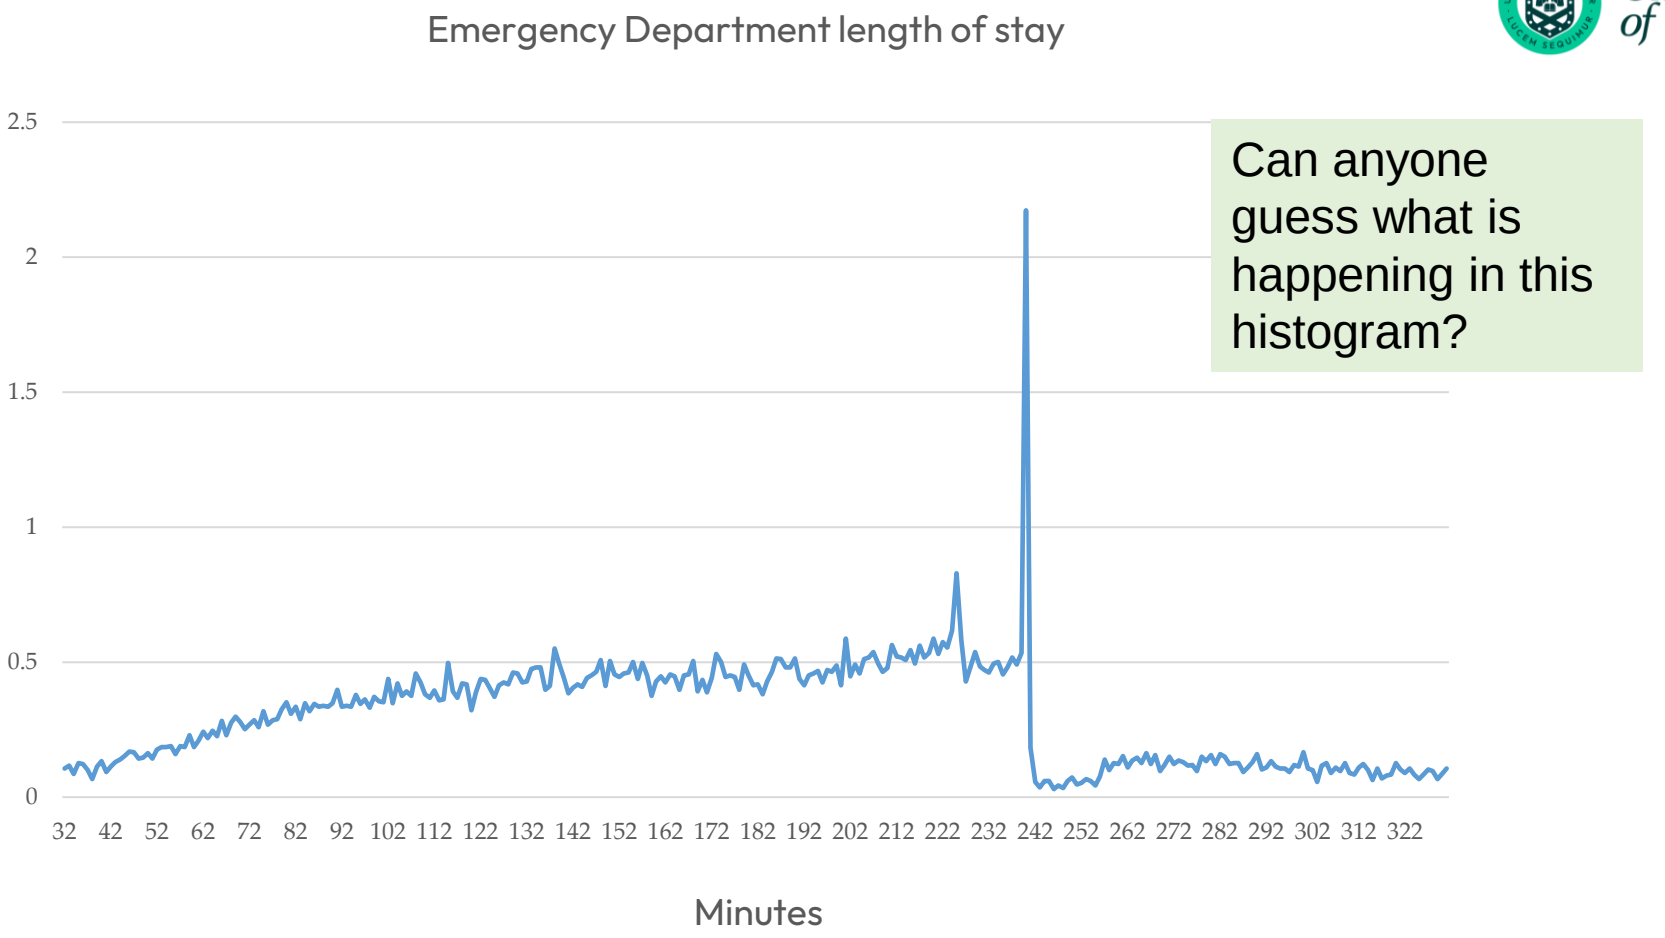

Can anyone guess what is happening in this histogram?

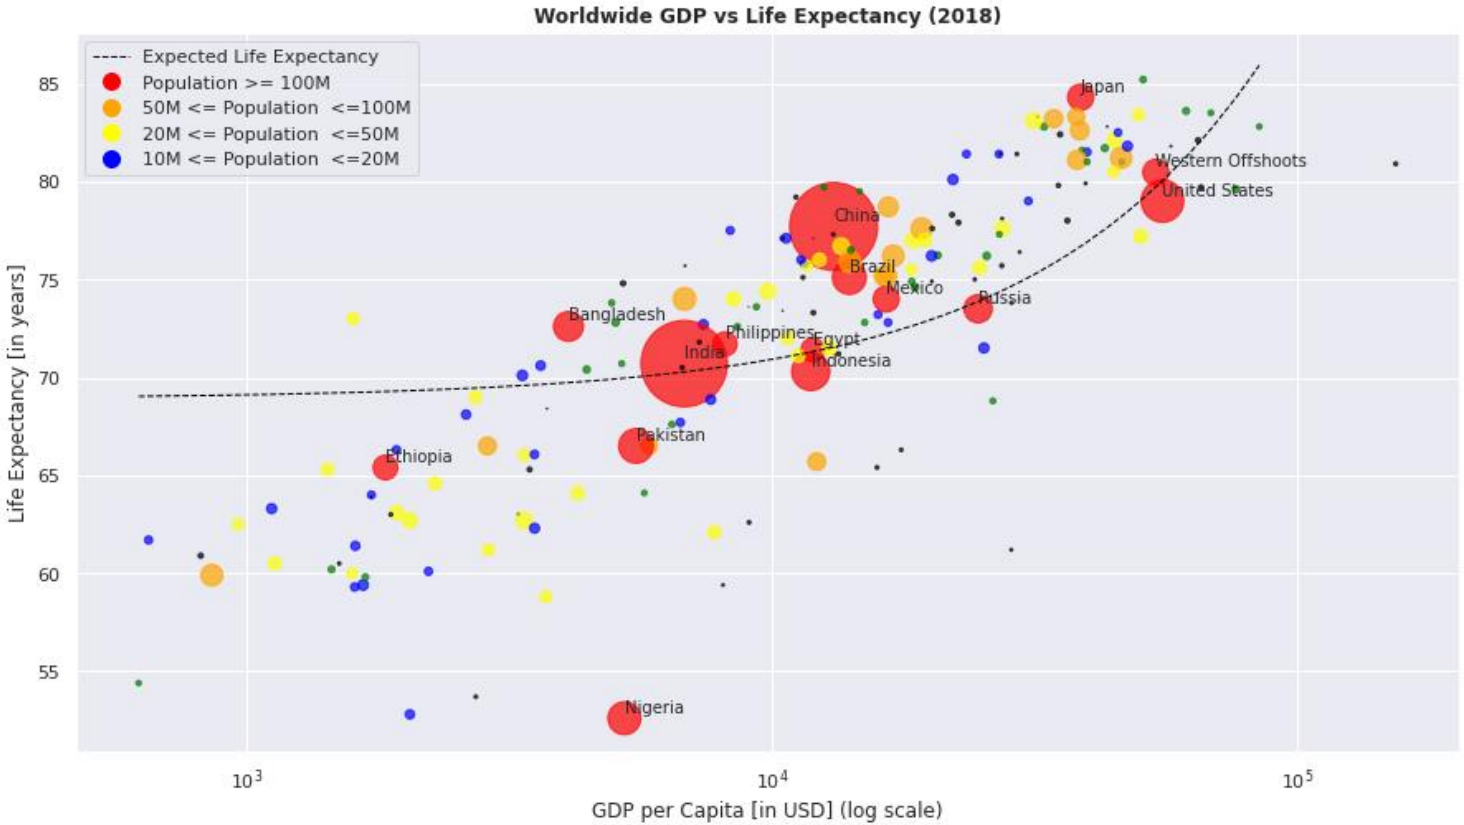

Scatterplots

Statistics

Correlation: the strength of a relationship between two variables. If two variables are correlated, as one changes in value, the other changes in the same direction.

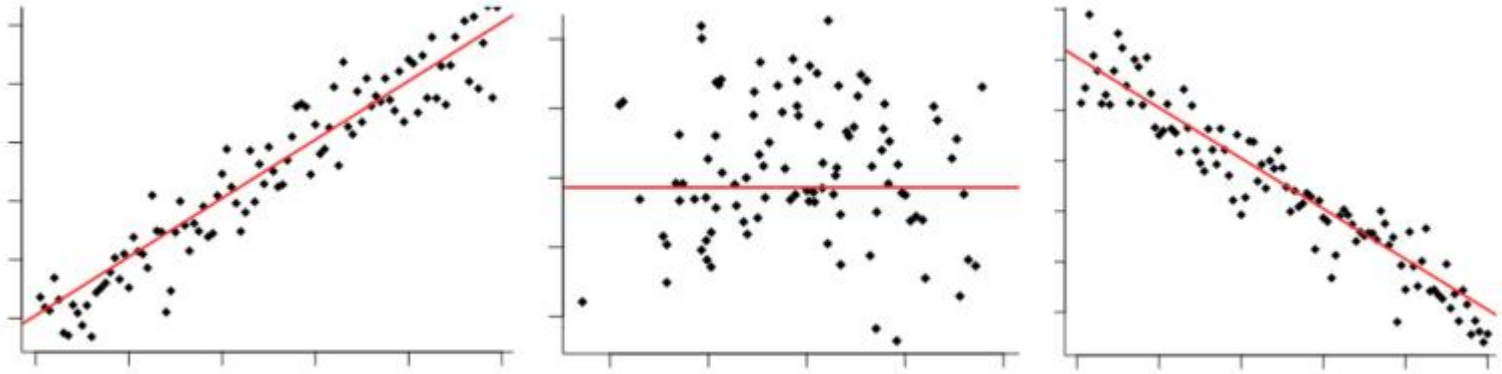

Anscombe's quartet

I		II		III		IV	
x	y	x	y	x	y	x	y
10.0	8.04	10.0	9.14	10.0	7.46	8.0	6.58
8.0	6.95	8.0	8.14	8.0	6.77	8.0	5.76
13.0	7.58	13.0	8.74	13.0	12.74	8.0	7.71
9.0	8.81	9.0	8.77	9.0	7.11	8.0	8.84
11.0	8.33	11.0	9.26	11.0	7.81	8.0	8.47
14.0	9.96	14.0	8.10	14.0	8.84	8.0	7.04
6.0	7.24	6.0	6.13	6.0	6.08	8.0	5.25
4.0	4.26	4.0	3.10	4.0	5.39	19.0	12.50
12.0	10.84	12.0	9.13	12.0	8.15	8.0	5.56
7.0	4.82	7.0	7.26	7.0	6.42	8.0	7.91
5.0	5.68	5.0	4.74	5.0	5.73	8.0	6.89

