SRI SATHYA SAI INSTITUTE OF HIGHER LEARNING
(Deemed to be University)
Department of Mathematics and Computer Science

LECTURE NOTES — BOOK 3

Designing Tables and Graphs to Enlighten
Source text: Show me the Numbers (Stephen Few)

UDSC-402: Data Analysis and Visualization	Instructor: Dr. D Bhanu Prakash	Date: 29-08-2026

Chapter Map

Ch.	Title	Central Question Answered
1	Introduction	Why does the design of quantitative displays matter at all?
2	Simple Statistics to Get You Started	What kinds of numbers and relationships are we displaying?
3	Differing Roles of Tables and Graphs	Table or graph — which do I choose?
4	Fundamental Variations of Tables	What can a table show, and how should it be structured?
5	Visual Perception and Graphical Communication	How does the eye–brain system actually read a graph?
6	Fundamental Variations of Graphs	What can a graph show, and which graph fits which message?
7	General Design for Communication	What design rules apply to tables, graphs and text alike?
8	Table Design	How do I make a table quick and accurate to read?
9	General Graph Design	What two rules must never be broken in a graph?
10	Component-Level Graph Design	How do I design each individual part of a graph?

Chapter 1 — Introduction

1.1 The Core Problem

- Tables and graphs are used everywhere in organisations, yet almost no one has been trained to design them. Poor design is therefore the norm, not the exception.

- Contrary to popular wisdom, numbers cannot always speak for themselves. Inattention to the display of quantitative data produces large but hidden costs — time wasted struggling to understand numbers instead of acting on them.
- Leveraging numbers moves us out of assumption, feeling, guesswork, gut instinct, intuition and bias, and into reliable fact based on measurable evidence.
- Karen A. Schriver: "Poor documents are so commonplace that deciphering bad writing and bad visual design have become part of the coping skills needed to navigate the so-called information age."

1.2 The Purpose of Quantitative Displays

> ***The purpose of quantitative tables and graphs is to provide the reader with important, meaningful and useful insight. That is it — not to entertain, not to indulge in self-expression, and not to make boring numbers interesting through flash and dazzle.***

- Edward R. Tufte: "The overwhelming fact of data graphics is that they stand or fall on their content, gracefully displayed… Above all else show the data."
- Tufte again, on dressing numbers up: "If the statistics are boring, then you've got the wrong numbers."
- The goal in the design of quantitative information is excellence, because the means to attain it are readily available.
- The rise of tables and graphs is a consequence of the personal computer and spreadsheet software (e.g. Lotus 1-2-3, Excel). Software made production easy — perhaps too easy — and something produced with a computer wrongly acquires an air of authenticity and quality.

1.3 Diagnosing a Bad Display — the Worked Examples

Example in the Book	What Is Wrong	The Remedy
3-D coloured QTD-sales bar graph (Fig. 1.1)	Great visual impact but almost no message; 3-D, shading, textured background; no unit of measure, no date, no comparison to plan or to prior period; a great deal of ink saying very little.	A small, plain table (Fig. 1.2) carrying current sales, % of total, % of quarter plan, projected quarter-end sales and projected % of plan — more information, less ink.
PBS/Pew table of favourable-vs-unfavourable opinion of the U.S.A. (Fig. 1.3)	A dense table used for a message that is about pattern, not look-up. No picture emerges however long you struggle.	Three side-by-side horizontal bar graphs (Fig. 1.4) sorted by favourable opinion. The story comes alive; Jordan's missing "undecided" bar pops out as empty space.
Nine-product clustered vertical bar graph (Fig. 1.5)	Relies entirely on its title to declare its purpose; the design subverts the	Divide and conquer: eight small graphs (Fig. 1.6), one per competing product, each with a

Example in the Book	What Is Wrong	The Remedy
	reader's ability to see what the graph is for.	black reference line for SlicersDicers.
Pie chart of market share, plain and then 3-D (Figs. 1.7, 1.8)	You cannot read one company's share, cannot rank it, and 3-D perspective further distorts perception.	A sorted horizontal bar graph with values printed beside the bars (Fig. 1.9): "Our Company" ranks second at exactly 13.46%.

1.4 Key Concepts Introduced

- Affordance — a characteristic of an object that reveals how it is supposed to be used and makes it easy to use in that way (a teapot has a handle; a push-door has a push-plate). A well-designed graph must, by its design alone, suggest its own use. (Donald A. Norman, The Design of Everyday Things, 1988.)
- Divide and conquer — replacing one over-crowded graph with a series of small, related graphs, each carrying one comparison.
- Reference line — a line (e.g. the black SlicersDicers line in Fig. 1.6) that acts as a powerful affordance, telling the reader what to compare against what.
- Eloquence in simplicity — the recurring theme of the whole book.

1.5 Scope of the Book

The book is limited to content that is (1) practical and applicable to typical organisational reporting, (2) focused on communication, and (3) focused on effective design.

Tables and graphs of quantitative business information serve four purposes:

Purpose	Meaning
Analysing	Using a display to discover the message that is hidden in the data.
Communicating	Using a display to pass a message on to others. (This book's focus.)
Monitoring	Using a display to track performance — speed or quality of manufacturing, for example.
Planning	Using a display to predict and prepare for the future.

Intended readers: anyone whose work involves presenting quantitative information — financial, business, data and decision-support analysts, managers, graphic artists and everyone else in between.

1.6 Content Preview — the Verbal/Visual Distinction

- Charts is the larger family of display methods. It includes tables, graphs, diagrams (which illustrate a process or set of relationships) and maps (which depict information geospatially).

- Graphs are perceived almost entirely by the visual system. Lines, bars and other objects positioned in a 2-D space formed by perpendicular axes communicate visual patterns and relationships. Seeing pattern is a natural function of visual perception.
- Tables interact primarily with the verbal system. Information in a table is processed sequentially — reading down columns or across rows, comparing one pair of numbers at a time.
- Neither is inherently better; each is better than the other for particular communication tasks, and both play a vital role.

REMEMBER Examination point: "Graphs are visual, tables are verbal" is the single sentence that explains almost every table-versus-graph rule in Chapter 3.

Chapter 2 — Simple Statistics to Get You Started

Numbers are neither intrinsically boring nor interesting; impact flows from the significance and relevance of the message they contain. Because numbers are the substance of tables and graphs, we begin with the statistics that matter.

2.1 Quantitative Relationships

- Quantitative stories always feature relationships. A number is of no use unless it measures something important.
- Quantitative stories contain two types of data: quantitative values (which measure things, always expressed in a unit of measure — \$200 = the quantity 200 + the unit dollars) and categorical items (which divide information into useful groups and identify the things being measured).
- Relationships are of two kinds: simple associations between quantitative values and the categorical items that label them; and more complex associations among multiple sets of quantitative values (e.g. marketing e-mails sent vs. orders received).
- Do numbers always represent quantitative values? No. Order numbers, years and sequence numbers are categorical. Test: "Would it make sense to add these numbers up or perform any mathematical operation on them?" If no, they are categorical.

2.1.1 Relationships Within Categories (four types)

Type	Definition	Example
Nominal	Values in a single category are discrete and have no intrinsic order ("in name only").	Sales regions North, East, South, West.
Ordinal	Categorical items have a prescribed order; any other order (except reverse) is rarely meaningful.	First/second/third; small/medium/large; best salesperson downward.
Interval	Categorical items are a sequential series of numerical ranges that subdivide a larger	Order size: ≥0 & <1,000; ≥1,000 & <2,000; … Also months of the

Type	Definition	Example
	range of quantitative values, arranged smallest to largest.	year — time is a quantitative scale measuring duration, so months are treated as (approximately equal) intervals.
Hierarchical	Multiple categories associated as separate levels in parent-to-child connections; each item has exactly one parent but may have many children.	Division → Department → Group; Product Line → Product Family → Product.

REMEMBER Trap: months look ordinal but are formally an interval relationship (a time series). Even though months are not equally long, for reporting purposes we treat them as equal intervals.

2.1.2 Relationships Between Quantities (three types)

Type	Definition	Example
Ranking	The order in which categorical items are displayed is based on the associated quantitative values, ascending or descending.	Top five sales orders of the quarter by revenue.
Ratio	Compares two quantitative values by dividing one by the other, producing a number that expresses their relative quantities.	Part-to-whole market share as % of total; month-over-month change in expenses.
Correlation	Compares two paired sets of quantitative values to see whether increases in one correspond to increases or decreases in the other.	Employee tenure vs. productivity; marketing e-mails vs. orders.

Note: percentages are preferred for part-to-whole comparisons because they supply a common denominator and a common frame of reference — and not just any denominator, but the nice round 100, which makes comparison easy.

2.2 Numbers That Summarize (Aggregation)

Summarization (also called aggregation) reduces large sets of numbers to a few numbers so readers can comprehend and assimilate the story. Counts and sums are the most common aggregations; the four measures below are the less familiar but more instructive ones.

2.2.1 Measures of Average

> ***An average is a single number that represents the centre of an entire set of numbers.***

Average	How Calculated	Behaviour and Use
Mean (arithmetic mean)	Sum all values, divide by the count.	Takes every value into account, no matter how extreme — therefore very sensitive to outliers. Correct when total financial impact matters.
Median	Sort the values and take the middle one (for an even count, average the two middle values).	Not at all sensitive to extremes; expresses what is typical. It is the 50th percentile — a percentile expresses the percentage of values falling on or below a particular value.
Mode	The value that appears most often.	Rarely useful. A set with no repeats has no mode; two repeated values = bimodal; more than two = multimodal. The most common value need not be near the centre.
Midrange	Add the highest and lowest values and divide by two.	Fastest to calculate but far too sensitive to the two extremes unless values are uniformly distributed. Use only for a rough on-the-spot estimate.

