




4 The Mechanics of Visual Perception

Vision science is a vast inter-disciplinary domain of study, which incorporates insights from diverse fields such as molecular genetics, neuroscience, ophthalmology, optometry, physiological optics, psychology, cognitive science and so on.

For our purposes, however, it will be sufficient to remark that human visual perception is **fragmented** – our eyes are constantly scanning our environment, absorbing a multitude of stimuli and **seeking patterns** in them, as in the image below.



4.1 Background	64
A Simple Example	64
Pre-Attentive Attributes	66
4.2 Gestalt Principles	68
Simplicity	69
Continuation	69
Proximity	70
Similarity	71
Focal Point	71
Correspondence	72
Duality	73
Common Fate	74
Closure	74
Uniform Connectedness	75
Combinations	75
4.3 Gestalt Guidelines	76
4.4 Examples	77

Figure 4.1: On the left: *Girl from the Volga* (*Ogonek* magazine, 1959); on the right: record of eye movement during free examination of the photograph with both eyes, for 3 minutes [36].

On a chart, this process may look like the one shown in Figure 4.2.

We can see that the eye movement is somewhat **erratic**, but it is clearly organized: the first thing we notice is the graph, followed by the accented colour and the legend text, the chart title and subtitle, the axes, and finally back again to the graph itself for good measure.

In this chapter, we will discuss how good chart design can help lead the reader's eye and keep chaos to a minimum.

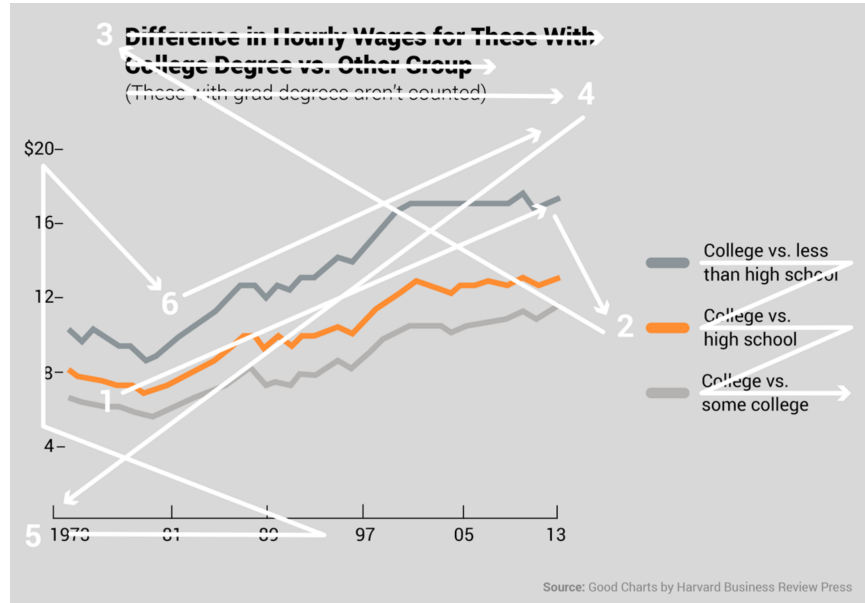


Figure 4.2: Eye movement on a chart [Good Charts, Harvard Business Review].

4.1 Background

Visual perception is achieved using **pre-attentive processes**, which are fast, instinctive, efficient, and with which we gather information from the bottom-up, moving our perception from features to patterns to objects, and **attentive processes**, which are slow and deliberate, in a top-down approach, allowing our perception to move from objects to patterns to features.

A Simple Example

The task of finding the number of “6s” in the image below is representative of tasks which engage attentive processes (the following example is based on a discussion in the excellent [4]) This is a time-consuming task, requiring us




Figure 4.3: How many 6s are there in this image?

to scan the image digit-by-digit, and row-by-row, while keeping track of the number of instances that are encountered. We view the image as a **collection** of digits, which can only be differentiated by their respective features (in this case, their shape), and a tremendous amount of effort is required to identify the feature of interest (the shape “6”).

There are ways to modify the image so that we accomplish the task by engaging pre-attentive processes instead. How many “6s” can be found in the image below? This time, the task can be achieved with a minimum amount

2869408609876
9348586748676
2967303986739
3967496749674

of effort. It is the **features** of the digits that most readily jump at us.¹ All that is required to answer the question is to count the number of highlighted digits: in this case, ten. The **contrast** provided by the use of different colour intensities breaks down the need for attentive processes, so to speak.

If this was all there was to it, there would be no need for this chapter: the best way to display information would simply be to provide the right contrast, and the important highlights of any chart/table would be readily available to the audience.

Part of the challenge, of course, is that the decision to highlight the “6s” was made **after** we asked you to look for “6s” in the image. If, instead, we had asked for the number of “8s”, the modification shown above would prove at best useless, and at worst, distracting (by introducing a distinguishable feature, the colour intensity, which is not linked to the task).²

One way out of this trap is to introduce a different colour for the “8s” in the image, as below. The red “8s” are still fairly easy to spot (and count), but they

2869408609876
9348586748676
2967303986739
3967496749674

are competing for attention with the black “6s”.