Source

Property	Value
Mean of *x* in each case	9 (exact)
Sample variance of *x* in each case	11 (exact)
Mean of *y* in each case	7.50 (to 2 decimal places)
Sample variance of *y* in each case	4.122 or 4.127 (to 3 decimal places)
Correlation between *x* and *y* in each case	0.816 (to 3 decimal places)
Linear regression line in each case	$y = 3.00 + 0.500x$ (to 2 and 3 decimal places, respectively)

Anscombe's quartet

I		II		III		IV	
x	y	x	y	x	y	x	y
10.0	8.04	10.0	9.14	10.0	7.46	8.0	6.58
8.0	6.95	8.0	8.14	8.0	6.77	8.0	5.76
13.0	7.58	13.0	8.74	13.0	12.74	8.0	7.71
9.0	8.81	9.0	8.77	9.0	7.11	8.0	8.84
11.0	8.33	11.0	9.26	11.0	7.81	8.0	8.47
14.0	9.96	14.0	8.10	14.0	8.84	8.0	7.04
6.0	7.24	6.0	6.13	6.0	6.08	8.0	5.25
4.0	4.26	4.0	3.10	4.0	5.39	19.0	12.50
12.0	10.84	12.0	9.13	12.0	8.15	8.0	5.56
7.0	4.82	7.0	7.26	7.0	6.42	8.0	7.91
5.0	5.68	5.0	4.74	5.0	5.73	8.0	6.89

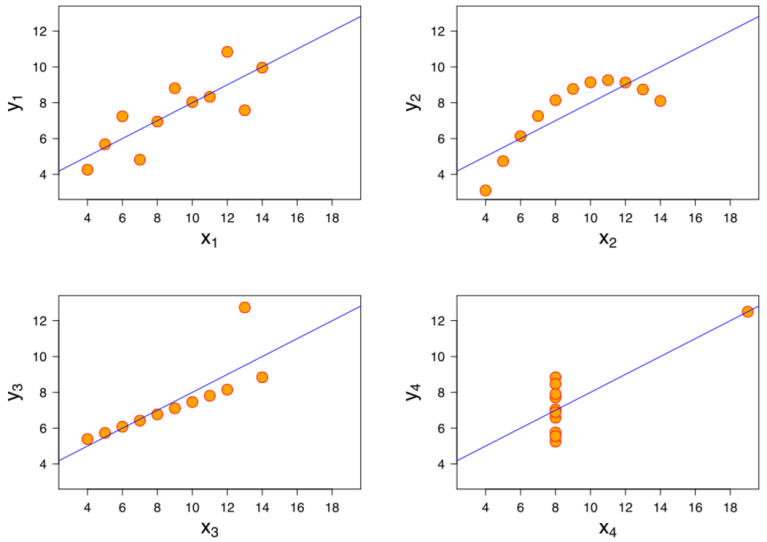

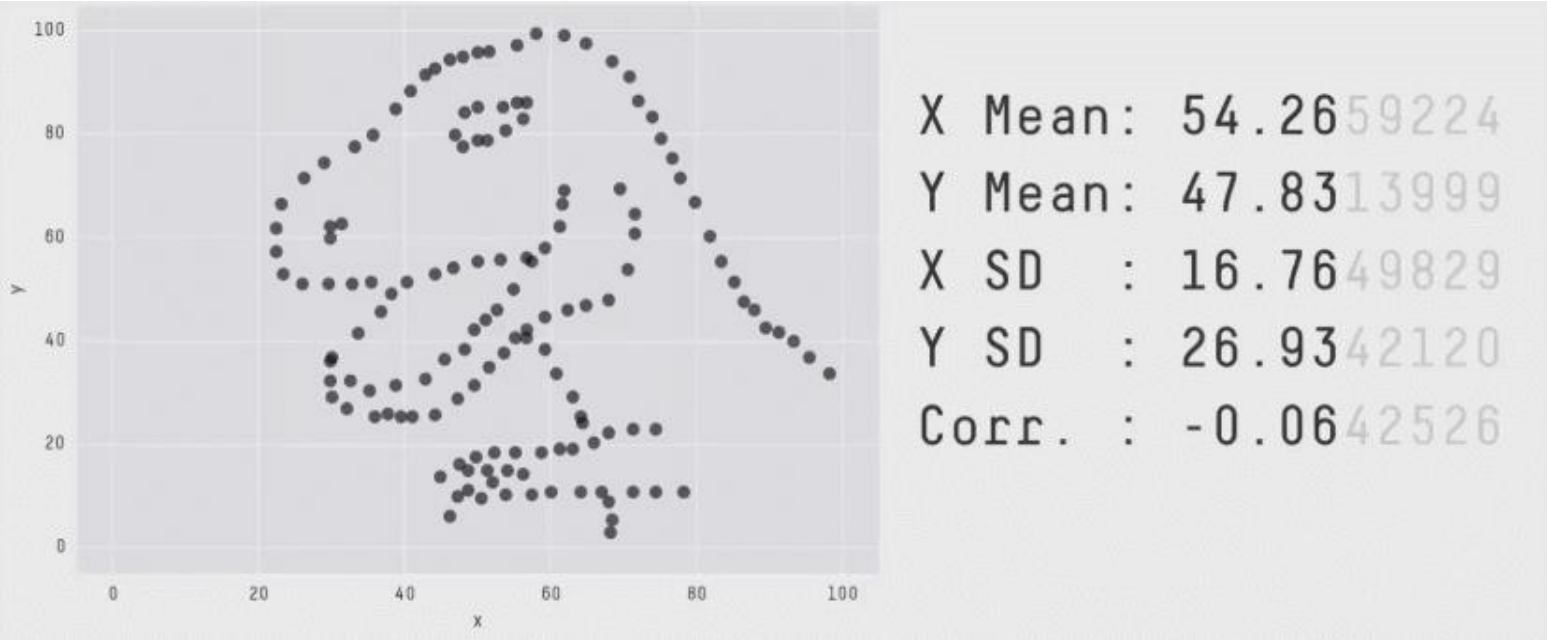

University of Exeter

Visual vocabulary

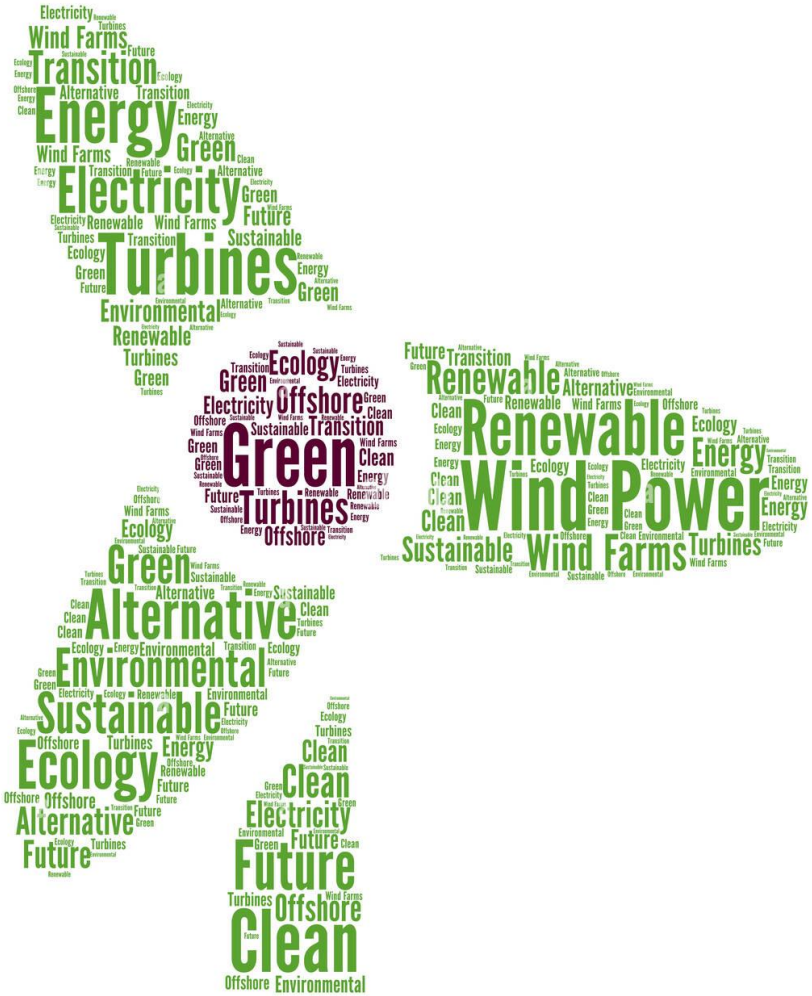

Visually Encoding Data

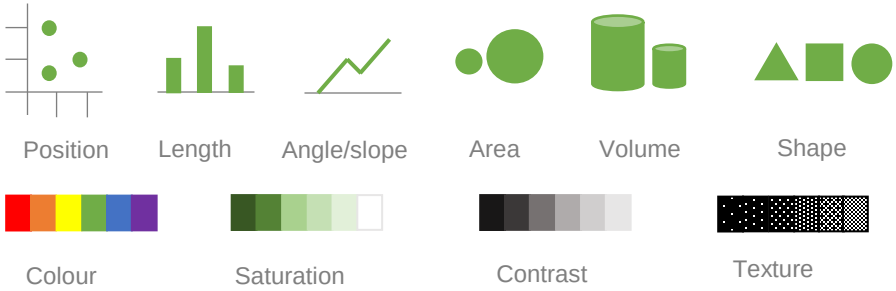

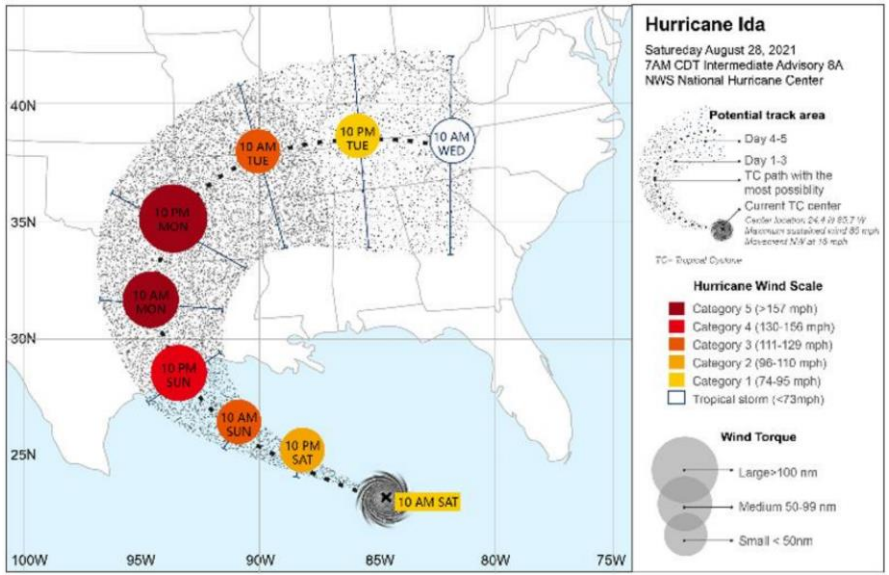

Data storytelling allow us to use data sets to convey insights effectively.