The classic illustration: 13 marketing salaries containing a Vice-President at $475,000 (an outlier — a value located far from most of the values). The mean is $100,000 but the median is $62,000. To describe a typical salary, use the median. To describe the department's comparative impact on expenses, use the mean, because ignoring extremes would undervalue financial impact.

2.2.2 Measures of Variation

Two sets can share an identical average yet one be tightly grouped and the other volatile. In Few's worked example, Warehouse A and Warehouse B both ship 12 orders with sum 51, mean 4.25 and median 4.5 — yet A always ships in 3–5 days while B ranges from 1 to 10 days. Only variation reveals the difference.

Measure	How Calculated	Strength / Weakness
Spread (range)	Highest value minus lowest value. (A = 2 days, B = 9 days.)	Everyone understands it — its strength. But it relies on only two values and is heavily affected by extremes, so different distributions can share the same spread.
Standard deviation	(1) Compute the mean; (2) subtract each value from the mean; (3) square each difference; (4) sum them; (5) divide by the number of values — or by n − 1 for a sample; (6) take the square root.	Measures variation relative to the mean and uses every value. Warehouse B = 2.59 days vs. Warehouse A = 0.83 days. The bigger the standard deviation, the

Measure	How Calculated	Strength / Weakness
		greater the range of variation relative to the mean.

Normal distribution and the empirical rule. If measurements of almost any event are arranged from lowest to highest, most fall near the centre and fewer fall farther out. A frequency polygon of such data traces a bell-shaped curve called a normal distribution. For a normal distribution:

- 68% of values fall within ±1 standard deviation of the mean;
- 95% fall within ±2 standard deviations;
- 99.7% fall within ±3 standard deviations.

Business situations worth measuring for variation: selling price of products by region or representative; time to manufacture, answer calls or resolve problems; employee compensation for the same job across departments; cost of goods across vendors; departmental expenses.

2.2.3 Measures of Correlation

> ***The linear correlation coefficient measures the direction (positive or negative) and degree (strong or weak) of the linear relationship between two paired sets of quantitative values.***

It answers three questions: Does a correlation exist? Is it strong or weak? Is it positive or negative? Interpretation rules:

- All values fall between +1 and −1.
- 0 indicates no linear correlation; +1 a perfect positive linear correlation; −1 a perfect negative linear correlation.
- The greater the absolute value, the stronger the linear correlation.
- A variable is simply something that can take multiple values (e.g. employee productivity).

The scatter plot is the graph designed for this purpose: read what happens on the Y axis in relation to the X axis. Upward from left to right = positive; downward = negative. Tight grouping about the line = strong; bouncing about = weak. A straight line of best fit through the centre of the points is called a linear trend line or, formally, a regression line.

> **REMEMBER** Critical limitation: a linear correlation coefficient cannot describe a non-linear relationship — e.g. productivity that rises with tenure to a point and then falls. That is still a real relationship, and the remedy is managerial (new incentives or retraining at mid-tenure), not statistical.

2.2.4 Measures of Ratio

Whereas correlation measures the relationship between multiple paired sets of values, a ratio measures the relationship between a single pair of values. Ratios can be expressed in four ways:

1. As a sentence — "Two out of every five customers who access our website place an order."
2. As a fraction — 2/5.
3. As a rate — 0.4.

4. As a percentage — 40%.

Rates and percentages are the most concise and therefore the most useful in tables and graphs. Conventional forms exist: the book-to-bill rate is expressed as a rate (5 ÷ 4 = 1.25); profit margin is expressed as a percentage (25%). A powerful use of ratio is to set a reference value to 1 or 100% and express all other values relative to it — this removes the need for readers to do arithmetic in their heads.

2.2.5 Measures of Money

- Inflation — the value of money is not static. Sales of $100 million ten years ago and $200 million today is not a 100% real increase. Comparisons of monetary value across time are most accurate when adjusted for inflation.
- Multiple currencies — 100,000 U.S. dollars ≠ 100,000 British pounds ≠ 100,000 Japanese yen. Values must be converted to a single common currency before combining or comparing them, or results can be wrong by a large order of magnitude.

2.3 Summary at a Glance

Type of Summary	Method	Note
Average	Mean	Measures the centre in a manner equally sensitive to all values, including extremes.
	Median	Measures the centre in a manner insensitive to extreme values.
Variation	Spread	Simple, relies entirely on highest and lowest values, only roughly defines a range.
	Standard deviation	Rich expression of the distribution of values across the entire range.
Correlation	Linear correlation coefficient	Indicates whether a linear correlation exists and, if so, its direction and strength.
Ratio	Rate or percentage	Measures the direct relationship between two quantitative values.

Chapter 3 — Differing Roles of Tables and Graphs

Occasionally neither is right: when the message is only a number or two, a sentence or bullet does the job — "Q2 sales = $1,485,393, exceeding forecast by 16%." There is eloquence in simplicity.

3.1 Tables Defined

- Information is arranged in columns and rows.

- Information is encoded as text (words and numbers).
- It is the arrangement that characterises a table, not the presence of grid lines. A single set of values occupying one column or row is merely a list, not a table.
- Tables need not be quantitative — a meeting agenda (time, topic, facilitator) is a table. Tables have been in use for almost two millennia and are readily understood by almost anyone who can read.

3.2 When to Use Tables

> ***Tables make it easy to look up individual values.***

Use a table when one or more of the following is true:

5. The display will be used to look up individual values.
6. It will be used to compare individual values, but not entire series of values to one another.
7. Precise values are required.
8. The quantitative information involves more than one unit of measure.
9. Both summary and detail values are included.

Why these work: text is understood directly, without translation, so tables offer precision (27.387483) that visual encoding cannot; a graph normally carries a single quantitative scale in a single unit of measure, whereas a table can carry units, dollars and percentages side by side; and a table can hold monthly detail with an annual total in one display.

3.3 Graphs Defined

- Values are displayed within an area delineated by one or more axes.
- Values are encoded as visual objects positioned in relation to those axes.
- Axes provide scales — quantitative and categorical — used to label and assign values to the visual objects.
- Essentially, a graph is a visual display of quantitative information. The horizontal axis is the X axis, the vertical the Y axis.

3.4 A Brief History of Graphs

- The precursor of the graph is the map — the earliest known map is about 4,300 years old, drawn on a clay tablet, representing northern Mesopotamia. Longitude and latitude form a quantitative scale of sorts.
- René Descartes, La Géométrie (1637), adapted the grid to encode numbers as coordinate positions in a 2-D grid, founding an entire field of mathematics. Quantities such as displacement were graphed against time as far back as the 11th century (Massironi, via Ron Rensink).
- William Playfair, a Scottish social scientist of the late 18th century, invented the bar graph, was perhaps the first to use a line to show change through time — and, on an off day, also invented the pie chart.

- Edward Tufte, The Visual Display of Quantitative Information (1983), is the leading modern advocate of excellence in graphic design; William S. Cleveland, The Elements of Graphing Data, is the outstanding technical resource.

3.5 When to Use Graphs

> ***Graphs are used to display relationships among and between sets of quantitative values by giving them shape.***

The patterns a graph makes visible include:

- Trend — the overall nature of change during a period of time.
- Sub-patterns within the trend — e.g. the "hockey stick", where sales fall in the first month of each quarter and climb to a peak in the last, typically caused by quarterly sales bonuses.
- Comparison of entire sets of data (Q1 vs. Q2 vs. Q3), not merely pairs of values.
- Exceptions — the uncharacteristic August dip in international sales caused by customer holidays.
- Similarities and differences between two whole series, both overall and at particular points.

Graphs also handle very large data sets: thousands of records become instantly intelligible in a scatter plot.

3.5.1 The Tufte Disagreement

Tufte: "Tables usually outperform graphics in reporting on small sets of 20 numbers or less. The special power of graphics comes in the display of large data sets." Few disagrees slightly: even with only eight numbers, a graph of job satisfaction by income, education and age made the diverging pattern (older, high-earning employees without a college degree) obvious where the table hid it.

REMEMBER The essential qualification: displaying information in a graph rather than a table does not by itself make meaningful patterns visible. Figures 3.14 and 3.15 show the same eight numbers; only the first was designed to feature the pattern. We must design the graph to feature evidence of the particular story we are trying to tell.

3.6 Summary at a Glance

• The display will be used to look up individual values. • It will be used to compare individual values. • Precise values are required. • The quantitative values include more than one unit of measure. • Both detail and summary values are included.	• The message is contained in the shape of the values (patterns, trends and exceptions). • The display will be used to reveal relationships among whole sets of values.

Chapter 4 — Fundamental Variations of Tables

This chapter answers what a table can display, then how it can be structured, then links the what to the how.