You may also eventually notice that the “8s” and the “6s” occur together more often than might be expected if the digits appeared completely randomly: with a few exceptions, a red “8” is immediately followed by a black “6”. But it takes a special effort to see these as “86s”; the colour difference is difficult to overcome.

In comparison, what are the “items” that you naturally recognize in the image below?

Figure 4.4: How many 6s are there in this image?

1: Their shape **and intensity** allow for easy differentiation.

2: Alternatively, if the “6s” have not been highlighted correctly, using the colour intensity as a pre-attentive process can lead to erroneous results: there are, in fact, *eleven* “6s” (the third-last digit of the 4th row was not highlighted above).

2869408609876
9348586748676
2967303986739
3967496749674

Figure 4.5: How many 8s are there in this image?

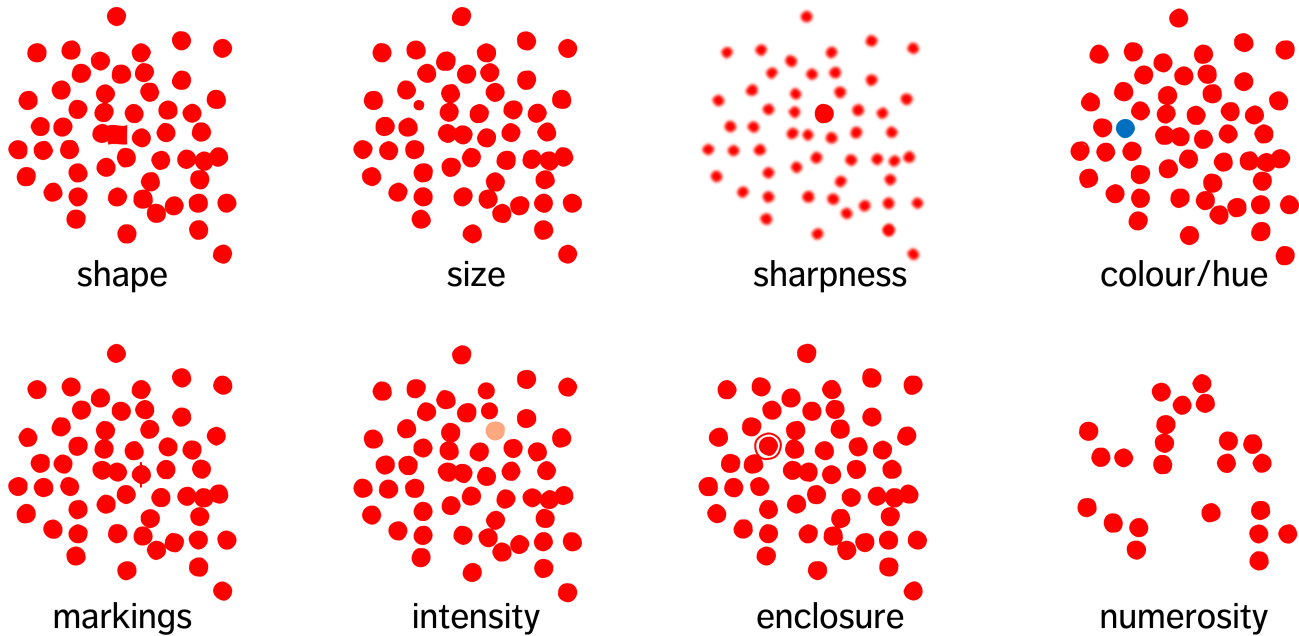


Figure 4.6: Pre-attentive features; how easily is the focal point observation identifiable in each display? (based on [9]).

2869408609876
 9348586748676
 2967303986739
 3967496749674

The colour invariance makes it nearly impossible to see the “6s” and the “7s” for the “67s”.

Pre-Attentive Attributes

The previous discussion provides some evidence to suggest that colour is a stronger trigger for pre-attentive processes than shape is. But might this only be the case for aligned digits?

When observations are represented by small two-dimensional markers, 8 modifiable features are available: shape, size, sharpness, colour (or hue), markings, intensity, enclosure, and numerosity. Not all features are created equal when it comes to pre-attentive processes, however: how easily can you identify the uniquely “interesting” observation in each of the examples from Figure 4.6?

Some of these features can be **combined** to make it easier to highlight an observation: it would be extremely surprising for someone to wrongly identify which observation is meant to be THE focal point in Figure 4.7.




Figure 4.7: Can you find the focal point which is identified through a combination of colour, size, shape, enclosure, and numerosity?

There is such a thing as “too much of a good thing”, however. As an example, we have seen that colour is a good differentiator (the black “6s” and red “8s” in the preceding images, for instance), but too many colours make the resulting display just as hard to parse as one with no colour whatsoever (see Figure 4.8).

2869408609876
 9348586748676
 2967303986739
 3967496749674

Figure 4.8: How many 6s, 4s, 9s, etc. are found in this image?