Visual data encoding helps people to quickly understand data by enhancing **pattern recognition**, **reducing cognitive load** by making intuitive insights, and **preventing misinterpretation**.

The choice of encoding type can significantly impact how the audience perceives and interprets the data story.

Effective encoding captures attention, clarifies complexities, and emphasizes key messages.

Ma, Q., & Millet, B. (2022). Design guidelines for hurricane risk forecast to non-expert users. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* (Vol. 66, No. 1, pp. 2031-2035). Sage CA: Los Angeles, CA: SAGE Publications.

Why do people hate Pie Charts?

Pie charts aren't *bad* visualizations. They just need to be used appropriately.

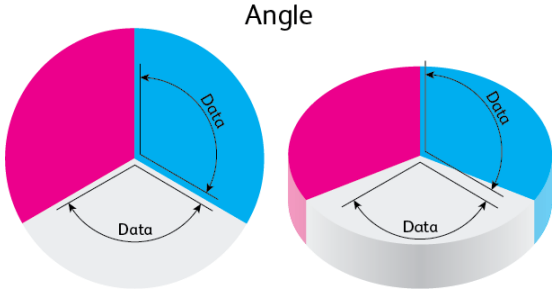

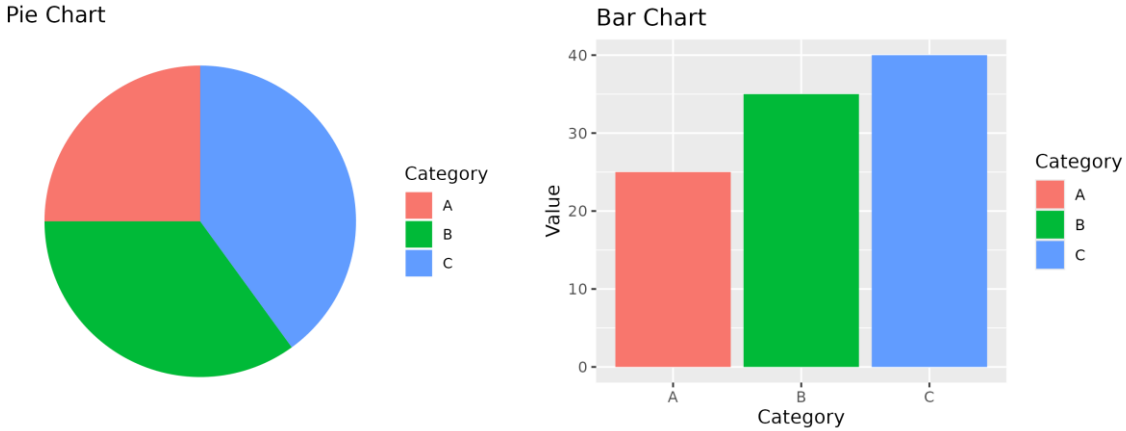

The data in a pie chart is encoded in **the angle** of the slices.

You may want to say it's encoded in the area, but if you create a pie chart by hand, what's the first thing you need to do?

How to Choose Between a Bar Chart and Pie Chart | Atlassian

Looking at Pie Charts….

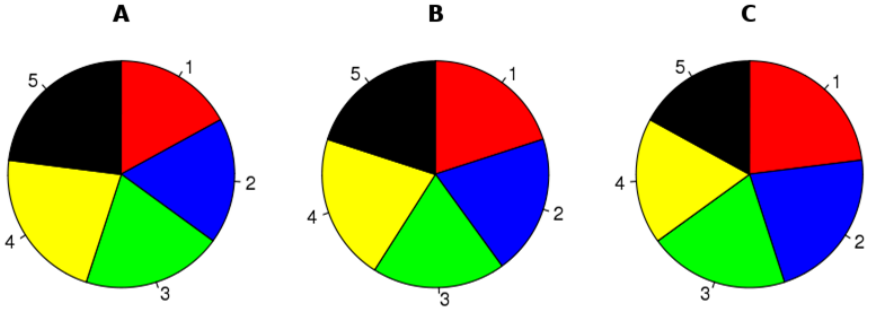

Looking at Pie Charts….

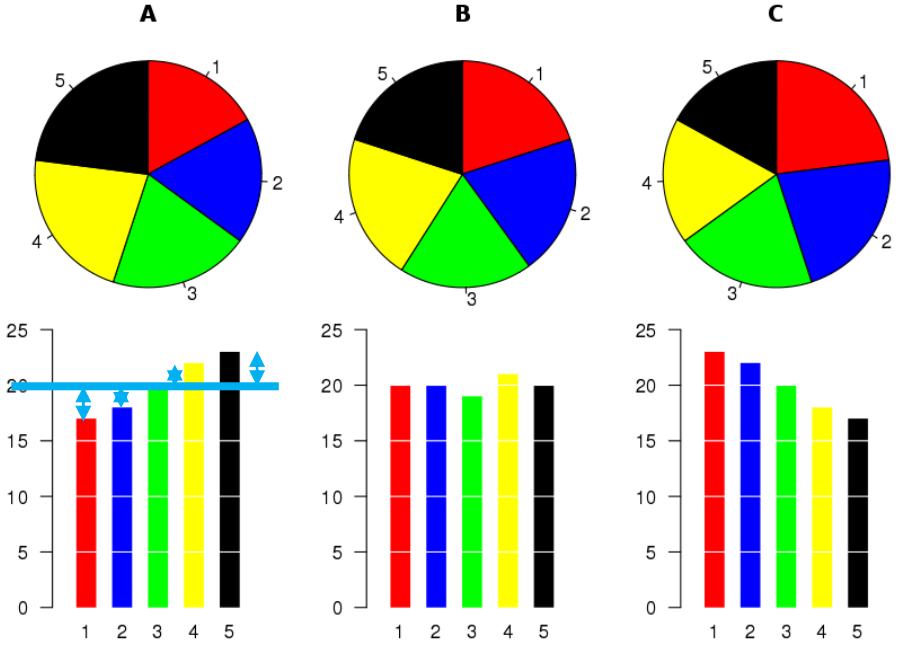

Source

Looking at Pie Charts....

If you really need to compare the differences, then pie charts aren't what you need.

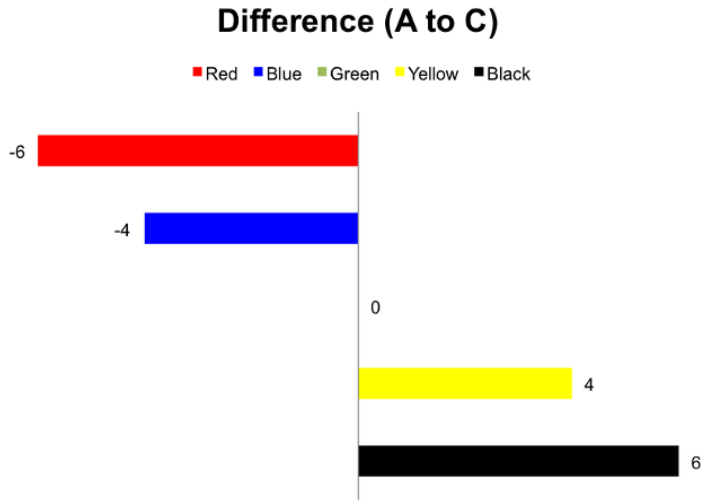

Looking at Pie Charts….

How important is it to see those small differences?

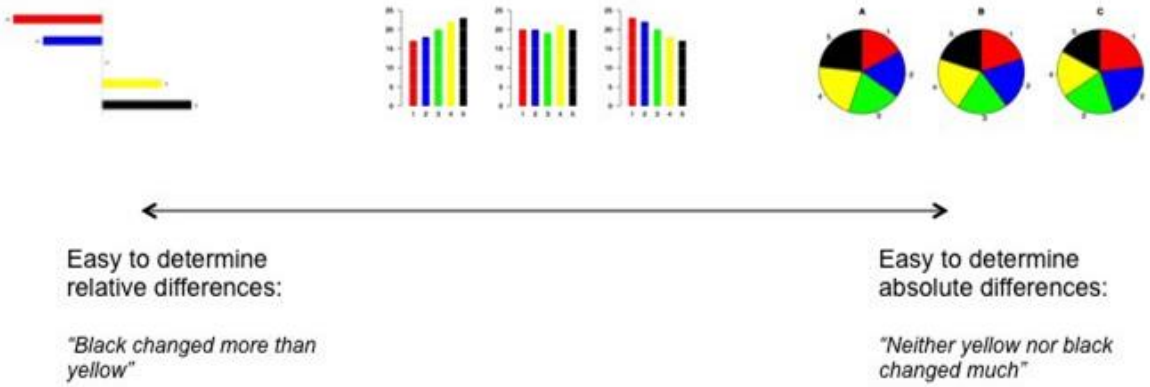

University of Exeter

Proportions. Tree maps.

Tree maps are a very useful way of visualization proportions. The data is encoded as area, and it's easier to break it up.

See this interactive map of population.

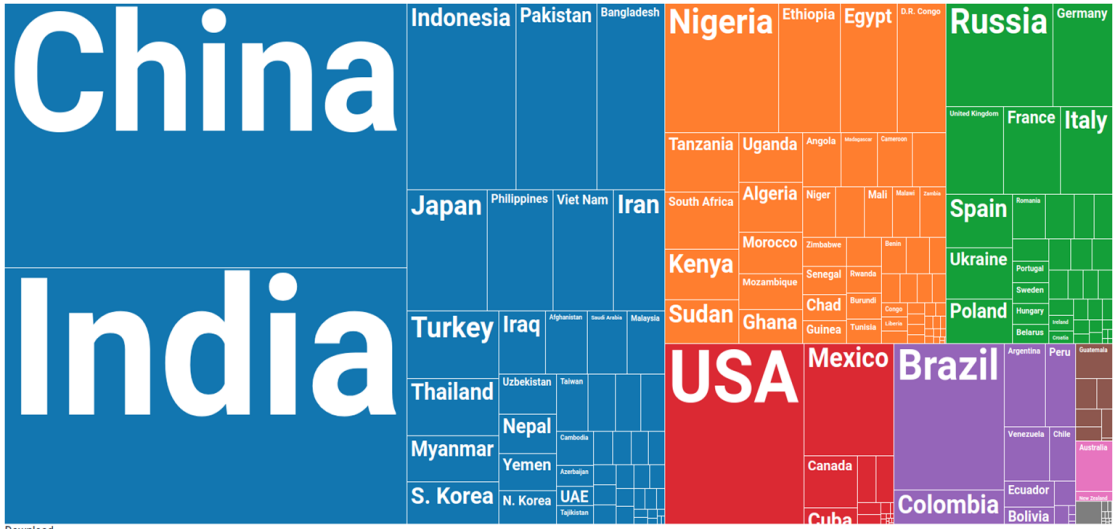

Proportions. Waffle plots.