4.1 The Five Relationships Tables Display

Primary Function	Relationship Type	Relationship	Example
Look-up	Quantitative-to-categorical	Between a single set of quantitative values and a single set of categorical items	QTD sales per salesperson — a simple list for look-up.
		Between a single set of quantitative values and the intersection of multiple categories	Sales by salesperson (category 1) and month (category 2) — a multi-dimensional view.
		Between a single set of quantitative values and the intersection of multiple hierarchical categories	Sales by product line → product family → product.
Comparison	Quantitative-to-quantitative	Among a single set of quantitative values associated with multiple categorical items	Comparing Mandy's February sales with Terri's — same layout as look-up, but the primary activity is comparison.
		Among distinct sets of quantitative values associated with a single categorical item	Sales, returns and net sales per salesperson.

4.2 Structural Variations

Structure	Definition	Notes
Unidirectional	Categorical items are laid out in one direction only — either across the columns or down the rows.	May contain more than one set of categorical items and still be unidirectional, provided all of them run the same way (e.g. Dept and Expense Type both down the rows).
Bidirectional	More than one set of categorical items, laid out both across the columns and down the rows. Also called a crosstab or pivot table.	Quantitative values sit in the body, bordered by categorical items across the top and along the left. To find a value, look where the relevant column and row intersect.

Efficiency argument: the same expense data occupies a 3 × 8 = 24-cell unidirectional grid but only a 4 × 5 = 20-cell bidirectional grid — and the bidirectional version additionally supplies a total for every expense type. It is smaller and contains more information.

4.3 Matching Relationship to Structure

Relationship	Unidirectional	Bidirectional
Single set of quantitative values ↔ single set of categorical items	Yes — the natural form.	No — a second category is required for a crosstab.
Single set of quantitative values ↔ intersection of multiple categories	Yes, but consumes more space and repeats labels.	Yes — the ideal form; most compact.
Single set of quantitative values ↔ intersection of hierarchical categories	Yes — hierarchy reads naturally left to right down the rows.	Yes, when one level of the hierarchy runs across the columns.
One set of quantitative values across multiple categorical items (comparison)	Yes.	Yes — often better, because values to be compared end up adjacent.
Distinct sets of quantitative values for a single categorical item	Yes — each measure gets its own column.	No — the measures are different units, not items of one category.

REMEMBER Principle to remember: the layout of a table is primarily determined by the relationship the table is meant to feature. Chapter 8 then refines that layout into a readable design.

Chapter 5 — Visual Perception and Graphical Communication

Vision is our most powerful and efficient channel for information. To design graphs that work, we must know how the eye and brain decode them. Do not skip this chapter — everything in Chapters 7 to 10 is an application of it.

5.1 Mechanics of Sight

- Sensation is a physical process that occurs when a stimulus is detected. Perception is a cognitive process — the interpretation of that stimulus to make sense of it.
- The chain is: Stimulus → Sensory organ (eye) → Perceptual organ (brain) → Iconic memory → Working memory → Long-term memory.

5.1.1 The Eye

Part	Function
Cornea	The protective covering on the surface of the eye.
Iris	A muscle working like a camera shutter, enlarging or reducing the opening.

Part	Function
Pupil	The opening itself — small in bright light, larger in dim light.
Lens	Thickens to focus on near objects, thins for distant objects.
Retina	The rear inner surface, coated with millions of nerve cells called rods and cones.
Rods and cones	Rods sense dim light and record in black and white; cones sense brighter light and record in colour. Cones are of three types, specialised for blue, green or red.
Fovea	A small central area of the retina where cones are densely concentrated — images focused here are seen with extreme clarity.
Optic nerve	Carries the neural signals to the brain.

Despite the millions of bits our eyes process every second, we see only a little of what is out there, and only a fraction of that in detail. During a saccadic eye movement we effectively see nothing.

5.2 The Brain — Three Memories

The computer analogy: keyboard buffer → RAM → hard disk corresponds to iconic memory → working memory → long-term memory.

Memory	Characteristics	Design Consequence
Iconic memory (visual sensory register)	Holds each snapshot for less than a second. Processing here is automatic and unconscious — this is preattentive processing: extremely fast recognition and nothing more.	If you want something to stand out, encode it with a preattentive attribute that contrasts with its surroundings (e.g. red text among black). If you want objects seen as a group, give them the same preattentive attribute.
Working memory	Temporary (seconds to a few hours with rehearsal) and severely limited: only three or four chunks at a time. Attentive processing happens here — conscious, sequential, slow, and capable of learning and understanding.	A legend with ten different colours or symbols forces constant back-and-forth reference. Design displays so that many data values combine into a few larger, coherent patterns (chunks).
Long-term memory	Stores chunks through an intricate neural network of associations; holds our ability to recognise images and detect meaningful patterns.	Little direct design impact; we mainly need preattentive attributes and respect for working-memory limits.

REMEMBER Preattentive vs. attentive: counting the 5s in four rows of grey digits takes seconds (attentive, serial). Print the 5s in black among light-grey digits and they appear instantly (preattentive, parallel). Visual perception evolved as an aid to survival and is fundamentally oriented toward action.

5.3 Attributes of Preattentive Processing

Colin Ware groups preattentive attributes into form, colour, spatial position and motion. Few drops motion and reduces the list to those relevant to static tables and graphs:

Category	Attributes	Notes
Form	Length, Width, Orientation, Shape, Size, Enclosure	Length = distance along an object's dominant dimension (the length of a bar). Width = variation along the secondary dimension.
Colour	Hue, Intensity	Hue is what we normally call colour. Intensity covers the other two attributes of colour — saturation and lightness.
Spatial position	2-D position	Our perception of space is primarily two-dimensional; we perceive vertical and horizontal position clearly and accurately, depth much less well.

5.3.1 The HSL Colour System

- Hue is described numerically in degrees from 0° to 360° around a colour wheel: red 0°/360°, yellow 60°, green 120°, cyan 180°, blue 240°, magenta 300°.
- Saturation measures the degree to which a hue fully exhibits its essence, expressed as 0%–100%.
- Lightness (brightness) measures how dark or bright a colour appears, expressed as 0% (fully dark) to 100% (fully light).
- Any colour can be described numerically by these three measures.

5.4 Which Attributes Encode Quantity?

Type	Attribute	Quantitatively Perceived?
Form	Length	Yes
	Width	Yes, but limited
	Orientation	No
	Size	Yes, but limited
	Shape	No

Type	Attribute	Quantitatively Perceived?
	Enclosure	No
Colour	Hue	No
	Intensity	Yes, but limited
Position	2-D position	Yes

All of these attributes can encode categorical differences; only length and 2-D position encode quantity accurately. **Rule: whenever you can use length or 2-D position to encode quantitative values, do so.**

- Why not area? Given a small circle worth 1, the large circle in Few's illustration is worth 16 — but in a class of 70 students answers typically range from 5 to 50. We can tell one area is bigger, not by how much.
- Why not colour intensity? We can tell one object is darker than another provided the difference is large enough, but we cannot assign a value to the difference.
- Why not hue? Different hues are perceived only as categorically different, never as more or less. (3-D position is preattentive but cannot be used on a flat surface, where it relies on illusory cues of light, shadow, occlusion and size.)

5.5 Perceptual Effects of Context

- Our visual receptors measure differences, not absolute values. A candle is dim in sunshine and brilliant in a pitch-black cellar.
- Four squares that are all exactly 50% grey appear to differ when placed against different surrounding intensities. Practical consequence: never use a gradient fill in a graph background, or identical bars on the left and right will appear different.
- The same effect applies to length — the two horizontal lines of the Müller-Lyer style illustration are identical but appear different because of the shapes at their ends. Even when we know it is an illusion, the misperception persists.
- Text/background combinations: black on white works well; white on black works well; red or dark blue on white works to a lesser degree; yellow on white and blue on black fail. Do not memorise combinations — look at them.

5.6 Limits to Distinct Perception

- Preattentive symbols become less distinct as the variety of distractors increases (Ware): it is easy to spot a hawk among pigeons, harder in a sky full of varied birds.
- Research limits: no more than about **eight distinguishable hues**, about **four orientations** and about **four sizes**. Other attributes should also be limited to a few values.
- Using more forces slow attentive processing — and then readers hit the three-to-four-chunk limit of working memory anyway.

- Preattentive processing usually cannot handle more than one visual attribute at a time. Picking out "all black circles" is easy; picking out "dark grey squares" when both shape and intensity vary requires attentive processing.
- Chosen values must be far enough apart along the continuum to be distinct — four carefully chosen greys work; four adjacent greys do not.

Few's nine distinct hues, in order: **gray, blue, orange, green, pink, brown, purple, yellow, red.**

- Cultural context matters: in most Western cultures red signifies danger or heat; in China red represents good fortune.
- Tufte: colours found predominantly in nature and not overly vibrant (medium grey, green, blue, brown) are soothing and easy on the eyes — ideal for anything that should not stand out. Fully saturated, bright hues demand attention and should be reserved for highlighting.
- Our ability to distinguish colours decreases with the size of objects. Maintain **three palettes**: light-but-distinct colours for large objects such as bars; darker versions for small objects such as lines and points; bright versions for objects that must stand out.
- Colour-blindness: roughly the commonest form prevents distinguishing red from green (both appear brown). Use red or green, not both, unless their intensities differ sufficiently.
- A good source of data-presentation palettes is ColorBrewer (www.ColorBrewer.org).