It would appear, then, that in visualization and dashboard design as in life, “a little goes a long way.”

Pre-attentive attributes are the domain of **iconic** (brief) memory (see Section 5.1, *Visualization and Memory*). They

- help to define a hierarchy of focus, and
- push non-message impacting components into the background.

For the most part, we use pre-attentive attributes to help **emphasize the story**. The main challenge is that highlighting one aspect of a chart can make other aspects harder to see.³

Before we move on to the next topic (Gestalt principles), we should note that there are other ways to facilitate the task of counting the number of digits, *as long as the order in which the digits appear is not crucial*.

For instance, we could re-arrange the digits as follows:

22333334444445
 66666666666777
 77777888888899
 99999999000

3: This can provide a justification for replicating the same base chart on multiple occasions in a dashboard; this may be acceptable if each replicate highlights a different aspect of the chart.

Figure 4.9: Ordering the digits to count them, with colours.

We could also place each digit's instances on a different row, to obtain relative frequencies:



Figure 4.10: Re-arranging the digits to determine the relative frequencies of the digits: were there more 7s than 8s? Note the effect of changing from a varying-width font (left) to a fixed-width font (right).

Of course, if the order of the digits matter (if they form a time series, say), then such solutions dramatically change the nature of the dataset and its insight, and are contra-indicated.

4.2 Gestalt Principles

“Gestalt” is the German word for “unified whole”. The **Gestalt principles** are the “laws” that govern human perception. They were first “identified” during the inter-war years by psychologists Wertheimer, Koffka, and Kohler in an attempt to describe how we group elements together, how we recognize patterns and simplify complex images when we perceive objects, in short, how we “gain meaning from the chaotic stimuli” surrounding us [37], and how our minds, seeking order, inform what we see by “perceiving a series of individual elements as a whole” [ibid].⁴

Designers often use these principles to organize content on charts, dashboards, websites, etc, in order to make them **aesthetically pleasing**, and **easy to parse** and **understand**. In this section, we discuss the following:

- simplicity
- continuation
- proximity
- similarity (invariance)
- focal point
- isomorphic correspondence
- figure / ground duality
- common fate
- closure
- uniform connectedness
- combinations

4: The expression “the whole is greater than the sum of the parts” provides a good example of Gestalt-style thinking.

Simplicity

The brain has a marked preference for **simplicity** – it tends to process simple patterns faster than patterns that are more complex. The lesson, then, is simple (!): arrange the data in your charts **simply** (and logically), whenever and wherever it is possible to do so.




Figure 4.11: Simplicity: who is the best performing salesperson in the sales force? Who had the fewest sales?

In the chart on the left of Figure 4.11, the reader has to look at each measurement sequentially while keeping track of the highest and lowest values, side-stepping the distraction provided by each bar's distinct colour, before identifying Nutty as the best performer, and Amy and Fred as having the fewest sales (tie).

The chart on the right is simplified by first removing the distinct colours and ordering the sales figures from highest to lowest, making it is easy to spot the best performer (left-most observation) and the lowest sales figures (right-most observations).

In this instance, the objective (finding extreme observations) is aligned with ordering the bars from largest to smallest (or *vice-versa*), so the simplicity principle readily kicks-in; it is not too difficult to imagine situations where that is not the case (such as time series or other charts where the sequence of events/data is crucial and must be reported in the right order).

Continuation

The simpler chart above is easier to parse, partly because the ordering creates a simple line to follow: it goes down slowly and **continuously** from left to right, without much vertical variability. In contrast, the line in the original chart zig-zags quite a lot, increasing the **cognitive load** on the reader.




Figure 4.12: Continuation: who is the best performing salesperson in the sales force? Who had the fewest sales?

Why should that be the case? According to Duke University's A. Bejan [38]:

“Scanning left to right is five times faster than scanning up and down and that is largely due to the left or right eye taking over

when the opposite eye gives up. When you scan vertically, it's like having just one eye. The eyes are also arranged on a horizontal axis, which happens to fit in with the landscape.

[...] The visual scene that is easiest to scan is one where the horizontal axis is wider than the vertical. This is true for a gazelle scanning the African horizon for predators, or early human hunter gatherers searching their territory for food.

Animal vision should be configured in a way that seeing and scanning should be the fastest and the easiest. And when the proportions allow this to be done, it should be a source of pleasure because of its past evolutionary associations with finding food or a mate.”⁵

5: This suggests that horizontal charts should be preferred to vertical charts, assuming the available space in the report/-dashboard allows it.

Proximity

Objects and shapes that are in **proximity** to one another appear to form groups; the effect generated by the group is perceived as more “powerful” (more organized?) than that generated by the separate elements, creating the **illusion** of shapes, even if the various elements are not contiguous, as in Figure 4.13.




Figure 4.13: Proximity: on the left, 9 grey squares; on the right, 9 grey squares and/or 1 grey super-square?