[Five alternatives to the pie chart](#).

Waffle plots are great when you have a small amount of data. Or you just reduce a percentage to 100 boxes.

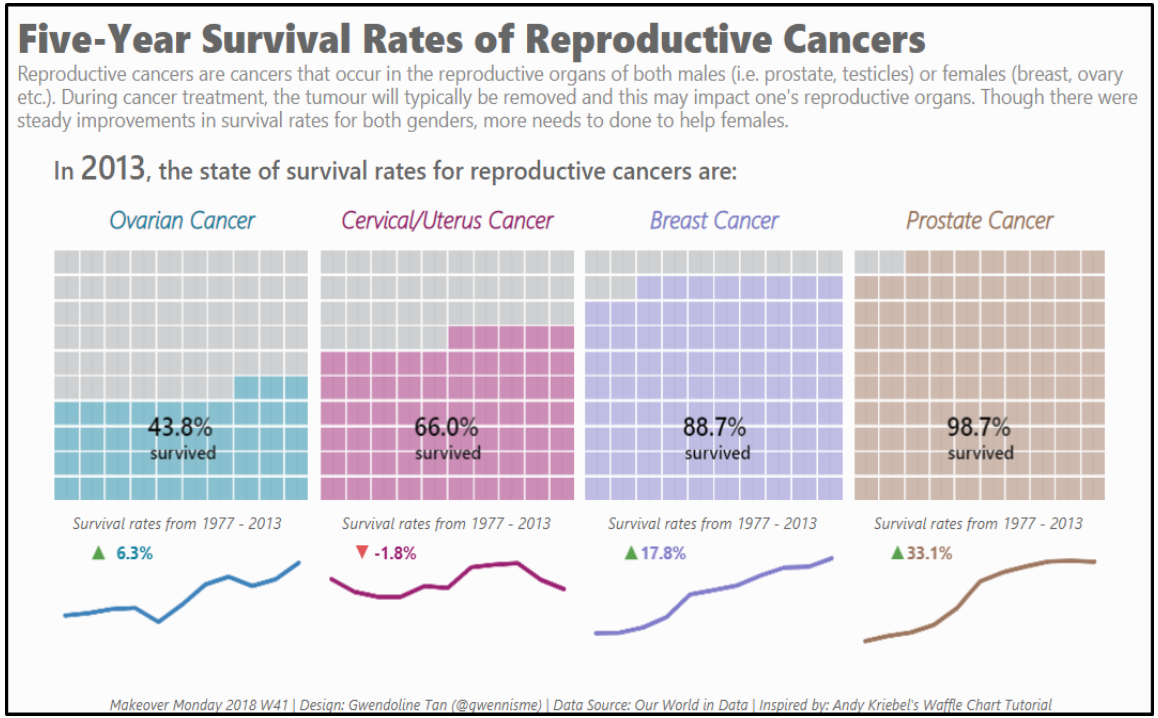

Proportions. Isotype/pictograms.

Reduce the population to 100 people. Then colour them proportionally.

[Introduction to SAMueL-1 (Stroke Audit Machine Learning) — SAMueL Stroke Audit Machine Learning 1 (samuel-book.github.io)]