5.7 Limits to the Use of Contrast

> ***Contrast is most effective when only one thing is different in a context of other things that are the same. When everyone in the room shouts, no one can be heard.***

Our brains insist on making sense of any difference we see, so only meaningful differences should be included. In the "Funding by Source" example, a single bar in a different colour dominates attention precisely because nothing else differs.

5.8 Gestalt Principles of Visual Perception

The Gestalt School of Psychology began in 1912 to uncover how we perceive pattern, form and organisation. Gestalt is simply the German word for "pattern". These principles describe how we group what we see.

Principle	Statement	Application to Tables and Graphs
Proximity	We perceive objects that are close to one another as belonging to a group.	Space data slightly closer horizontally to make readers scan by row; closer vertically to make them scan by column. White space is the primary structuring device of a table.
Similarity	We group objects that are similar in colour, size, shape or orientation.	Use hue, intensity, size, shape or orientation to encode categorical groupings — kept to a few, clearly distinct values. Can also drive row-wise or column-wise scanning in a table.

Principle	Statement	Application to Tables and Graphs
Enclosure	Objects enclosed by a border or a common field of colour are perceived as belonging together.	Borders and background fills group or emphasise regions. A strong enclosure is unnecessary — subtle works. We even perceive illusory enclosures that are not really there.
Closure	We perceive open structures as closed, complete and regular whenever we can reasonably do so; we dislike loose ends.	A graph's data region can be defined by a single subtle pair of X and Y axes rather than a heavy full border with fill — minimising the visual weight of supporting components.
Continuity	We perceive objects as belonging together if they are aligned or appear to continue one another.	Left-aligned bars share an obvious common baseline, so a vertical axis line becomes unnecessary. In tables, distinct alignments distinguish division names from department names and remove the need for vertical grid lines — even when columns touch.
Connection	We perceive objects connected by a line as part of the same group. (An extension of continuity rather than an original Gestalt principle.)	Lines are the strongest way to connect points in a graph; points without lines are much harder for the eye to connect, and lines additionally reveal the overall shape of the data.

REMEMBER Relative perceptual strength: **enclosure > connection > proximity and similarity.** Choose the weakest principle that does the job, so grouping happens without visual clutter.

Chapter 6 — Fundamental Variations of Graphs

6.1 Graphical Means for Encoding Quantitative Values

Object	Description	Strength / Verdict
Points	Any simple, small geometric shape marking a location. Used in scatter plots and, against a categorical scale, in dot plots.	Encode by 2-D position. Do not require a zero-based scale, so they suit subtle differences among large values. KEEP.
Lines	Connect a series of values, or summarise a set as a trend line / line of best fit (which may be a curve if the pattern is curvilinear).	Let us view an entire series as a single pattern from left to right — ideal for change through time. KEEP.

Object	Description	Strength / Verdict
Bars	A line given width, becoming a rectangle. Width is meaningless; value is encoded by length and by the end position against the scale. Vertical bars are sometimes called columns.	Great visual weight makes individual values stand out distinctly. MUST begin at zero. KEEP.
Boxes	A rectangle encoding an entire range of values, with a mark for the centre (usually the median).	The one way to compare many distributions at once. KEEP.
Shapes with 2-D areas	Pie charts and the wider family of area graphs; also bubbles (points of varying size) in bubble plots.	Visual perception cannot accurately assign values to 2-D areas. AVOID — except bubbles when only approximate values are needed and the third variable earns its complexity.
Varying colour intensities	Light-to-dark encoding of a value, e.g. bubble colour for patient age, or choropleth shading.	Perceived quantitatively only to a limited degree. AVOID unless length and 2-D position are unavailable (e.g. on maps).

The pie chart verdict. A pie does have an axis — the circle's perimeter — but it is never labelled because doing so makes the chart unreadable. Pie charts encode values three ways (2-D area of slices, angles at the centre, arc lengths at the edge) and none of the three is easy to decode or compare. Ordering six similar slices by size is nearly impossible; the same values as bars are trivially ordered. Few does not use pie charts and recommend abandoning them.

REMEMBER With rare exceptions, rely on just four objects — **points, lines, bars and boxes** — because they exploit the only two attributes we perceive accurately: length and 2-D position.

6.2 Visual Attributes That Encode Categorical Items

Attribute	Applies To	Assessment
2-D position	Points, lines, bars, boxes	The most common and clearest means; no ambiguity about which month a bar belongs to.
Hue	Points, lines, bars, boxes	Second most effective. Requires a legend, but distinct hues let the reader attend to one series at a time. Fails if objects are very small.
Point shape	Points only	Dots, squares, triangles, asterisks, diamonds. Distinguishable, but not as fast as colour. Useful when position and colour are already taken.

Attribute	Applies To	Assessment
Fill pattern	Bars and boxes only	Last resort. Harder to distinguish than colour and can cause moiré vibration — a dizzying effect, worst with vivid lines running in different directions. Use only when colour printing is impossible; distinct greys usually work better.
Line style	Lines only	Solid, dashed, dotted. Breaks in a line undermine perception of a continuous series. Reserve for non-colour printing; varying colour intensity or thickness usually works better.

6.3 The Eight Relationships Graphs Display

Relationship	Definition	Trigger Words in Your Message
Time series	A series of values featuring how something changed with time.	change, rise, increase, fluctuate, grow, decline, decrease, trend
Ranking	How values relate sequentially, sorted by size low-to-high or high-to-low.	larger than, smaller than, equal to, greater than, less than
Part-to-whole	How individual values that make up a whole compare to one another and to the whole, as ratios summing to 100% or to 1.	rate or rate of total, percent of total, share, accounts for X percent
Deviation	How one or more sets of values differ from a reference set, expressed directly as the difference.	plus or minus, variance, difference, relative to
Distribution	How values are distributed across their full range from lowest to highest. For a single set this is a frequency distribution.	frequency, distribution, range, concentration, normal curve / bell curve
Correlation	Whether two paired sets of values vary together, in which direction and to what degree. Correlation may indicate causation but need not — both may be driven by external factors.	increases with, decreases with, changes with, varies with, caused by, affected by, follows
Geospatial	Where values are located, usually on a map (non-geographic spatial displays such as a building floor plan also exist).	geography, location, where, region, territory, country, state, city
Nominal comparison	The simplest: a set of discrete values in no particular order, merely to be read and compared.	"this is bigger than that", "these two are far bigger than all the others"

Five of these types come from Gene Zelazny's taxonomy in Say It with Charts (component, item, time-series, frequency-distribution and correlation comparison); Few added nominal comparison, deviation and geospatial.

6.4 Graph Design Solutions — Relationship by Relationship

6.4.1 Nominal Comparison

- Values are discrete and unconnected, so encode them in a way that emphasises distinctness — **bars**, vertical or horizontal.
- When all bars are long and similar in length, the differences are hard to see; bars cannot fix this because they require a zero base. Replace them with **points (a dot plot)**, which can use a narrowed scale.
- Usable objects: vertical bars, horizontal bars, points.

6.4.2 Time Series

- Time must run **horizontally, left to right, along the X axis**; the quantitative scale goes on the Y axis. This immediately eliminates horizontal bars and horizontal boxes.
- Vertical bars work only when you wish to emphasise and compare individual points in time; their visual weight distracts from the overall shape.
- Points alone do not convey the sequential nature of time — floating dots give no sense of continuity. Adding a line makes the flow of time visible.
- Lines with or without points are the standard choice; points are needed only when exact positions matter.
- Exception: when values were collected at **irregular intervals**, omit the line and show points only (a dot plot). Connecting irregular observations implies a pattern of change that may be quite unlike the true daily pattern.
- Usable objects: vertical bars (individual emphasis), lines (pattern of change), points only (irregular intervals), vertical box plots (distributions changing through time).

6.4.3 Ranking

- Bars reinforce the individuality of values and their relative sizes; sorting them by size makes the ranking message crystal clear.
- Rule of thumb — to highlight the highest values: sort descending; vertical bars with the highest on the left, horizontal bars with the highest on top. To highlight the lowest values: sort ascending; lowest bar on the left or on top.
- The convention rests on the fact that most cultures treat top and left as the beginning, probably because of left-to-right, top-to-bottom reading.
- Usable objects: bars (vertical or horizontal, only with a zero-based scale); points (especially when the scale must not begin at zero).

6.4.4 Part-to-Whole

- Percentage is normally the best unit — the whole equals 100%.

- Pie charts: rejected. Stacked bars: only slightly better than a pie. In a single stacked bar only the first segment can be read directly off the scale; every other segment requires mental subtraction, and adjacent segments of similar size cannot be compared.
- The fix is to **unstack** the bars into individual bars, and then **sort by size** so that similar values sit next to one another. This combines a ranking and a part-to-whole relationship in one graph.
- Because bars serve many purposes, make the part-to-whole intent explicit: give the graph a title stating that the bars are parts of a whole; make the bars **touch** (people notice the missing gaps); and print the values as text, including "100%" as the total.
- Justified use of stacked bars: when the whole must be expressed in a unit other than percentage (e.g. U.S. dollars by channel) while still giving an approximate sense of the parts (regions).
- Usable objects: bars (vertical or horizontal); stacked bars (only for the case above).