Groupings can be created through proximity that support the chart’s priorities. As an example, the same 16 sales observations are displayed in both charts of Figure 4.14; grouping the observations by region and by quarter, respectively, varies the chart’s focus.

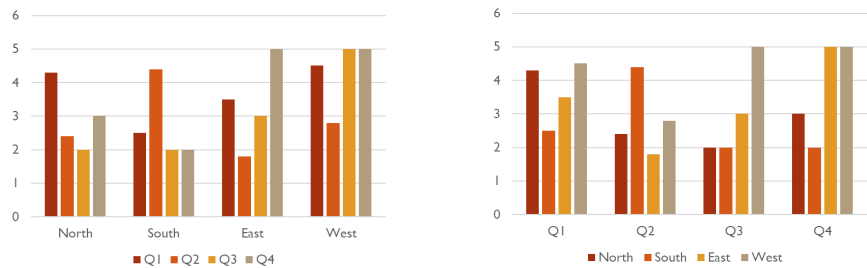



Figure 4.14: Proximity can be used to highlight different aspects of a datasets.

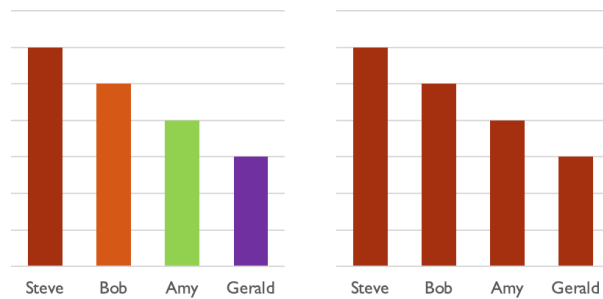
The effect of proximity is best exemplified by finding a title for the charts: the chart on the left could best be named “Sales per region, per quarter”, whereas the chart on the right, “Sales per quarter, per region”: the first layer of the focus differs from one chart to the other. If we understand the chart’s priorities, we can create proximity groupings that support them.

Similarity (Invariance)

When stimuli **physically resemble** one another, human brains view them as being part of the same category or object; stimuli that do not are viewed as part of a different object.⁶ We can use similar characteristics to establish relationships and to encourage groupings of objects.




In the examples above, similarity dominates over proximity: we see rows before we see columns; even though the last disk in the second row, say, is further from the first disc in that same row than it is from the sixth disc in the first row, it is clearly part of the same “row”.⁷



In the charts above, the different colours (on the left) leave us thinking that they must *mean* something; perhaps there are multiple charts on the page, and all elements relating to Steve are shown in red, to Bob in orange, to Amy in green, and to Gerald in purple. If that is not the intended interpretation, if the colours are simply used as decoration, making them all similar keeps the distractions to a minimum.⁸

Focal Point

In opposition to similarity, the focal point principle states that (a few) **distinctive-looking objects** can create focal points. To highlight one salesperson’s performance, for instance, we could make their bar graph colour different, as in the figure below.



6: Similarity and proximity often come together to form a **visual hierarchy**: either principle can dominate the other, depending on their application and combination.

Figure 4.15: Similarity can be used to create effects that overcome proximity.

7: We can still **see** columns, but it requires a tremendous amount of concentration and cognitive effort on the reader’s part to do so.

Figure 4.16: Similarity can be used to simplify complicated charts and reduce the cognitive load.

8: In charts, similarity often translates to markers of the same colour (or shape, size, and so on), being seen as representing observations from the “same” group: similarity in the chart makes it easier to spot **similarity in the data**, which reduces the overall cognitive load.

Figure 4.17: The focal point directs the reader’s eyes.

It would take a special individual indeed to be able to take a look at the chart on the right without at some point focusing on Peter's entry – focal points trigger **iconic memory**, which we will briefly discuss in Section 5.1 (*Visualization and Memory*).

In general, we can use the various pre-attentive attributes discussed above to create and/or highlight focal points, with an important caveat: if every observation is made special, then no observation is special – focal points work best when only a small number of observations are set apart.

Enclosures provide another way to bring out focal points; they work best when the observations of interest are close to one another, so that the box (or curve) is clearly there, but remains **unobtrusive**.




Figure 4.18: The enclosure focal point also directs the reader's eyes.

9: The other bars are only noticeable once we have cleared the enclosure hurdle.

In the above chart (on the right), it is clear that we are meant to focus our attention on Peter's and Bob's bars.⁹

Isomorphic Correspondence

People being people, they like people and stories about people (we discuss this further in Chapter 7, *Stories and Storytelling*). Stories spread less gingerly (and are easier to understand) when they use **storytelling tropes**, which is to say, shortcuts and conventions shared by members of a specific culture.

We interpret and respond to visual displays based on shared experiences (in particular, when it comes to chart colours): accounting practices and years of experience with spreadsheet software have made red a colour associated with **negative financial balances**, for instance (in opposition to black). As tempting as it might be to do so, we cannot assume that this meaning is universal.