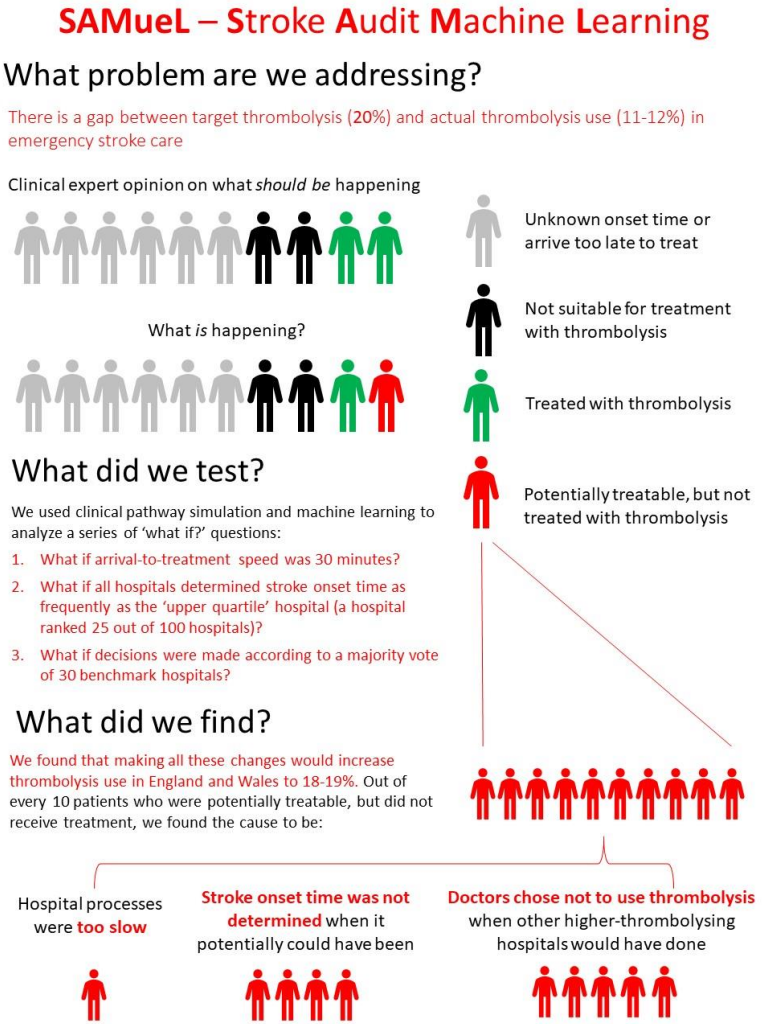

Chart Suggestions—A Thought-Starter

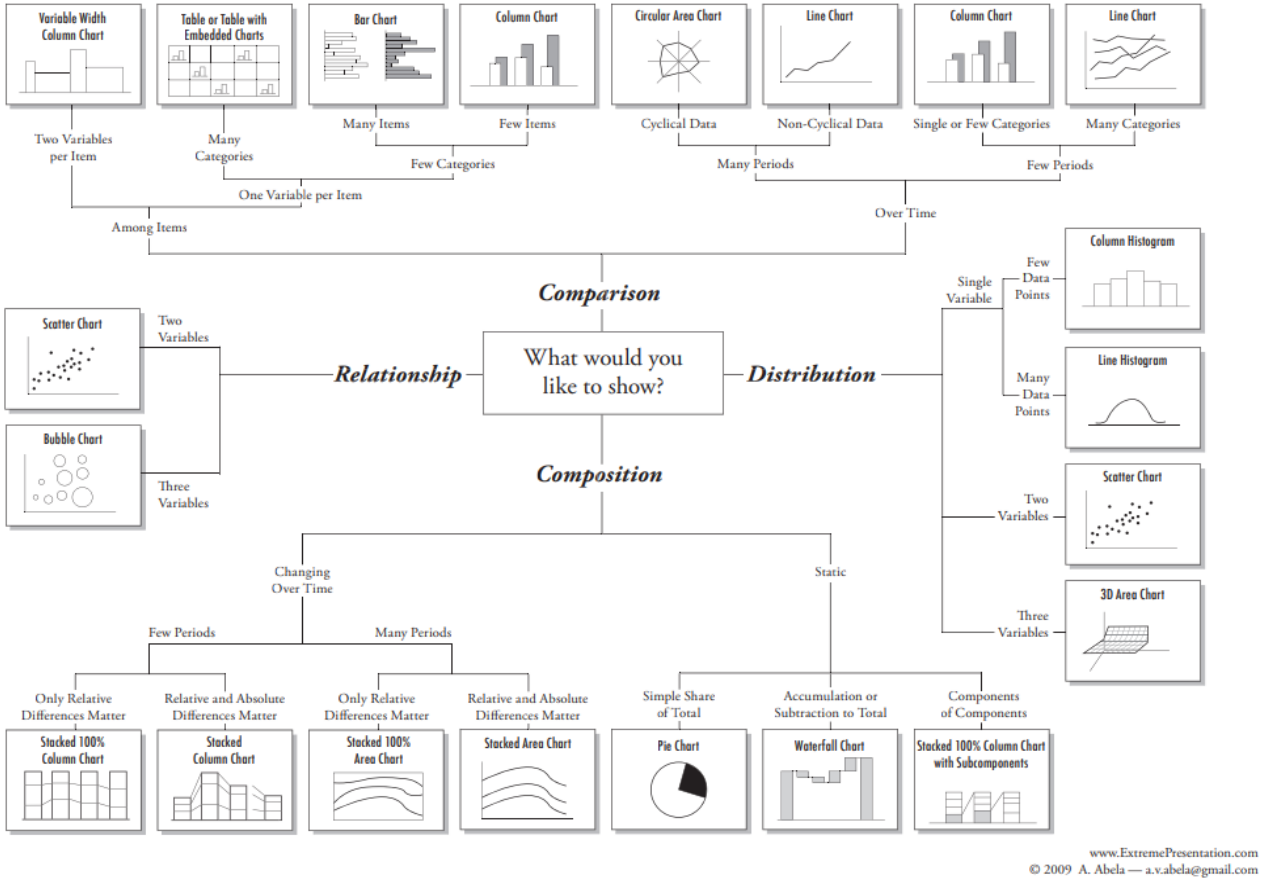

Reference: The Chart Chooser

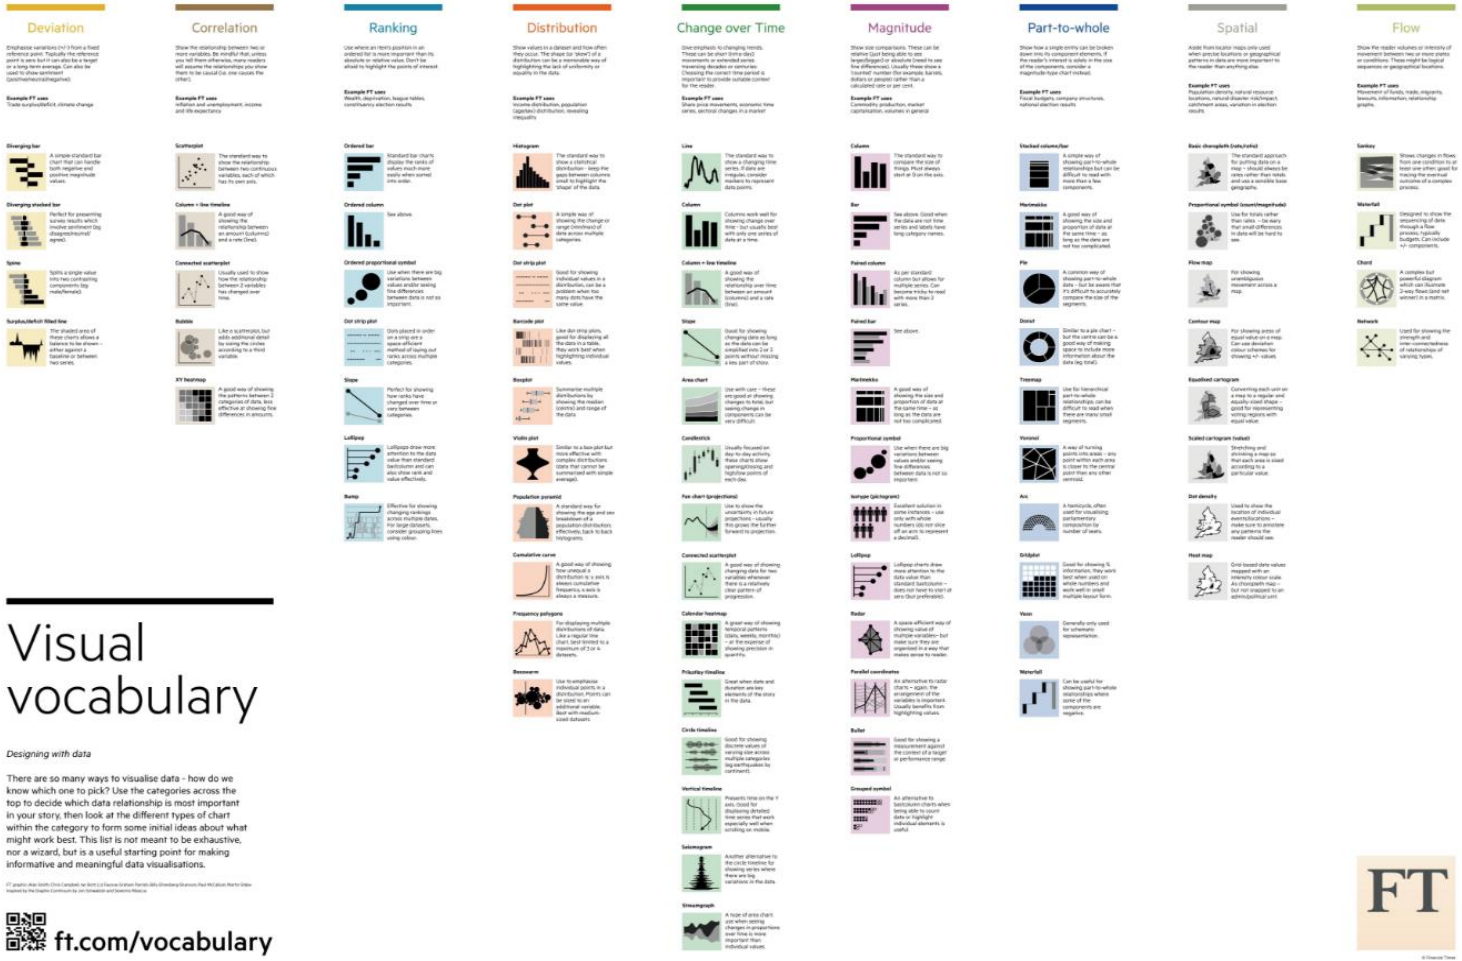

There are a huge range of different ways of showing data.

The Financial Times created this chart of a Visual Vocabulary to help give names to the different methods.

Source.

The Hans Rosling Gapminder Tool

For a good example of compressing multiple dimensions of data into a single visualization, see Hans Rosling's Gapminder tool.

Hans' talk is linked on the ELE page as well.

Gapminder has continued on without Hans' input since his passing and continued to help put world development and human experience data in context.

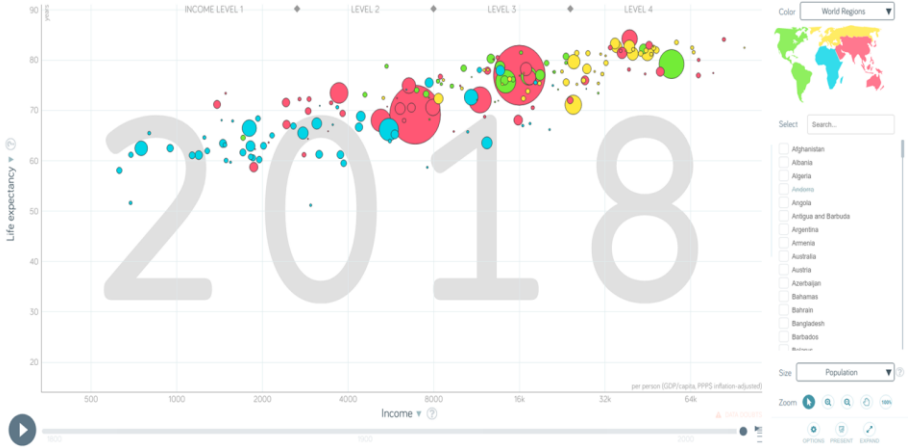

University of Exeter

Principles of visualisations!

1. Clarity and Simplicity
2. Accuracy
3. Relevance
4. Consistency
5. Hierarchy and Emphasis
6. Effective Use of Colour
7. Labelling and Annotation
8 . Sufficiency
9. Interactivity
10. Storytelling
11. Chart Selection
12. Data-Ink Ratio
13. Audience Consideration
14. Accessibility
15. Testing and Feedback
16. Ethical Considerations
17. Credible Data Sources

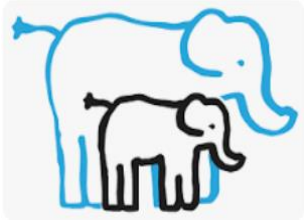

Reference: Rosling 2019

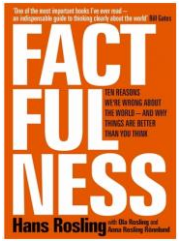

Rosling, H. (2019). Factfulness, Flammarion. Available electronically.

Principles of visualisations

1. Clarity and Simplicity: Keep the visualization simple and easy to understand. Avoid clutter, excessive decorations, and non-essential elements. The audience should be able to grasp the main message quickly.

2. Accuracy: Ensure that the data presented is accurate and that the visualization correctly represents the data. Misleading visualizations can harm the credibility of your information.

3. Relevance: Focus on the most important data and insights. Eliminate distractions and irrelevant details. Highlight what matters.

4. Consistency: Use consistent colours, scales, and terminology throughout the visualization. This helps the viewer make meaningful comparisons and understand the data more easily.

5. Hierarchy and Emphasis: Use visual hierarchy to guide the viewer's attention. Important elements should be more prominent, and less important elements should be de-emphasized.

6. Effective Use of Colour: Choose colours purposefully. Use colour to convey information, not just for decoration. Consider colourblind-friendly palettes. Too many colours can be confusing.

7. Labelling and Annotation: Clearly label data points, axes, and any relevant features. Annotations help provide context and explanations for the data.

8. Sufficiency: Provide enough data points to make the visualization informative but not overwhelming. Avoid overplotting, which can make the data hard to interpret.

9. Interactivity: For digital visualizations, interactivity can allow viewers to explore data in more detail. However, make sure it enhances understanding and doesn't create confusion.

10. Storytelling: Arrange data and visual elements in a logical sequence to tell a story. Help the viewer understand the narrative or insights you want to convey.

11. Chart Selection: Choose the right type of chart or graph for the data. Bar charts, line charts, pie charts, scatter plots, and others have different strengths for different types of data.

12. Data-Ink Ratio: Maximize the data-ink ratio, which is the proportion of ink (or pixels in digital formats) used to represent data compared to the total ink used in the visualization. Reduce unnecessary ink.

13. Audience Consideration: Understand your audience's background and familiarity with the subject matter. Adjust the complexity and terminology of the visualization to match the audience's level of expertise.

14. Accessibility: Ensure that your visualization is accessible to all, including individuals with disabilities. Use alt text for images, provide text descriptions, and follow accessibility guidelines.

15. Testing and Feedback: Test your visualization on potential users and gather feedback to make improvements. Different perspectives can help identify issues and areas for enhancement.

16. Ethical Considerations: Be mindful of the ethical implications of your data visualization, especially when dealing with sensitive or controversial topics. Avoid distorting or misrepresenting data.

17. Credible Data Sources: Clearly cite and reference the data sources used in your visualization to establish trustworthiness.

University of Exeter

Chart Junk

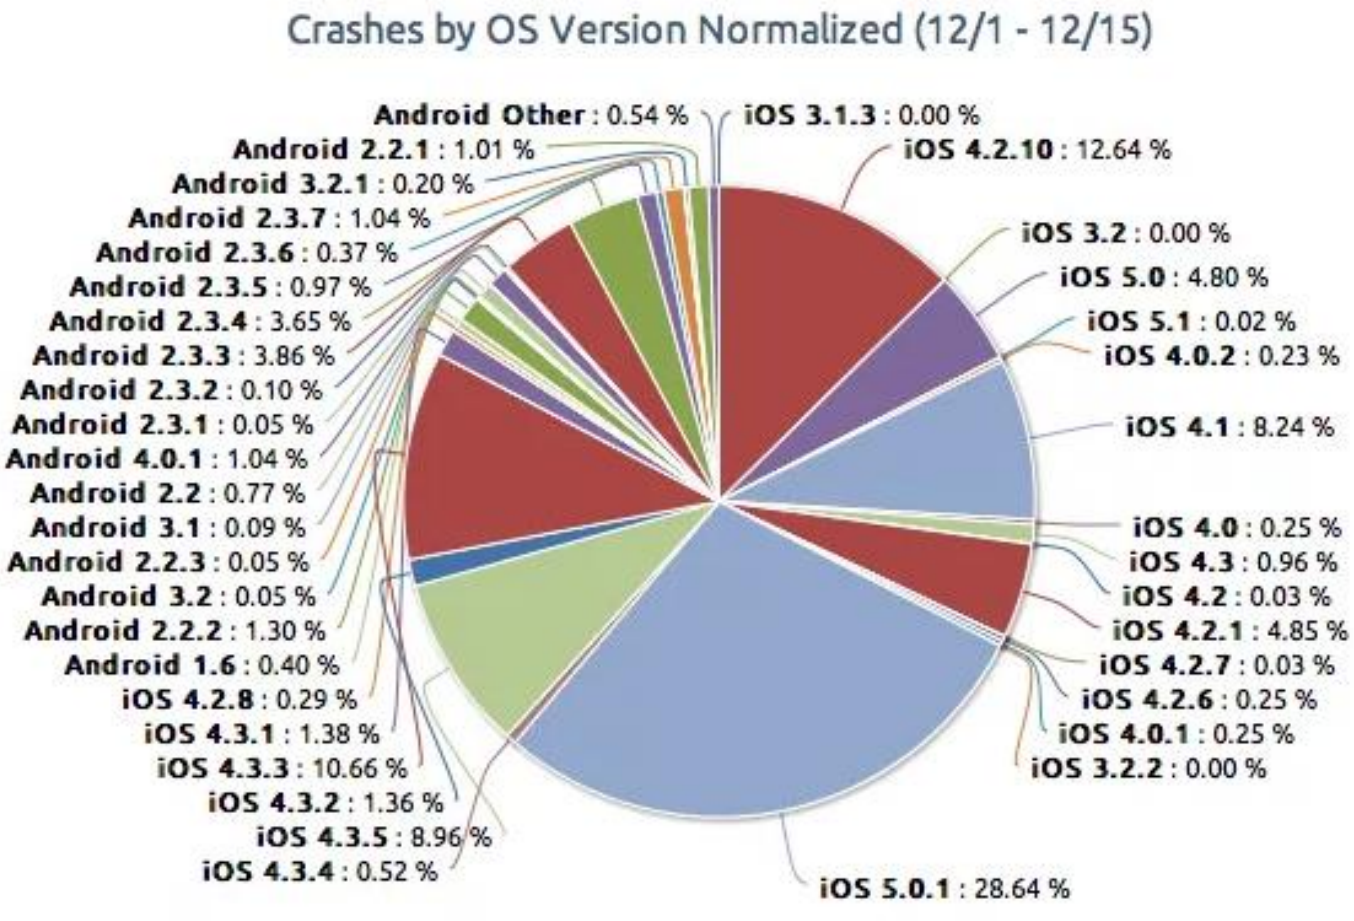
Crashes by OS Version Normalized (12/1 - 12/15)
Android Other : 0.54 %
Android 2.2.1 : 1.01 %
Android 3.2.1 : 0.20 %
Android 2.3.7 : 1.04 %
Android 2.3.6 : 0.37 %
Android 2.3.5 : 0.97 %
Android 2.3.4 : 3.65 %
Android 2.3.3 : 3.86 %
Android 2.3.2 : 0.10 %
Android 2.3.1 : 0.05 %
Android 4.0.1 : 1.04 %
Android 2.2 : 0.77 %
Android 3.1 : 0.09 %
Android 2.2.3 : 0.05 %
Android 3.2 : 0.05 %
Android 2.2.2 : 1.30 %
Android 1.6 : 0.40 %
iOS 4.2.8 : 0.29 %
iOS 4.3.1 : 1.38 %
iOS 4.3.3 : 10.66 %
iOS 4.3.2 : 1.36 %
iOS 4.3.5 : 8.96 %
iOS 4.3.4 : 0.52 %
iOS 3.1.3 : 0.00 %
iOS 4.2.10 : 12.64 %
iOS 3.2 : 0.00 %
iOS 5.0 : 4.80 %
iOS 5.1 : 0.02 %
iOS 4.0.2 : 0.23 %
iOS 4.1 : 8.24 %
iOS 4.0 : 0.25 %
iOS 4.3 : 0.96 %
iOS 4.2 : 0.03 %
iOS 4.2.1 : 4.85 %
iOS 4.2.7 : 0.03 %
iOS 4.2.6 : 0.25 %
iOS 4.0.1 : 0.25 %
iOS 3.2.2 : 0.00 %
iOS 5.0.1 : 28.64 %