6.4.5 Deviation

- Show the difference **directly** rather than forcing the reader to subtract mentally. A graph of actual vs. plan headcount is inferior to a graph of variance from plan.
- Deviations may be expressed as raw amounts or as percentages — both are valid but tell different stories. Raw numbers show overall organisational impact; percentages compare departments' relative performance.
- Highlight the unfavourable direction (e.g. positive headcount variance in red) so it draws attention.
- Deviation is frequently combined with time series; then encode with a **line** to represent the continuous, flowing nature of time (e.g. each month's sales as a percentage of January's, or percentage change from the prior month).
- Usable objects: bars (vertical or horizontal); lines (when the deviation runs through time).

6.4.6 Distribution

Display	Object	When to Use
Histogram	Bars	A single distribution where the emphasis is on the number or percentage of occurrences in each interval. The most common distribution display. Bars conventionally have no spaces between them, which suggests the continuous rather than discrete nature of the scale.
Frequency polygon	Line	A single distribution where the emphasis is on the overall shape. It also works superbly for cumulative distributions. Slight disadvantage: harder than a histogram for comparing one interval with another.
Strip plot	Points	Shows every individual value rather than aggregating into intervals. Best for small data sets where each value's exact position matters. Overlapping points are solved by jittering

Display	Object	When to Use
		(repositioning coincident points slightly) or, for large sets, by transparency, so denser colour reveals concentration.
Box plot (box-and-whisker plot)	Boxes	Comparing many distributions, or showing how a distribution changes through time. Invented by John Tukey of Princeton.

Anatomy of a five-value box plot — it communicates seven facts: the highest value; the lowest value; the spread (100% of the values, high minus low); the median (50th percentile), marked inside the box; the midspread (the middle 50% of the values); the 75th percentile (the value at or above which the highest 25% reside); and the 25th percentile (the value at or below which the lowest 25% reside).

- Reading practice: if the median sits nearer the bottom of the range, more people fall toward the lower end; a long upper whisker means values thin out as they rise.
- Most readers have never learned to read box plots. A simplified **three-value box plot** (low, median, high) needs only a small legend and a few words of explanation.
- Box order matters. Reordering male and female boxes changes whether a line drawn through the medians is smooth or jagged; jaggedness is jarring and increases the salience of the discrepancy. Both versions tell the truth — one brings this aspect of the truth to attention more forcefully.
- Boxes may be horizontal as well as vertical; strip plots may show several rows or columns to compare multiple distributions.
- Multiple distributions: lines (frequency polygons, limited to a few lines), boxes, points (strip plots). When the number of values per group differs greatly, plot percentages rather than counts so the shapes are comparable.

6.4.7 Correlation

- Both axes carry quantitative scales, so the natural display is the **scatter plot** — in effect two frequency distributions at once, one from horizontal positions and one from vertical positions.
- Add a **trend line** to make the direction and strength of the correlation visible.
- Alternative for audiences who cannot read scatter plots: the **table lens** (invented by Ramana Rao and Stuart Card) — two bar graphs side by side, one per variable, with the rows sorted by the first variable. If the second set of bars rises or falls consistently with the first, a correlation exists.
- Trade-off: scatter plots are superior overall and can hold thousands of points; a table lens is more intuitive and can show more than two variables at once, but only a limited number of bars fit in a given space.
- Usable objects: points (scatter plot); bars (table lens).

6.4.8 Geospatial

- On a map, our two best encodings are unavailable: 2-D position is already used for geographical location, and objects cannot share a common baseline. We must fall back on size, colour intensity and width.
- Points of varying size — to pinpoint specific locations (e.g. 911 calls in a section of Seattle) or to aggregate whole regions. When only one variable is shown, size is usually better than intensity because colour discrimination weakens as objects get small.
- Points or areas of varying colour intensity. Applying intensity directly to geographical regions produces a **choropleth map**.
- Lines of varying thickness or colour intensity — for values associated with routes (traffic on roads, hurricane tracks).
- Combining all three lets a single map carry three quantitative variables (e.g. circle size = sales, circle intensity = profit, state colour = population).

6.5 Summary at a Glance — Relationship × Object

Relationship	Points	Lines	Bars	Boxes
Nominal comparison	Yes — dot plot, when the scale cannot begin at zero	Avoid	Yes — horizontal or vertical	Avoid
Time series	Only when intervals of time are irregular	Yes — emphasis on overall pattern	Vertical only — emphasis on individual values	Only when showing distributions changing through time; vertical
Ranking	Yes — dot plot, especially when the scale cannot begin at zero	Avoid	Yes — horizontal or vertical, sorted by size	Only when ranking multiple distributions
Part-to-whole	Avoid	Only to show how parts of a whole changed through time	Yes — horizontal or vertical (unstacked and sorted)	Avoid
Deviation	As a dot plot when the scale cannot begin at zero	Useful when combined with a time series	Yes — always vertical when combined with time series	Avoid
Distribution (single)	Strip plot — emphasis on individual values	Frequency polygon — emphasis on overall pattern	Histogram — emphasis on individual intervals	Avoid

Relationship	Points	Lines	Bars	Boxes
Distribution (multiple)	Strip plot	Frequency polygon — limit to a few lines	Avoid	Box plot
Correlation	Scatter plot	Avoid (except as a trend line)	Table lens — horizontal or vertical	Avoid
Geospatial	Vary point sizes to encode values	To mark routes	Avoid	Avoid

Chapter 7 — General Design for Communication

These practices apply equally to tables, graphs and text. There are two fundamental objectives: **Highlight** and **Organize**. We highlight important information to give it a clear voice without distraction; we organize information to lead readers through it in a manner that promotes understanding and use.

7.1 Highlight

> ***Above all else show the data. — Edward Tufte***

Data-ink ratio (Tufte): the amount of ink that presents information compared to the total amount of ink. The object is not to eliminate all non-data ink — some supporting components are always needed — but to reduce it to no more than what is necessary to make the data ink understandable.

7.1.1 Reduce the Non-Data Ink

10. **Subtract unnecessary non-data ink.** Ask of every component: "Would the data suffer any loss of meaning or impact if this were eliminated?" If the answer is no, remove it. Resist keeping things because they are cute or because you worked hard on them.
11. **De-emphasize and regularize the remainder.** Push what remains into the background: thin lines, soft neutral colours such as light grey. Make every repeated supporting component look precisely the same everywhere, because the eye is drawn to contrast and the brain assigns meaning to any difference.

Displays consist of **three visual layers**: (1) data as the prominent top layer, (2) non-data items as the middle layer, and (3) the background. When everything stands out, nothing stands out.

Saint-Exupéry, quoted by Few: "In anything at all, perfection is finally attained not when there is no longer anything to add, but when there is no longer anything to take away." Elegance comes from the Latin *eligere* — to choose out, to select carefully.

7.1.2 Enhance the Data Ink

12. **Subtract unnecessary data ink.** Not all information is equally important. Summarise when detail is unnecessary; remove what is peripheral to the reader's purpose. The more you earn readers' trust by giving them only what they need, the more attention they will pay to everything you give them.

13. **Emphasize the most important data ink** using preattentive attributes.

Attribute	Values Useful for Emphasis
Width	Thicker lines — including boldfaced words and numbers — stand out more than thinner ones.
Orientation	Slanted (italic) text stands out, assuming upright type is the norm.
Size	Bigger objects, words and numbers stand out more than smaller ones.
Enclosure	Items enclosed by lines or background fill stand out more than those that are not.
Hue	A hue different from the norm stands out.
Colour intensity	Dark or bright items stand out more than light or pale ones.

REMEMBER Each step of highlighting produces greater simplicity. Your message may be complex, but its design should be so simple that to your readers it is nearly invisible.

7.2 Organize

Readers organise what they see whether you help them or not. Fail to organise deliberately and the result is not unorganised information but information organised in a way that does not support its story — possibly conveying the wrong message. Three sub-objectives:

7.2.1 Group

- Grouping happens on several levels: overall message → topics → modes of expression (tables, graphs, text) → categorical vs. quantitative data → individual categorical items and their values.
- The simplest technique — **proximity** — is often the best: place related content together and separate topic groups with white space.
- When related topics must also be tied across separate displays (e.g. bookings vs. billings appearing in one graph and two tables), select another Gestalt principle or preattentive attribute. Choose carefully: colour intensity would imply that one is more important, so if they are equal in importance use different **hues** instead.
- Conventions: tables use proximity and continuity to organise columns and rows; graphs use similarity (common hues or shapes) and connection (a line joining points).

7.2.2 Prioritize

Ask every time: "What are the important numbers here?" Some numbers are so much more important than others that a few seconds spent on them yields benefits that years spent on the rest could never equal.

Quantitatively Perceived Attribute	In Tables and Text	In Graphs and Objects
Width	Boldfaced text	Thicker graph lines; wider bars
Size	Bigger tables; larger fonts	Bigger graphs; bigger symbol shapes
Colour intensity	Darker or brighter colours	Darker or brighter colours
2-D position	Positioned at the top, at the left, or in the centre	Positioned at the top, at the left, or in the centre
Contrast-Based Attribute	**In Tables and Text**	**In Graphs and Objects**
Orientation	Italics	Data points oriented differently from the norm
Shape	Any font different from the norm	Any symbol shape different from the norm
Enclosure	Border or fill colour behind tables, rows, columns or particular values	Border or fill colour behind graphs or particular values
Hue	Almost any hue different from the norm	Almost any hue different from the norm
2-D position	Any position out of vertical or horizontal alignment with the norm	Any position out of vertical or horizontal alignment with the norm

- We are extremely sensitive to alignment — the slightest misalignment jumps out and we try to impose meaning on it. Keep edges aligned unless you intend something to stand out.
- Indentation may be used deliberately to establish hierarchy, but indent far enough to make the intention clear.
- **Pointers** — arrows, asterisks, check marks — reliably draw attention but are not subtle. Use with discretion to avoid clutter.