Indeed, **context matters**. For instance, red is also associated with:

- a required **stop** for traffic lights, in opposition to green or yellow;
- **heat** and **warmth** on water faucets and air conditioning units, in opposition to blue;
- **left-leaning political positions** (right-leaning in the U.S.);
- **English Canada**, in opposition to French Canada (blue, again);
- **marriage** in Chinese culture, etc.

Shapes may also be imbued with "meaning", as is the case with Quebec and the Maritime provinces **traffic lights** ↻ (although the conventions change from one jurisdiction to another – *caveat exactoris*).¹⁰

10: 30 years removed from Latin classes, but we think this might mean "beware other drivers"?



With an appropriate understanding of conventions (which requires pre-knowledge of the chart's audience), we can convey additional (and subliminal) meaning by using shapes and colours judiciously.

In the chart below, the combination of the horizontal dashed line and the use of the colour red suggests that the smaller bar values are linked to **negative outcomes**.





Figure 4.19: A lifetime of experience tells us that the red columns represent unexpected events or disastrous results.

Figure / Ground Duality

The focal point principle works well when we clearly perceive chart elements as either being **figures** (in focus, foreground) or as part of the **(back)ground**. In general, objects in the foreground are **promoted** in importance by the brain, while background objects are **demoted**.

This is not to say that we do not see the latter, but rather, that it may take more time before our brains get around to “reading” them. This is desirable property of charts, as it reduces the cognitive load required to parse a chart: strong contrast between the foreground and the background makes it easier to distinguish between the two types of objects.¹¹

In the chart below, the low contrast between the figure and the background (on the left) increases the cognitive load: with some effort, we can tell that we are supposed to look at the bars, but that is much more demanding than what we see in the second and third chart.



11: The background may or may not be the chart's **canvas**.

Note that dark and light are not necessarily synonymous with background and foreground – the **contrast** between them is what matters.¹²

Contrast is also a good solution to a number of **accessibility** problems, a topic which we will discuss shortly (see Chapter 6, *Accessibility Considerations*).

Figure 4.20: Figure / ground duality in practice: not enough contrast (on the left), just enough contrast (middle, right).

12: Dark backgrounds and light foreground (such as text, for instance) make it easier to read at night; numerous applications give users the option to reverse these.

Common Fate

When lines or shapes **come together** (either in the sense of direction or location), a **relationship** between the underlying data observations is implied, whereas **askew** lines or shapes are perceived as unrelated or somewhat less related. We can thus use direction and/or movement to establish or negate relationships.

In the chart below, for instance, the line for “C” seems to belong to a different group of observations than those for “A” or “B”.




Figure 4.21: Common fate on line graphs: whatever is being measured and displayed in these charts, we would not be surprised to find out that groups A and B correspond to varieties of a similar product (cake mixes, say), while group C is linked to something altogether different (such as chain saws).

Closure

The human brain is a ridiculously effective **closure** machine: we see complete figures or forms even when the stimulus that reaches us is **incomplete**, if it is partially hidden by other objects, or if part of the information needed to make a complete picture in our minds is unavailable – it is why we can jump to conclusions so quickly,¹³ or why we can recognize all types of cats and dogs very early on even if we have only ever seen a few specimens.¹⁴

This reaction apparently stems from our mind’s natural tendency to **recognize familiar patterns**, filling in any information that may be missing along the way: do we need to explicitly tell you what you are looking at, below?



Closure is thought to have evolved from ancestral **survival instincts**: a partial sighting of a predator should trigger the mind to **automatically** complete the picture, so that we have enough time to react to potential danger even if not all the necessary information is available.

The main application of closure to charts is that it removes the need for **bounding boxes**.

13: Note that there is no mention of our conclusions being valid, only that we reach them quickly.

14: Strangely, this is something that our best A.I. still struggle with [Shaw ↗]:



Figure 4.22:

11:54 am · 12 Jun 2017
11.6K Retweets 1,825 Quote Tweets 20.9K Likes

Uniform Connectedness

When objects are **connected** to each other directly through line segments or curves, but also because they have the same colour or shape, or yet again if they overlap one another, we tend to perceive them as being part of a single uninterrupted object, separate from the other “unlinked elements.



Figure 4.23: In these charts, the connected line seems like the true trend, and the unconnected observations appear to be anomalies or outliers; of course, the data is exactly the same in both cases [39].

Applications could include using the same bullet shapes, colours, or numbering system for list items to promote the impression of groups, or displaying items of a similar nature inside a frame or specifically-coloured rectangle.

Combinations

When we see how individual Gestalt principles can be used to organize and create **structure** in a chart, it can be tempting to combine two or more of them on the theory that we cannot have too much of a good thing.¹⁵

In Figure 4.24, the **proximity** of the observations creates what seems to be four clusters (left-most column).

Since the proximity of observations is often **dictated by observation measurements** (the (x, y) coordinates), that principle is often overshadowed by any additional highlighting attempt. In the second column, we use **similarity** (top row), **connectedness** (middle row), and **enclosure** (bottom row) to partially erase the four clusters and to create new groupings.