Ink to Data Ratio

- The "data-ink ratio" : Edward Tufte, a prominent expert in data visualisation.
- It refers to the proportion of ink (or pixels in digital formats) used in a visualisation that is directly related to representing the data, as opposed to ink used for labels, decorations, or non-essential elements.
- It encourages the minimisation of unnecessary ink to maximise efficient and clear visualisations, remove clutter, redundant, or distracting elements.
- It encourages designers to concisely convey the information the visualization is intended to communicate. This helps viewers quickly grasp the main message and insights from the data.

Principles of Data Ink

Above all else show data.

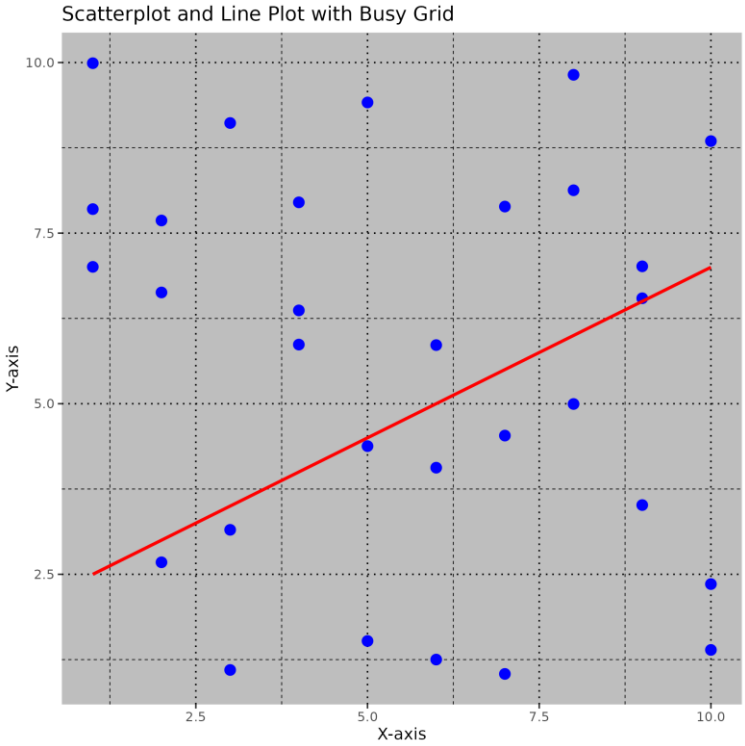

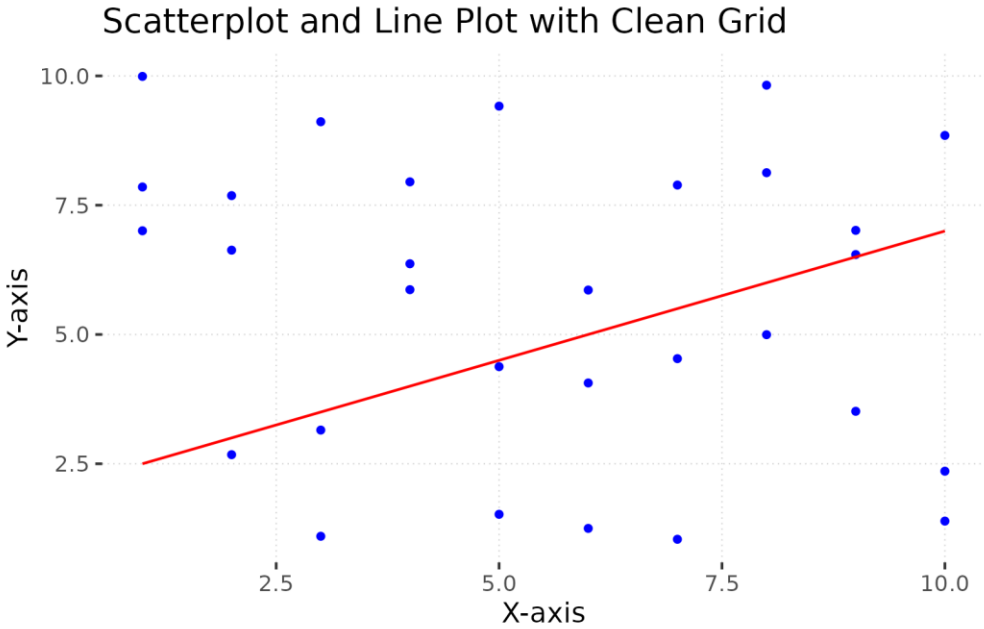

University of Exeter

Cool effects = Distorting data

What is the value for March? Can you tell?

There's some invisible tangent plane connecting to the "back" of the chart.

These charts add an extra dimension, they add complexity, but there's no information in that dimension.

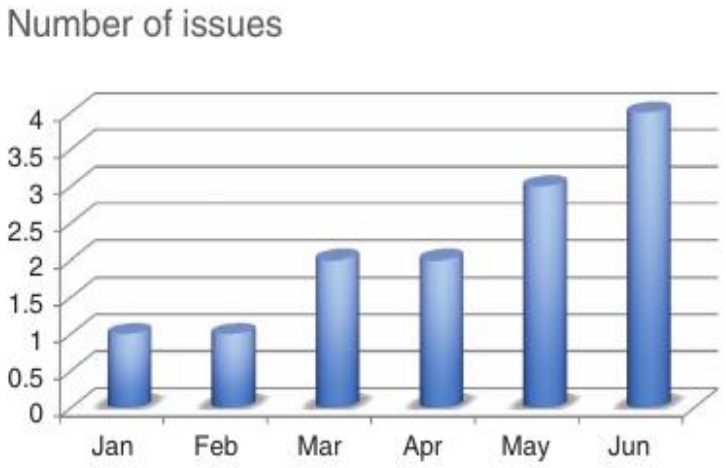

FIGURE 2.25 3D column chart

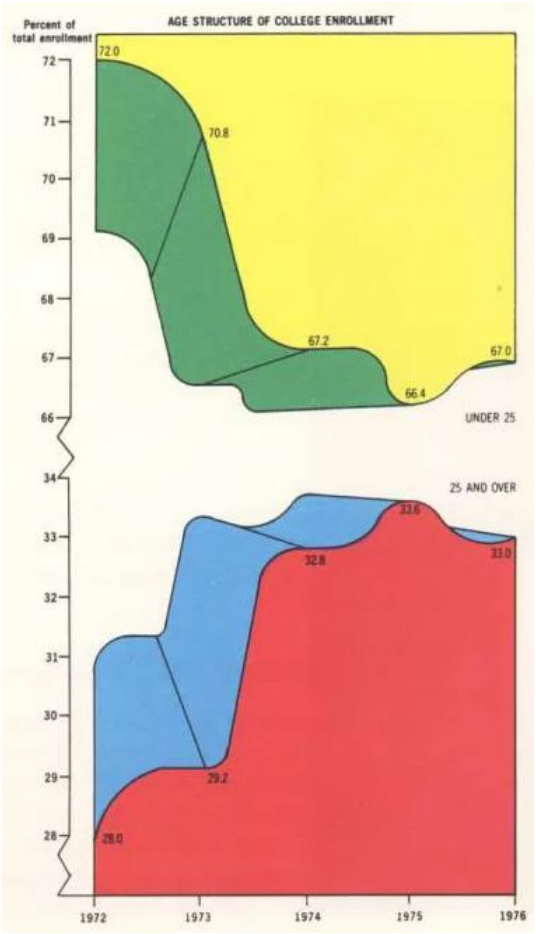

The interior decoration of graphics generates a lot of ink that does not tell the viewer anything new... Regardless of its cause, it is all non-data-ink or redundant data-ink, and it is often ***Chartjunk*** (Tufte, 2001).

Excessive and unnecessary use of graphical effects – colour, 3D effects and disguised redundancy to represent just five numbers.

Tufte (1983, p.118) says, *"This may well be the worst graphic ever to find its way into print."*

Tufte, E. R. (2001). The visual display of quantitative information, Graphics press Cheshire, CT.

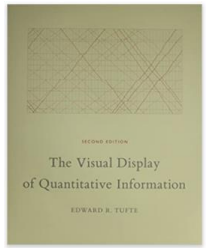

University of Exeter

Remove

to improve

(the **data-ink** ratio)

Created by Darkhorse Analytics

www.darkhorseanalytics.com

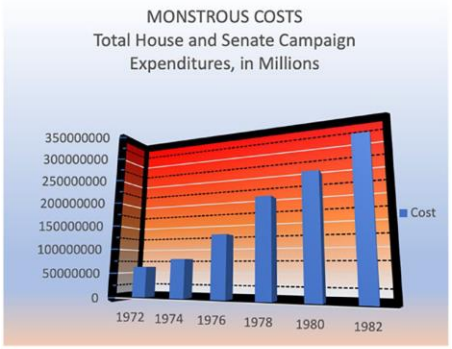

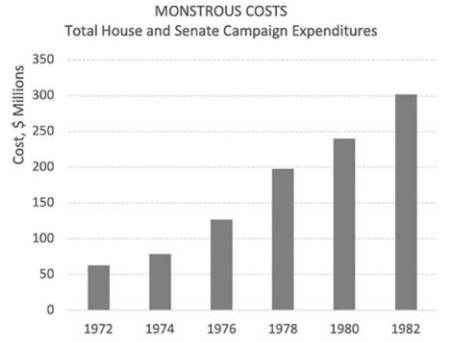

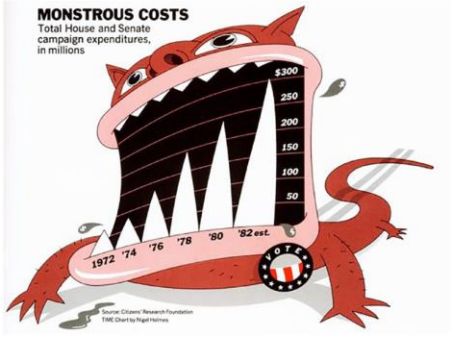

A “cluttered” visualization (top),

a minimalist “decluttered” version (middle),

and a version that incorporates pictorial embellishment (bottom).