7.2.3 Sequence

- The strongest sequencing method is location in 2-D space: readers scan left to right and top to bottom, column by column.
- This is strongest for text, which can only be read sequentially; it applies to graphs too, but any graph highlighted as important will pull the eye to itself first.

- When a particular order is essential, reinforce it with numbers (1, 2, 3), letters (A, B, C) or other sequential labelling.

7.3 Integrating Tables, Graphs and Text

Role of Text	What It Does
Label	Titles, column and row headers, axis labels, legends. Clear titles are vital data in themselves — people often choose a report from its title alone.
Introduce	Previews what readers will find, what they should especially notice, and what they should do with the information. It is your chance to say in writing what you would say in person.
Explain	Placed exactly where needed — e.g. a note under a time-series graph explaining that an April promotional campaign caused the May up-tick. Improve the design first; add text only if the design cannot solve it.
Reinforce	Repeat genuinely important information in a second form to increase the odds it gets through. Do not overdo it, or you waste readers' time and lose their confidence.
Highlight	Call attention to something in words without repeating it — useful when visually emphasising it would unfairly diminish other data.
Sequence	Instruct readers to navigate the contents in a particular order when visual methods alone cannot.
Recommend	Recommendations for action are best communicated in words.
Inquire	Vastly underrated. You can sometimes add more value by asking one important question than by providing hundreds of answers.

Text placement. Blend tables, graphs and text; place each unit exactly where it is needed. A legend placed far from the data it labels forces constant back-and-forth because the encodings cannot be held in working memory. If a long discussion spans several pages, reproduce the table or graph wherever it is needed. The tighter the integration, the better.

7.3.1 Required Text on Every Page

Question	What to Provide
What?	A good title describing the quantitative and categorical content without being long-winded. "Sales" is not enough; "2011 Bookings by Month and Region" says a great deal more.
When?	Two facts: the range of dates the information represents, and the point in time when the information was collected — the "as of" date. Reports covering the same period often differ simply because they were produced at different times.
Who?	Your name or your group's name plus a means of contact — not self-promotion, but so readers know whom to ask.

Question	What to Provide
Where?	Page numbers, best in the form "Page 13 of 197", which tells readers from the first page how long the report is.

REMEMBER Excerpts of multi-page reports get copied and circulated. If identifying information appears only at the beginning, a reader holding page 7 has no idea where it came from.

Chapter 8 — Table Design

Tables have existed since the 2nd century CE and became commonplace with spreadsheet software in the 1970s. Used properly and designed well they are fantastic; the practices needed to optimise them are simple and easy to learn.

8.1 Structural Components of a Table

Group	Components	Notes
Data (information) components	Categorical labels; quantitative values; complementary text	Complementary text plays the eight roles of Chapter 7 (label, introduce, explain, reinforce, highlight, sequence, recommend, inquire). The commonest uses are titles and headers.
Support components	White space and page breaks; rules and grids; fill colour	White space is the primary structuring device and rests on the Gestalt principle of proximity. Grids work on enclosure; rules work primarily on connection. Fill colour uses similarity and enclosure. The background is the surface on which all of it sits.

8.2 Table Terminology (examinable)

Term	Meaning
Body	The rectangular area containing the quantitative values only — a definition that highlights the centrality of the numbers.
Cell	The intersection of a column and a row.
Column headers	Labels appearing above the column of data they describe.
Row headers	Categorical labels appearing to the left of the values associated with them.
Spanner header	A column header that spans multiple columns.

Term	Meaning
Spanner rule	A rule underscoring a spanner header, spanning each column to which the header applies.
Rule	A line running only horizontally or only vertically — rules do not intersect.
Grid	Intersecting horizontal and vertical lines forming a matrix of rectangles around the data.
Footer	A row summarising the preceding rows. A group footer summarises a subset of rows.
Border	The outer boundary of the table.

8.3 Delineating Columns and Rows

> ***Use white space alone whenever space allows. When you cannot use white space, use subtle fill colours. When you cannot fill colour, use subtle rules. Avoid grids altogether.***

- White space is the least visible means available — it does its job without calling attention to itself.
- As rows grow wider, more vertical white space is needed to track across them. Two remedies: decrease horizontal white space between columns, or increase vertical white space between rows.
- **The practical limit is a 1:1 ratio** of data height to white-space height. In Few's example a 10-point font in a 12-point row (2 points of white space, 20% of the data height) is difficult to track; white space equal to the row height tracks easily; white space at 150% of data height wastes space and destroys data density.
- White space also steers scanning direction: more space between columns than rows makes readers scan down columns; the reverse makes them scan across rows.
- The problem with rules and grids is that they excessively **fragment the data and introduce clutter** — the eye perceives individual cells (Gestalt enclosure) instead of a seamless flow of information. Progressively removing grid lines makes tables progressively easier to read.
- Fill colour, when used, must be subtle. Its purpose is to group or highlight, not to decorate.

8.4 Arranging Data

8.4.1 Columns or Rows

- Arrange a set of categorical subdivisions **across separate columns** if they are few in number and the maximum number of characters is not too large.
- Arrange **time-series subdivisions horizontally** across separate columns (matching the left-to-right convention for time).
- Arrange **ranked subdivisions vertically** down the rows.

8.4.2 Groups and Breaks

- Use just enough vertical white space between groups to make the breaks noticeable.
- Repeat column headers at the beginning of each new group.
- Keep table structure consistent from group to group.
- When groups should be examined independently, start each on a new page.

8.4.3 Column Sequence

- Place sets of categorical subdivisions that run down the rows of a single column **to the left** of the quantitative values associated with them.
- Place hierarchically related categorical subdivisions **left to right to reflect the hierarchy** (Division, then Department, then Group).
- Place quantitative values that were calculated from another set of values **just to the right** of the column from which they were derived (e.g. Variance immediately right of Actual).
- Place columns containing data that should be compared **close together**.

8.4.4 Data (Value) Sequence

Whenever categorical subdivisions have a meaningful order — chronological, ordinal, or by rank — sort them in that order.

8.5 Formatting Text

Topic	Best Practice
Orientation	Avoid text orientations other than horizontal, left to right. Rotated or vertical headers slow reading dramatically.
Alignment	Align numbers to the right, keeping decimal points aligned. Align dates however you wish, but keep a consistent number of characters or digits for each part. Align all other text to the left. Centre non-numeric items only if they all have the same number of characters and the header is significantly wider.
Number formatting	Place a comma to the left of every three whole-number digits. Truncate whole numbers by sets of three digits when precision can be reduced to the nearest thousand, million or billion. When negative numbers are enclosed in parentheses, keep the negative numbers themselves right-aligned with the positive numbers.
Date formatting	Express months either as a two-digit number or a three-character word; days as two digits; years consistently as either two or four digits.
Number and date precision	Do not exceed the required level of precision. Extra decimal places are noise.
Font	Select a font that is legible and use the same font throughout the table. Serif or sans-serif both work if legible (Times New Roman, Palatino, Arial, Verdana); decorative faces (Academy Engraved, Poplar Std Black, Script,

Topic	Best Practice
	Synchro LET) do not. The best typefaces for tables vary the width of text characters but keep numerals equal in width, so digits line up in columns.
Emphasis and colour	Boldface, italicise or change font colour when useful to group or highlight — not for decoration.

8.6 Summarizing Values

- **Column and row summaries** frequently convey critical information. Whenever you design a table, ask whether they would be meaningful even if not requested; if the answer is yes or even maybe, and there is space, include them. Being at the edges, they are easy for readers to ignore.
- Summaries of values across multiple columns in a single row go in a column of their own, generally to the right of the values they summarise.
- **Group summaries** aggregate meaningful subsets rather than the whole column or row.
- **Make columns containing group summaries visually distinct from detail columns.** Otherwise readers mistake a "Q1 Total" or "Q1 Mo Avg" column for just another monthly value — an error that grows as the eye moves further from the headers. The minimum distinction (extra white space) is usually enough; keep it subtle unless the summaries genuinely deserve more attention.
- **Headers versus footers.** Summary rows are conventionally placed in footers, but when the summaries matter more than the details — and especially when the details run over many pages — place them in the group header instead, where they are found immediately.

8.7 Giving Page Information

- **Repeat column headers at the top of each page.**
- **Repeat the current row header at the top of each page.** In a grouped table, the group label (e.g. the state name) is printed only on the first row of the group — repeating it on every row would waste the reader's attention and destroy the visual cue that a new group has started. But by page three of Alabama's data the reader has forgotten which state it is. Repeating the current row header at the top of each page costs nothing and benefits readers a great deal.

Chapter 9 — General Graph Design

Tables are somewhat forgiving of design flaws because they encode data verbally. Graphs encode data as visual objects, so those objects must be prominent, accurate and clear. Two principles apply **exclusively** to graphs.

9.1 Principle 1 — Maintain Visual Correspondence to Quantity

Only **length** and **2-D position** can be interpreted easily and accurately. A bar twice as long as another is perceived as twice the value; if the shorter bar actually represents 1.75 or 0.5 rather than 1, something is amiss. Darrell Huff's How to Lie with Statistics (1954) was among the earliest warnings about deliberate manipulation.