We see:

- two types of observations in the first case (blue squares, orange discs);
- three in the second (observations linked by line segments), and again
- two in the last (the observations inside the enclosure, and those outside).

In the third column, we combine **connectedness with similarity** highlights in the top row, **similarity with enclosure** in the middle row, and **enclosure with connectedness** in the bottom row, to mitigated success. Arguably, the first combination provides the idea that there are two layers to the groupings: first, the ones provided by the connected line segments, then the one provided by the use of colour and shape; the other two are too confusing to parse properly (in this example, at least) – the hierarchy is lost.

And the combination of proximity with connectedness, similarity and enclosure creates a mish-mash where way too much is going on; this is not unlike what was seen in Figure 4.8.¹⁶

^{15:} Those of you of a certain age might remember how that worked out for *Smells Like Teen Spirit* in the early 90s. . .

^{16:} It would appear that the use of the Gestalt principles is governed by the meta-principle that “**less is more**”, as are so many aspects of the visualization process.

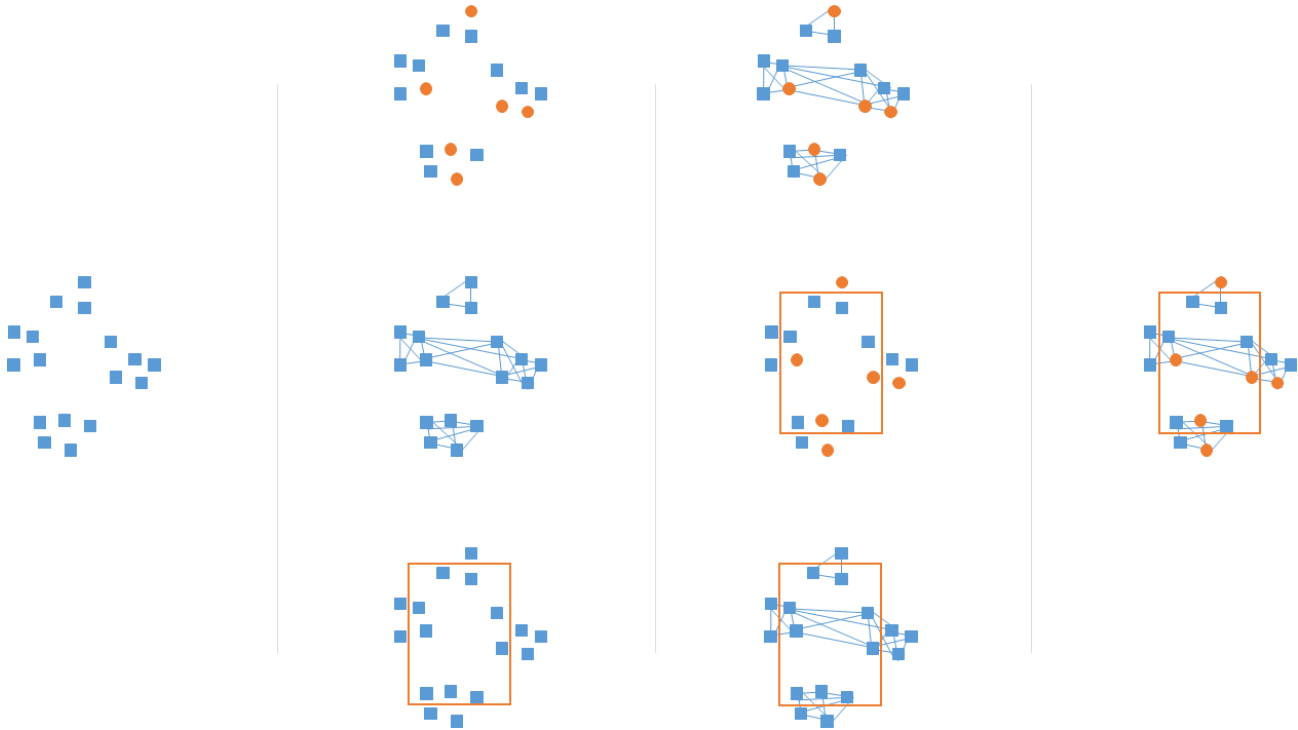


Figure 4.24: Use of various Gestalt principles (and combinations thereof) to recognize similarity in observations. Note the ambiguity as we move to the right.