The graph at the bottom was created by Nigel Holmes for *TIME Magazine* and was reprinted in his 1984 book, *Designer’s Guide to Creating Charts & Diagrams*.

Franconeri, S. L., Padilla, L. M., Shah, P., Zacks, J. M., & Hullman, J. (2021). The science of visual data communication: What works. *Psychological Science in the public interest*, *22*(3), 110-161.

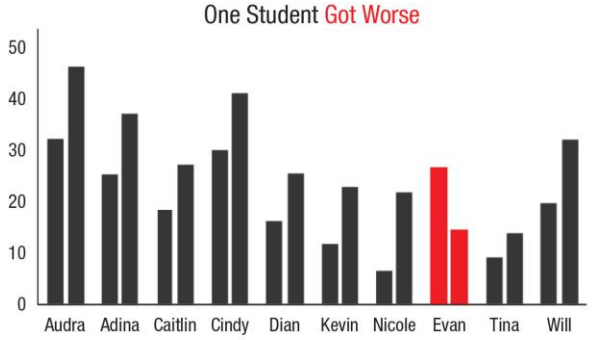

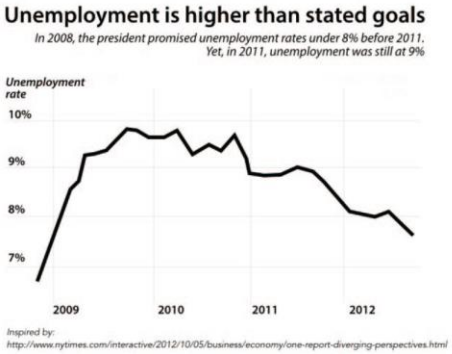

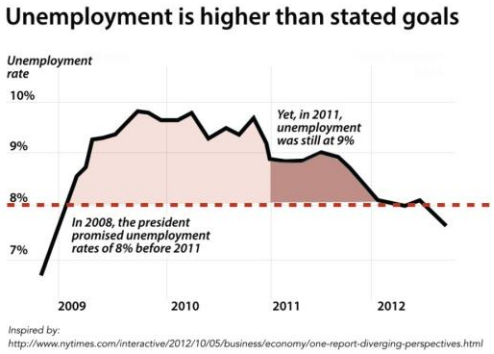

The graphic at the top illustrates a colour-highlighting technique suggested in business-oriented practitioner guides (e.g., Knaflic, 2015).

The graphs at the bottom (inspired by Bostock et al., 2012) are an adaptation of a graph by data journalists using grouping, highlighting and verbal annotation

Franconeri, S. L., Padilla, L. M., Shah, P., Zacks, J. M., & Hullman, J. (2021). The science of visual data communication: What works. *Psychological Science in the public interest*, *22*(3), 110-161.

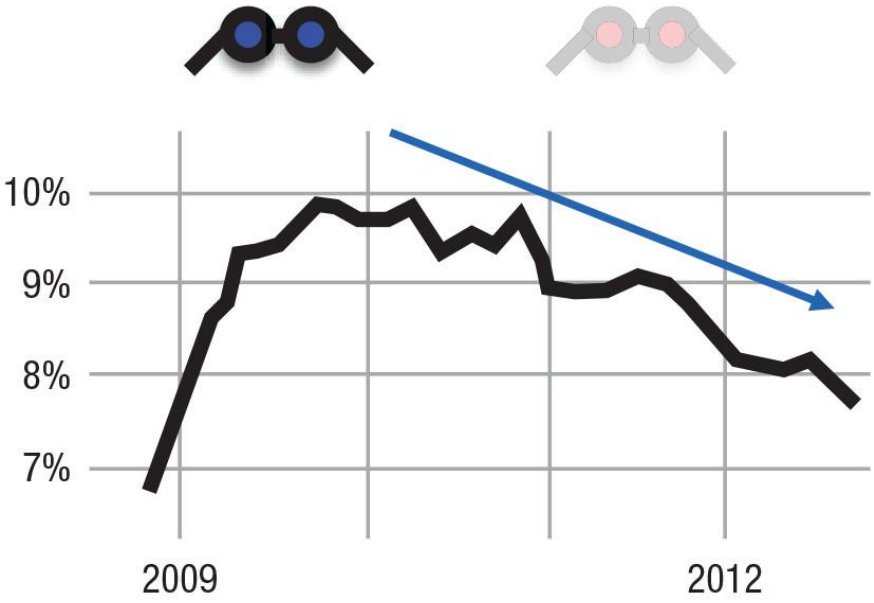

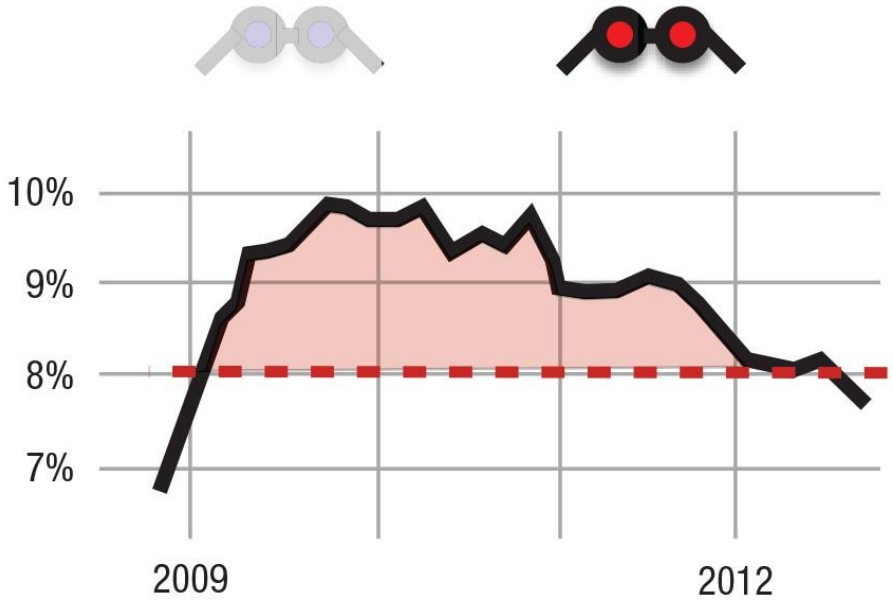

An example of emphasizing different perspectives in a single data set (inspired by Bostock et al., 2012). One data set can be seen with dramatically different perspectives, depending on which patterns an observer does and does not extract.