> ***The goal of design for communication is always to promote an accurate understanding of the truth.***

9.1.1 The "Sales Are Skyrocketing!" Case Study

A rise from $19,500,000 in July to $19,560,000 in December — less than one-third of one percent — was made to look like more than 200% growth. Five devices did it:

14. The Y-axis scale does not start at zero; it runs only from $19,475,000 to $19,560,000, so minor changes appear extreme.
15. The plot area is taller than it is wide, dramatically increasing the slope of the line.
16. The line is green — the colour of growth, health and dollars in English-speaking cultures — and bright green on black pops with visual impact.
17. The highest value ($19,560,001) is set as the top of the scale, so the line appears to run off the top of the graph.
18. The boldfaced axis label "Millions" sits in the prominent upper-left position beside the title, suggesting increases of millions.

An even worse version removes all Y-axis values and labels only the final data point, so there is no way at all to know the extent of the increase. Making it 3-D and manipulating the angle would exaggerate it further.

9.1.2 Correspondence to Tick Marks

- Always keep the distance between tick marks consistent with the difference in the values they represent. Equal distances must mean equal numeric intervals.
- Never place a gap in the values (1, 2, 7, 8, 9) even if no data falls in the missing range — this destroys the graph's visual integrity.
- **Scale breaks mislead even when marked.** With a break between 400 and 8,000, a dramatic rise in influenza cases from 2010 to 2011 is hugely under-represented, despite the scale starting at zero.
- When one graph cannot serve both the overall story and the fine detail, **use two graphs** — a zero-based graph showing the dramatic rise, plus a second, narrower-scale graph showing the earlier years in detail. Nothing is gained by squeezing into a single graph what several can present more effectively.
- (Logarithmic scales are the recognised exception where tick marks are not equidistant in value — see Chapter 10.)

9.1.3 Zero-Based Scales

- Setting the bottom of the scale above zero exaggerates differences visually. Usually avoid it.
- It is legitimate when you need a close look at small differences among large values — but then **alert your readers**, e.g. with a red note: "Attention: the dollar scale along the vertical axis has been narrowed to reveal the small but steady rise in sales since July."

- **Never eliminate zero when bars encode the values**, because a bar encodes value primarily through length and without a zero base the length no longer corresponds to the value. The MicroStrategy advertisement made an 83%-versus-63% difference in customer loyalty look like a six-fold difference by starting the scale at 60%.
- With positive and negative values, zero will not be the bottom of the scale but must still be the base from which bars extend up and down. Setting the base slightly below the lowest value is confusing and misleading.

9.2 Principle 2 — Avoid 3-D

Form of 3-D	Why It Fails
Depth added to data objects (bars, points, lines) with no third quantitative scale	Depth does not change the value. Bars merely occupy more space and become harder to tie to the scale; points become spheres or cubes that are harder to align; lines become thick ribbons with the same problem. It adds ink but no data — meaningless content the reader must process, wasting time and effort. Manipulating perspective and angle can make a steady rise from $100,000 to $121,000 look flat.
Depth added to the whole graph via a third (Z) axis, carrying a categorical or quantitative scale	Software applies an axonometric projection — tilting, rotating and adding perspective (e.g. tilted down 15°, rotated 20° clockwise, 30° of perspective). Some bars are always partially or entirely hidden, and it is nearly impossible to line objects up with the quantitative scale. Rotating to reveal hidden bars simply hides others, defeating the whole benefit of a graph. Grid lines on the walls, borders around bars and drop lines do not fix it — drop lines only add clutter and reduce reading to a slow series of look-ups. With lines instead of bars, fewer than half of Few's students could identify which line was the south region.

REMEMBER When you sincerely want to give readers a third variable, use **multiple related graphs in a series** (small multiples) rather than a third dimension.

9.3 Summary at a Glance

- Encode quantities to correspond accurately to the visual scale.
- Keep the distance between tick marks consistent with the difference in the values they represent.
- In most cases include zero in your quantitative scale, and alert readers when you do not. Always start the quantitative scale at zero when bars encode the values.
- Avoid 3-D displays of quantitative data.

Chapter 10 — Component-Level Graph Design

Graphs are built from components. The chapter is organised into primary data components, secondary data components and non-data components.

10.1 Anatomy of a Graph

Terms used throughout: Title · Subtitle · X axis and Y axis · X-axis title and Y-axis title · X-axis labels and Y-axis labels · Tick mark · Grid line · Legend · Data region · Graph border · Primary data object (point, bar, line or box) · Reference line · Trend line.

10.2 Primary Data Components

10.2.1 Points

The single objective is to keep points easy to perceive.

- When sets of points cannot be clearly distinguished: enlarge the points, or select objects that are more visually distinct.
- When points overlap so that some are obscured: enlarge the graph and/or reduce the size of the points, or remove the fill colours (hollow points).

10.2.2 Bars

Aspect	Practice
Orientation	Use **horizontal** bars when the categorical labels will not fit side by side beneath vertical bars. **Never** use horizontal bars for time-series values.
Proximity	Set the width of the white space separating bars that are labelled along the axis **equal to the width of the bars, plus or minus 50%**. Do **not** include white space between bars that are differentiated by a legend (they form one group). Do **not** overlap bars.
Fills	Avoid fill patterns (horizontal, vertical or diagonal lines). Use fill colours that are clearly distinct, and fairly balanced in intensity for data sets of equal importance. Use a more intense fill to highlight particular values.
Borders	Place borders around bars only when (a) the fill colour is not distinct against the background — then use a subtle grey border — or (b) you wish to highlight particular bars.
Base value	Always start bars at a baseline of zero.

10.2.3 Lines

- Distinguish lines using different **hues** whenever possible (rather than line styles).
- Include **points on lines only** when values at the same point in time on different lines must be precisely compared.

10.2.4 Boxes

Follow the principles for bar design, except when box plots are connected with a line to show change through time, which may require greater distance between boxes.

10.2.5 Useful Combinations of Objects

Objects should never be combined arbitrarily — combining bars for expenses with a line for revenue in a time series makes both patterns harder to see and harder to compare. There are three genuinely useful combinations:

Combination	Purpose	How It Works
Boxes + lines	Distributions through time	Boxes show each period's distribution; a line joins the centres (mean or median) to show how the centre moves through time.
Bars + lines (Pareto chart)	Ranked contributions to the whole	Named for Vilfredo Pareto, who found that ~20% of Italy's population owned 80% of the wealth — the 80-20 rule. Bars show each part ranked from greatest to least; a line accumulates the same values to show the cumulative percentage of the whole. Example: "people hear about us through the top three channels more than 80% of the time."
Bars + points	Uncluttered comparison	Bars carry the primary set (actual expenses); points carry the comparison set (budget). The primary values stand out, and the display is less cluttered and needs less space than two sets of side-by-side bars.

10.3 Secondary Data Components

10.3.1 Trend Lines

- A trend line describes the overall nature of a set of values; it does not represent actual values but summarises their essence.
- A straight line of best fit (linear trend line) in a **time series** is risky: it is based only on the values included in the graph. Adding a single earlier month (December 2010) reversed the direction of the trend in Few's example.
- Therefore **prefer a moving average (running average)** for change through time. A five-month moving average shows, for each month, the average of that month and the four previous months, softening jaggedness without the instability of a fitted line.
- In **correlations**, only use a straight line of best fit when the shape of the data is genuinely linear. Where response rises to a peak (four e-mails per month) and then falls, a curved line is needed to summarise the shape.
- Trend lines must always look different from actual values — lighter, dashed, thinner — so they are not mistaken for data. If the trend is the more important message, reverse the emphasis.

10.3.2 Reference Lines

- Reference lines provide context in the form of a meaningful comparison — e.g. last year's highest and lowest monthly revenues drawn across this year's line graph.
- They are especially useful for marking **a measure of the norm**: means, medians and standard deviations work well, making deviation from the norm immediately visible.
- Use reference lines to mark meaningful thresholds and regions.

10.3.3 Annotations

- By convention titles sit just above the data region and notes just below it. Notes may go inside the data region when they apply to particular data, provided they do not disrupt perception of the pattern.
- Uses: explaining an anomaly ("the result of a temporary reduction in sales staff"), marking events in time ("Marketing campaign begins / ends"), and labelling values that deserve to be featured (printing each line's highest month).
- Prominence should match importance. Reducing text from black to grey often prevents clutter — in the brain-mass-versus-body-mass scatter plot, grey labels can be read when wanted and ignored when studying the pattern.
- Use text to feature and comment on values directly when doing so is important to the story.

10.3.4 Scales

Scales are laid out along axes, dividing them into increments of equal length and assigning quantitative measures or categorical labels. Graphs rarely need more than two kinds:

Scale	Definition	Use and Cautions
Linear	The quantitative interval from one tick mark to the next is always the same; each tick adds a constant amount.	The default for almost all graphs.
Logarithmic (log)	Each value increases by a particular **rate** rather than a particular amount — each value is the previous value multiplied by the base. "The logarithm of a number is the mathematical power to which another number, called the base, must be raised to equal that number", e.g. $\log_{10}(10{,}000) = 4$.	Use log scales (1) to reduce the visual difference between data sets with very different magnitudes so they can be displayed together, and (2) to compare differences in value as percentages. Use only if readers do not need to compare values by the distances between them, and **always alert readers** that a log scale is in use. Do **not** use bars with a log scale — it is impossible to overcome the natural reading of a double-length bar as double the value; use points instead.