4.3 Data Visualization and Gestalt Guidelines

The Gestalt principles thus suggest the following tips in relation to data visualizations [4, 9, 39]:

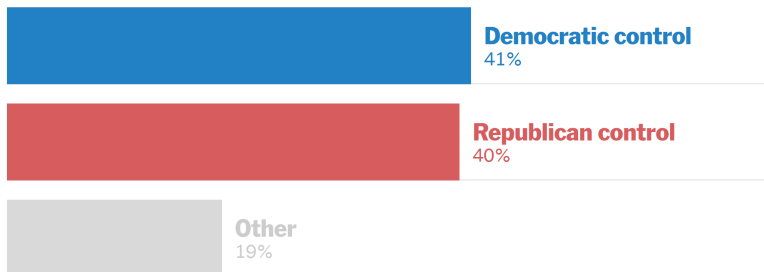
- the principle of proximity suggests that we can **do away with a legend** if we can place labels close to group of observations or curves;
- the principle of similarity suggests that we should **colour the labels** with the same colour as their corresponding observation groups;
- the principle of closure suggests that we do **not need borders around charts** since our brains automatically assume that the missing border is in fact there;
- the principle of uniform connectedness suggests that **outliers and anomalous observations** are easier to spot if they are kept disconnected from the main chart patterns;
- the principle of continuity suggests that we may be able to **remove the main axis** in a bar chart as all bars share a common baseline, and that chart elements should be aligned, when appropriate;

In effect, we can push readers to see what we think is important by removing distractions from charts; these suggestions are usually implemented when **decluttering** or when **evolving a storytelling chart** (which we will discuss in Sections 5.3 and 8.3, respectively).


4.4 Examples

Here are some examples of charts where at least one Gestalt principle has been clearly employed.¹⁷

What is your preference for the outcome of the 2022 congressional elections?



Based on a New York Times/Siena College poll of 849 registered voters from July 5-7.



Arrows on the map indicate the swing in each area - whether a neighbourhood went further towards Labour compared to 2017, or further to National. The size of the arrow shows how strong that swing was.

17: We make no judgement of their quality as data visualizations; they are simply charts in which the principles are easily identifiable.

Figure 4.25: Simplicity and continuity (a few ordered bars), use of isomorphic correspondence (blue: democrats; red: republicans) [40] (above); common fate (long-ish arrows mostly pointing left), use of isomorphic correspondence (red: Labour; blue: National) [41] (below).



Figure 4.26: Strong contrast (light text, dark enclosure), focal point (some Southeastern Africa countries plotted in burgundy), use of isomorphic correspondence (blue: water; green: land) [42].

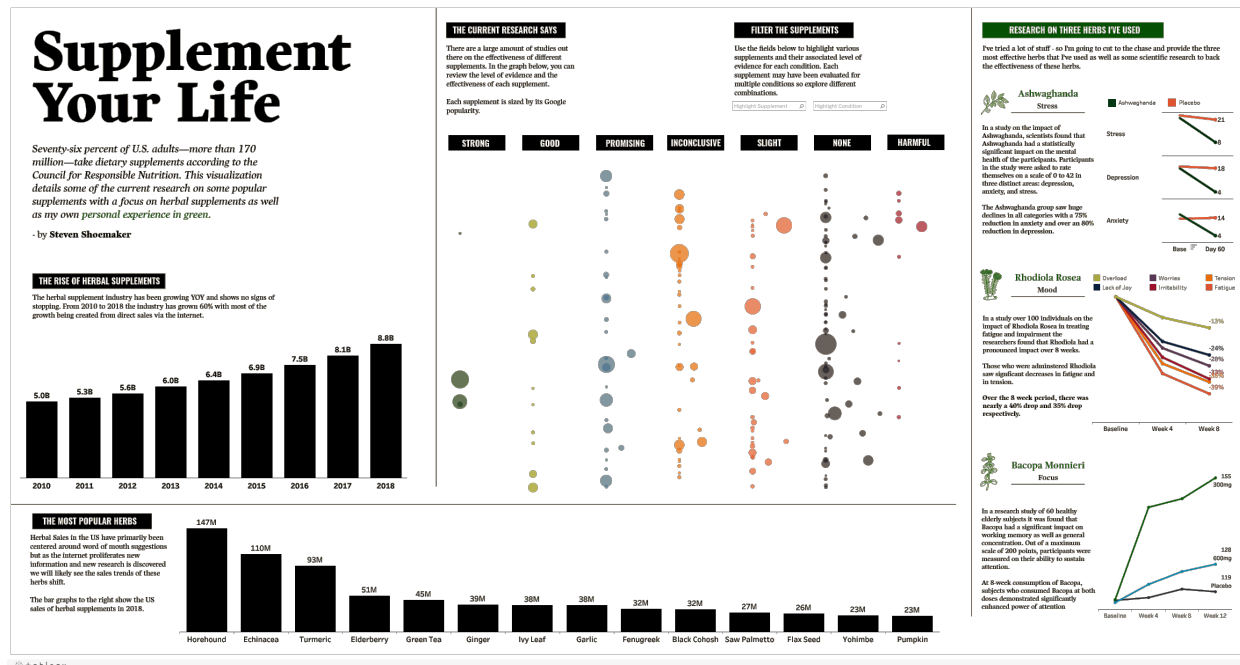


Figure 4.27: Figure / ground duality (black and white contrast), continuity (ordered bar charts) [Steven Shoemaker].