Using graphs to mislead

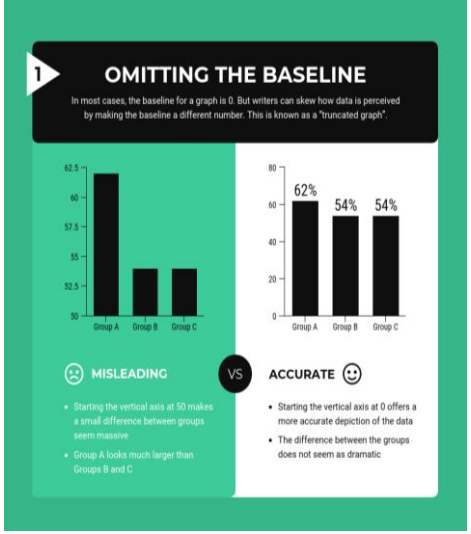

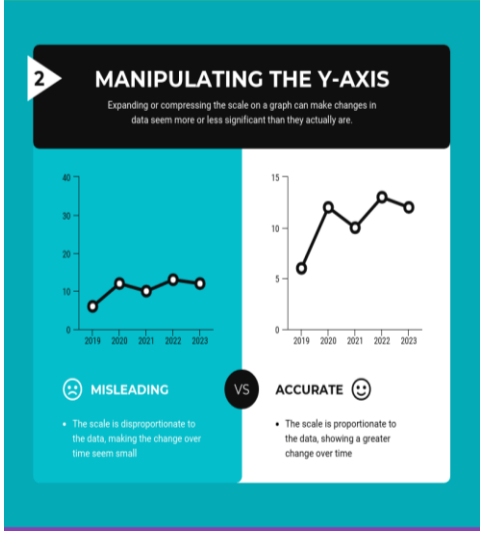

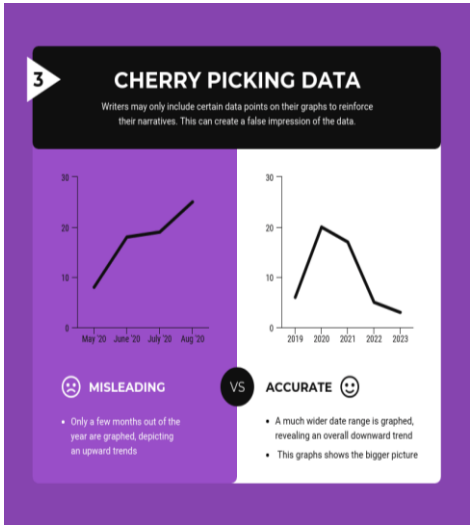

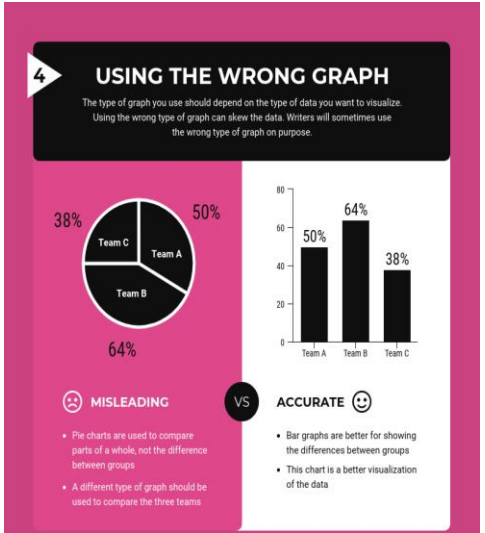

Source

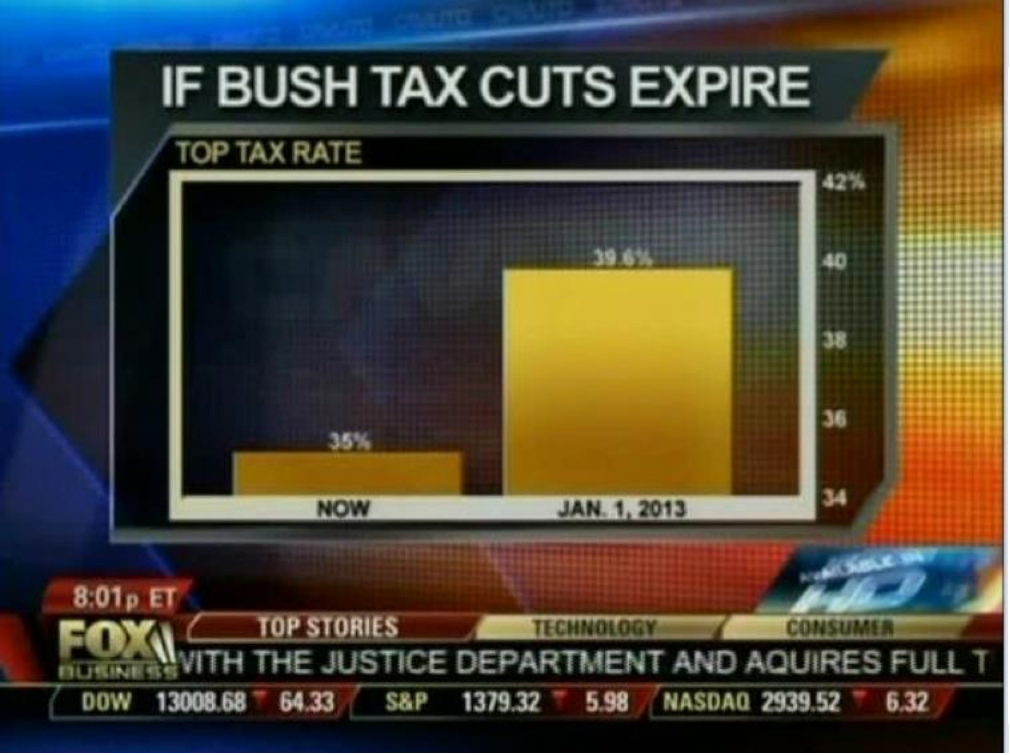

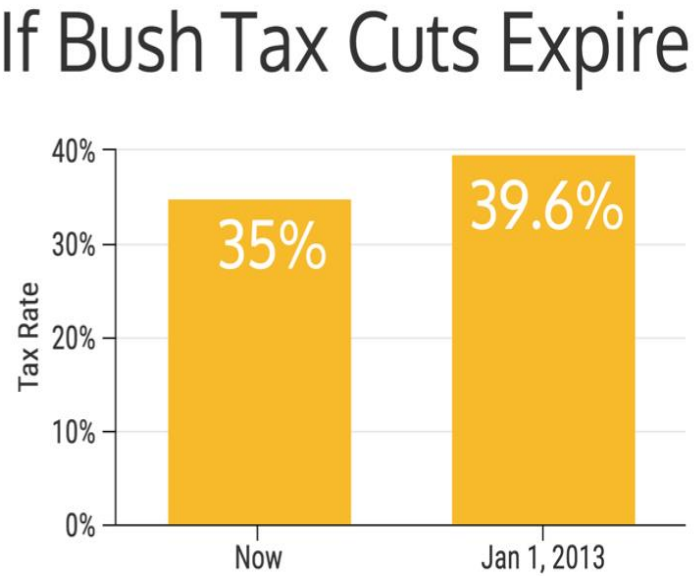

Deceptive Designs

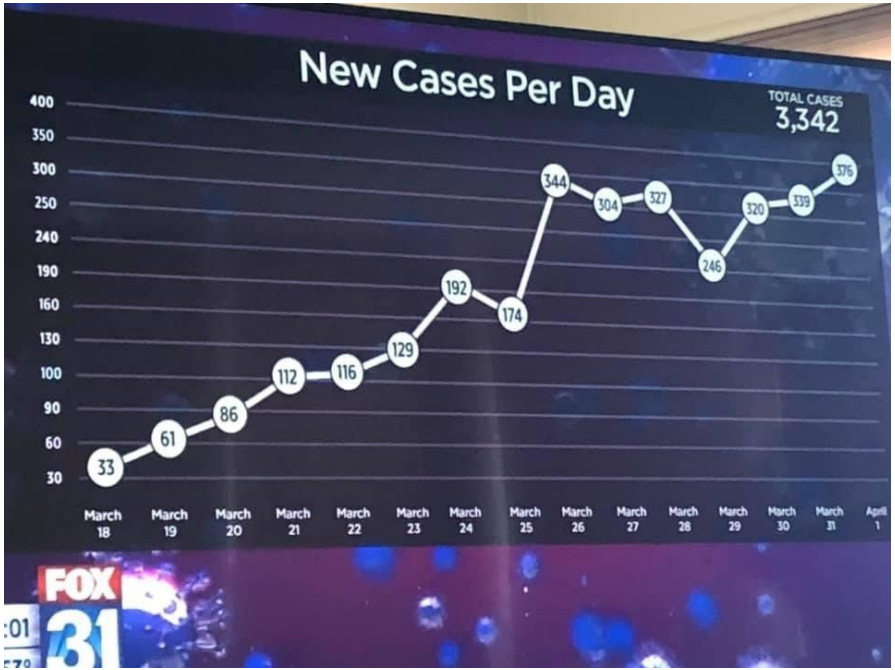

Deceptive Design

The gridlines are equally spaced on the page, but sometimes the same space represents 30 people, sometimes 10, and sometimes 50.

It seems to be completely arbitrary.

How to make that crazy Fox News y axis chart with ggplot2 and scales (freerangestats.info)

Deceptive Design

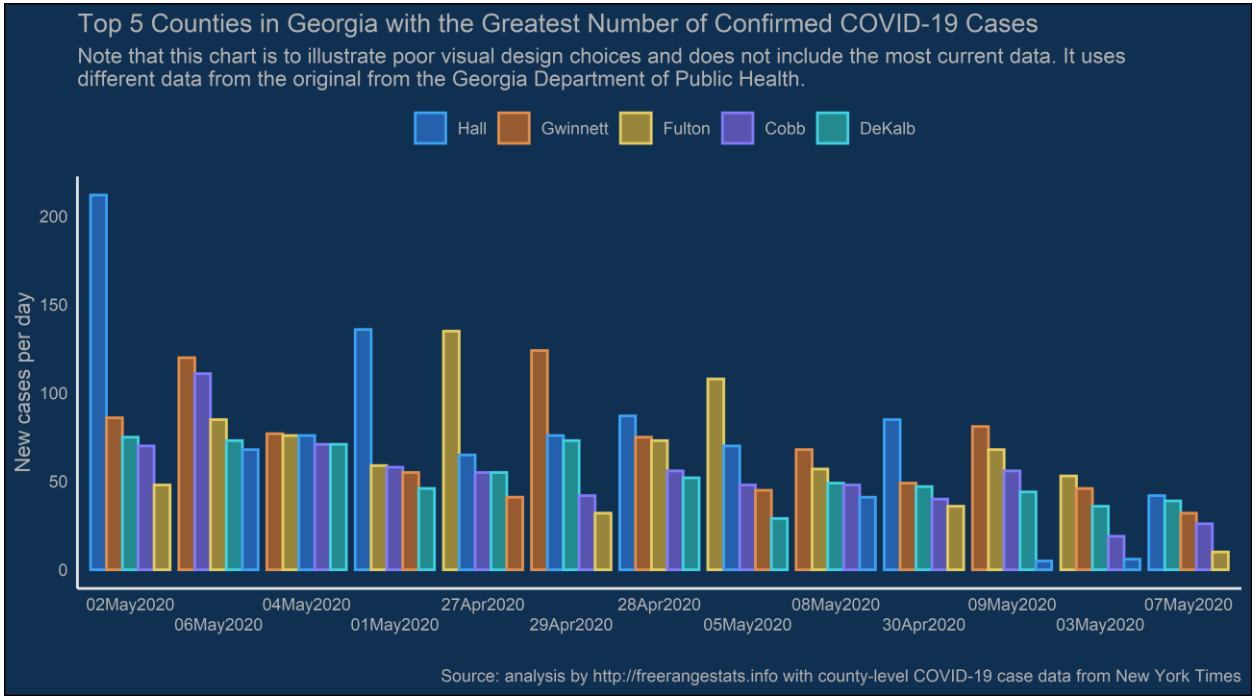

As a visualisation at least three things are wrong:

1. Dodged bar charts are rarely effective for making comparisons over time - it's difficult for the eye to follow;

2. Within each day's clump of bars, the counties are in a different order (highest to lowest, within the clump), reducing the meaning in the pattern in each clump;

3. The daily clumps of bars are not in chronological order.

Ordering bars within their clumps in a bar chart (freerangestats.info)

Absolute Precision Ranking for Seeing a Single Ratio

Visual estimation of the 1:7 ratio is noisier toward bottom

Highest

Position

1 7

Length

1 7

Area

1 7

Angle

1 7

Intensity

1 7

Lowest

Common Illusions That Distort Data

Caveats for the visual encoding in each row

99 98 a b 99 98 a b

Use caution with nonzero axes: Viewers tend to overestimate differences… even when the nonzero base is marked, as in the examples at left.

Stacked bar: Bars on baseline are position-coded = more precise perception.

The black & dark gray bars have the same value differences among them, but the differences are only visible across the black bars.

1 7

Sure, looks like a ~1:7 area ratio.

But beware: if the number is actually mapped to *Length*, the gray circles depict ~1:2.5 ratio.

The difference is larger for the lighter segments compared with the darker ones, right?
That is an illusion—the differences are identical.

Intensity values can look different depending their backgrounds.

Do not plot intensities on intensities.

Vision Is Powerful for Global Statistics

For each visualization, statistics are available quickly

Dot Plot

a b c d e f

Max height
Mean height
Min height

Stacked Bar

a b c d e f

Min
Max
Mean length of dark bars

Bubble Map

Mean Area
Min
Max

Slope Graph

Max
Mean Angle
Min

Heat Map

Min Max

Mean Intensity

Vision Is Sluggish for Comparisons

Isolating pairs with "larger second values" is tough…

So guide viewers to the right comparisons

a b c d e f

Tool: Shortcut comparisons by adding direct depictions of the deltas, as below

Δ#

"a, c, & e have increased"

Tool: Highlight and annotate the right comparisons for your viewers, as above.

Tool: You and your viewers will (generally) compare values that

(a) are close together or connected and
(b) have similar colors,

in that priority order.

For color heat maps, depict deltas as blue (+) & red (–)

[green/red is unsafe for colorblindness]

University of Exeter

The two columns on the left show a quick reference guide to channels that can depict data visually and common illusions for each channel.

The column in the centre presents a summary of how visual statistics are powerful.

The two columns on the right illustrate how comparisons are severely limited and present a set of design techniques that focus viewers on the "right" ones.

Franconeri, S. L., Padilla, L. M., Shah, P., Zacks, J. M., & Hullman, J. (2021). The science of visual data communication: What works. *Psychological Science in the public interest*, *22*(3), 110-161.

Next Week: Clusters and Similarity

- Read Data Science for Business, chapter 6 (book available electronically).
- Watch: StatQuest: K-Means clustering
- Watch: StatQuest: Principal Component Analysis (PCA) Step-by-Step
- Watch: StatQuest: PCA Main Ideas
- Play: Visualizing K-Means Clustering
- Play: Visualizing DBSCAN
- Play: Principal Component Analysis

University of Exeter

Any questions?

?