Worked examples: income vs. life expectancy for 179 countries (76 with incomes under $5,000) is unreadable on a linear X axis but clear on a log X axis; 2011 fruit sales with apples enormous and guavas tiny become readable as a dot plot on a log scale.

10.3.5 Tick Marks

- Mute tick marks in comparison to the data objects.
- Use tick marks with quantitative scales but not with categorical scales — except in line graphs where slightly more precision is needed.
- Aim for a balance: too many tick marks clutter the scale; too few make it hard to read values that fall between them.
- Avoid using tick marks to denote values at odd intervals.

10.3.6 Grid Lines

Thin, light grid lines may be used for three purposes only: to ease look-up of values; to ease comparison of values; and to ease perception and comparison of localised patterns. Anything heavier competes with the data.

10.3.7 Legends

- Use legends for categorical labels only when the labels are not associated with a categorical scale along an axis and cannot be attached directly to the data objects. Direct labelling is always preferable.
- Place legends as close as possible to the objects they label without interfering with other data.
- Render legends less prominent than the data objects they label.
- Use borders around legends only when necessary to separate them from other information.

10.4 Non-Data Components

10.4.1 Axes

- Two questions: how many axis lines, and what ratio of X to Y length?
- 2-D graphs conventionally use one vertical axis on the left and one horizontal axis along the bottom — sufficient in most cases.
- Moving the Y axis to the right makes values on the right easier to decode and draws the eye to the right half of the graph — useful when the final quarter matters most. Merely having visual content on one side that is missing on the other directs attention there.
- Four axes (a complete border) are useful mainly when the data region must be separated from surroundings that would otherwise compete for attention.
- An axis may be dropped where it hosts a categorical scale and horizontal bars encode the values, because the bars' edges trace the line the axis would draw. Vertical bars without a base, however, appear to float in space.
- **Aspect ratio** = the ratio of the data region's width to its height (width ÷ height). Where lines encode values, decreasing the aspect ratio increases the apparent degree of change: the same data at 2.0 looks moderate and at 0.5 "looks like a blast-off at Cape Canaveral". Both are accurate; they differ only in perception.
- **Never manipulate the aspect ratio to exaggerate or downplay change.** By convention, make time-series graphs wider than they are tall.

10.4.2 Data Region

Keep the background clean and light so that readers' eyes are drawn to the data. Background images make it difficult to focus on the data and must be avoided.

Consolidated Decision Rules

B.1 Choosing the Medium

If your message…	Use
is only one or two numbers	A sentence or a bullet point
will be used to look up or compare individual values, needs precision, mixes units of measure, or mixes detail and summary	A table
lives in the shape of the values — patterns, trends, exceptions — or compares whole sets of values	A graph
needs both	An integrated report: table + graph + text, tightly blended

B.2 The Encoding Hierarchy (Chapter 5 → 9)

Rank	Attribute	Use For
1	2-D position	Quantitative values (points, lines) and categorical items — the best at both.
2	Length (with a common baseline)	Quantitative values (bars, boxes). Requires a zero base.
3	Hue	Categorical items only. Max ~8 distinguishable hues.

Rank	Attribute	Use For
4	Colour intensity, size/area, width	Quantitative values only as a last resort — e.g. on maps. Perceived quantitatively but imprecisely.
5	Shape, orientation, enclosure, fill pattern, line style	Categorical items only; last resorts where colour is unavailable.

B.3 The Universal Design Checklist

Objective	Steps
Highlight	1. Subtract unnecessary non-data ink. 2. De-emphasize and regularize the remaining non-data ink. 3. Subtract unnecessary data ink. 4. Emphasize the most important data ink (width, orientation, size, enclosure, hue, colour intensity).
Organize	1. Group — proximity first; then similarity, enclosure, closure, continuity, connection. 2. Prioritize — quantitatively perceived attributes and contrast; watch alignment; use pointers sparingly. 3. Sequence — 2-D location first; reinforce with numbering when required.
Verify (graphs)	1. Does encoded quantity correspond to actual quantity? 2. Are tick-mark distances proportional to value differences? 3. Does the scale include zero (mandatory with bars)? 4. Is the display free of 3-D? 5. Is the aspect ratio honest?
Verify (tables)	1. White space first, then fill, then rules; never grids. 2. Numbers right-aligned, text left-aligned, decimals aligned. 3. Precision no greater than required. 4. Summaries visually distinct from detail. 5. Headers repeated on every page; What/When/Who/Where present.

Glossary

Term	Definition
Affordance	A characteristic of an object that reveals how it is meant to be used and makes it easy to use that way.
Aggregation	Summarisation — reducing large sets of numbers to a few numbers.
Aspect ratio	The ratio of a data region's width to its height. Strongly influences perceived degree of change.
Attentive processing	Conscious, sequential, slow visual processing in working memory; can result in learning and remembering.
Axonometric projection	The tilted, rotated, perspective-adjusted rendering software applies to 3-D graphs.
Bidirectional table	A table with categorical items laid out both across columns and down rows; a crosstab or pivot table.
Box plot / box-and-whisker plot	Tukey's display of a distribution: box for the midspread, centre mark for the median, whiskers to low and high values.
Bubble plot	A scatter plot using point size to encode a third variable.
Chart	The family of display methods that includes tables, graphs, diagrams and maps.
Choropleth map	A map that encodes aggregate regional values as colour intensities applied to the regions themselves.
Chunk	A meaningful unit of visual information held in working memory; only three or four fit at once.
Correlation	A comparison of two paired sets of values to see whether increases in one correspond to increases or decreases in the other.
Crosstab / pivot table	See bidirectional table.
Cumulative distribution	A frequency polygon accumulating occurrences from one end of the range to the other.
Data-ink ratio	The amount of ink presenting information compared to the total ink used (Tufte).
Deviation	A relationship expressing how values differ from a reference set of values.
Dot plot	A graph using points against a nominal or ordinal categorical scale; permits a non-zero scale.

Term	Definition
Drop lines	Support lines locating data objects against axes, especially in 3-D graphs; they clutter more than they help.
Empirical rule	In a normal distribution, 68%/95%/99.7% of values lie within 1/2/3 standard deviations of the mean.
Frequency distribution	The number of times something occurs within consecutive intervals across a quantitative range.
Frequency polygon	A line graph tracing the shape of a frequency distribution.
Gestalt	German for "pattern"; the school of psychology (from 1912) that described how we group what we see.
Hierarchical relationship	Multiple categories related as parent-to-child levels.
Histogram	A bar graph of a single distribution; bars conventionally touch.
HSL	The hue–saturation–lightness system for describing colour numerically.
Iconic memory	The visual sensory register holding each snapshot for under a second; the site of preattentive processing.
Interval relationship	Categorical items forming a sequential series of numerical ranges subdividing a larger range.
Jittering	Repositioning coincident points slightly so each can be seen in a strip plot.
Linear correlation coefficient	A single value from −1 to +1 giving the direction and strength of a linear relationship.
Log scale	A scale on which each value is the previous value multiplied by a base.
Mean / Median / Mode / Midrange	The four measures of average — see §2.2.1.
Midspread	The range of the middle 50% of a distribution (25th to 75th percentile).
Moiré vibration	The dizzying visual effect produced by fill patterns of vivid lines running in different directions.
Moving average	A running average of several consecutive values, used to show trend through time.
Nominal relationship	Discrete categorical items with no intrinsic order ("in name only").
Normal distribution	The symmetrical, bell-shaped frequency curve.

Term	Definition
Ordinal relationship	Categorical items with a prescribed order.
Outlier	An extreme value located far from most of the values in a data set.
Pareto chart	Bars ranked high-to-low plus a cumulative line; illustrates the 80-20 rule.
Percentile	The value at or below which a stated percentage of values falls; the median is the 50th percentile.
Preattentive processing	Automatic, unconscious, extremely fast visual recognition; tuned to specific attributes.
Ranking	Display order determined by the size of the associated quantitative values.
Ratio	A comparison of two values by division, expressible as a sentence, fraction, rate or percentage.
Reference line	A line marking a meaningful threshold, norm or region for comparison.
Regression line / line of best fit	A straight line drawn through the centre of the points in a scatter plot.
Rules and grids	Non-intersecting lines (rules) versus intersecting lines forming cells (grids).
Saccadic eye movement	The rapid jump of the eye between fixations, during which we effectively see nothing.
Scatter plot	A graph with quantitative scales on both axes, designed to display correlation.
Small multiples	A series of related graphs, each showing one comparison — the divide-and-conquer alternative to 3-D.
Spanner header / spanner rule	A header spanning several columns, and the rule that underscores it.
Spread	Highest value minus lowest value.
Stacked bar graph	Bars divided into segments representing parts of a whole; only the first segment reads directly off the scale.
Standard deviation	A measure of variation relative to the mean, computed from every value.
Strip plot	A distribution display plotting every individual value as a point.
Table lens	Side-by-side sorted bar graphs used to reveal correlation (Rao and Card).

Term	Definition
Trend	The overall nature of change during a period of time.
Unidirectional table	A table with categorical items laid out in one direction only.
Working memory	Temporary, capacity-limited memory (three to four chunks) where attentive processing occurs.