Global Happiness

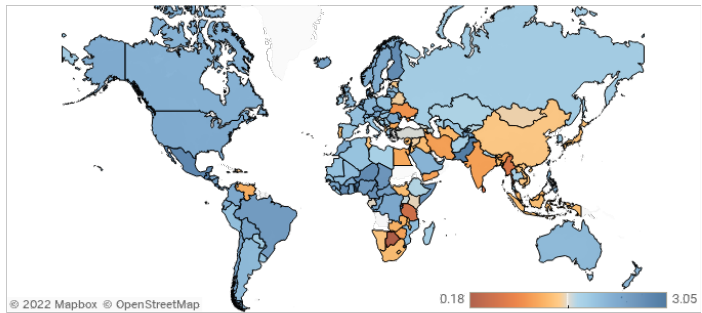
The **World Happiness Report** is an annual publication of the UN's Sustainable Development Solutions Network. It contains rankings of national happiness based on respondent ratings of their own lives, which the report correlates with various life factors. The 7 factors used for the Happiness Score are the following: Dystopia, Freedom to make Life Choices, GDP per Capita, Generosity, Healthy Life Expectancy, Perceptions of Corruption, and Social Support.

Take a look at the world map and below bar graphs to discover which countries have the highest happiness score based on these 7 factors.

Next, complete the quiz! Then explore the additional data on the last tab to see how relevant it is to the happiness score.



Dystopia by Country



Countries by Happiness Score Grouping

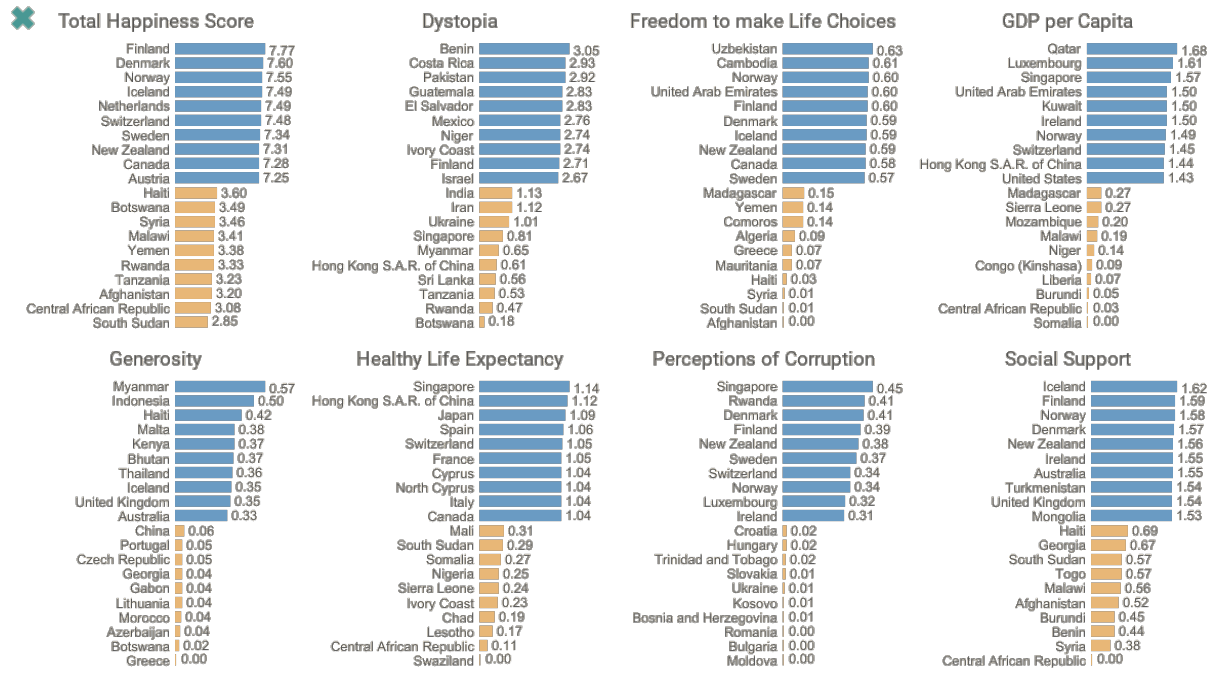




Tableau Developer:
Marit Imsdahl

Figure 4.28: Simple colour scheme, figure/ground duality (Latin America in black, background in the accent colour), enclosure of crucial title word [author unknown].

The inequality pandemic

The patient is Latin America, the most unequal region in the world. Throughout history, inequality has been a defining element of Latin-American societies. This can be perceived not only in the vast differences in wealth between rich and poor but also in the access to essential goods such as education or health. With the arrival of Covid-19, these inequalities have translated into more cases and deaths within the most impoverished areas in the region. To get a better understanding of the regional situation this visualisation focuses on the five biggest economies of Latin-America and their business centres: Mexico City (CDMX), Sao Paulo (SP), Bogotá (BOG), Buenos Aires (CABA) and Santiago (SCL).



Let's check on the patient ► **1. SYMPTOMS**



⊞ + a b | e a u

Figure 4.29: Similarity and continuity (ordered bars; blue: high scores; yellow: low scores), similarity (happiness score groupings colours) [Marit Imsdahl].