JDSSV

Journal of Data Science, Statistics, and Visualisation

May 2026, Volume VI, Issue III. *doi: 10.52933/jdssv.v6i3.172*

What is Text Doing in a Data Visualisation?

Paul Murrell
The University of Auckland

Abstract

This article discusses the role that text elements play in a data visualisation. Based on existing knowledge and frameworks for data visualization, we develop a framework that considers text as just another visual element that can be used to encode information, similar to bars and lines and data points. More specifically, we propose that data values can be encoded as the visual features of text, like the position or colour of text, as well as the letters, words, and sentences that make up the text. We make use of this framework to provide a deeper understanding of familiar uses of text, for example, in axes, legends, and titles. In addition, we show that this framework can be used to explore less familiar uses of text in data visualisation.

Keywords: Data visualisation, text.

1. Introduction

Figure 1a shows a grouped bar plot of the number of cats and dogs in different countries, according to the HealthforAnimals organisation (from September 2022). The choice of a grouped bar plot for showing these data can be justified in two ways:

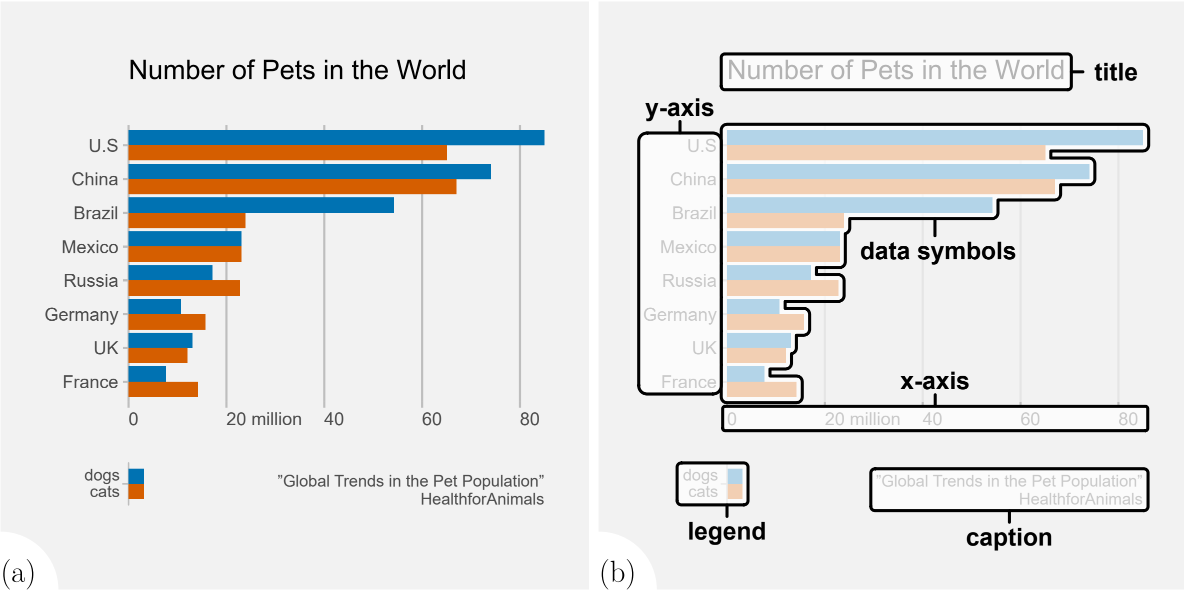

Figure 1: (a) A grouped bar plot showing the population of cats and dogs in major countries. This is a standard way to visualise count data broken down by two categorical variables. (b) The standard components of the plot: data symbols (the bars), a title and a caption, the x-axis and y-axis, and a legend.

1. We can use a data visualisation gallery that lists a variety of plots and dictates which plot to use based on the type of data variable that we have to display, for example, Wilke (2019, Chapter 5) or Wickham et al. (2023, Sections 1.4 and 1.5). In this case, we have one numeric variable (the number of pets) and two categorical variables (which country and which type of pet). The prescribed plot for that combination of variables is a grouped bar plot. We will refer to this as the *dogmatic* approach.

2. We can use a data visualisation framework that explains how different visual encodings work, for example, Munzner (2014), and select a plot based on visual encodings that work well. In this case, we have a numeric variable, and if we encode that as the lengths of bars, we can compare the number of pets accurately. We also have two categorical variables. If we encode the type of pet as the colours of bars, we can easily identify cats versus dogs. If we encode the country as the vertical positions of pairs of bars then we can easily identify different countries. The plot that emerges from that combination of encodings is a grouped bar plot. We will refer to this as the *reasoned* approach.

There are arguments for preferring a reasoned approach over a dogmatic one. For one thing, it is more satisfying to understand *why* we are making certain choices rather than just doing what we are told. This leads to a more intentional choice of plot. A reasoned approach is also more extensible. If we have a proper understanding of why some visual encodings work well, we can apply that knowledge to unfamiliar situations. A dogmatic approach cannot help us if the situation that we face is not covered by the existing rules.

Figure 1b shows a typical breakdown of the different components of a plot: data symbols, title, axes, legend, and caption (Cleveland 1994; Cairo 2019). The justifications given above for producing a bar plot are very focused on the data symbols. Both reasoned and dogmatic approaches are focused on how we are choosing to encode the data values. Choosing a bar plot arises from choosing bars as the data symbols that are used to encode the data values.

Can we also justify the choice of the other elements in Figure 1a? Can we justify the choice of the axes, legend, caption, and plot title? Text is a prominent component of all of these other elements. So the question of interest really is this: can we justify the different ways that we choose to use text in a plot?

As for data symbols, there are dogmatic approaches to choosing the text elements of a plot. For example, there are data visualisation guidelines that dictate that every plot must have a title (Schwabish 2021). Similarly, a standard edict is that there should always be an axis title and the values on a numeric axis should be evenly spaced (Robbins 2005).

However, where there is a reasoned approach to choosing data symbols, we are missing a reasoned approach to choosing the text elements of a plot. We are missing a satisfying explanation of how text works in a plot and we lack a basis for deciding how to use text in unfamiliar situations (Stokes and Hearst 2022; Hearst 2023).

In this article, we will develop a framework that explains how text works in a plot. We will use that framework to provide a reasoned justification for the choice of visual elements in Figure 1a and we will use the framework to explore some unfamiliar data visualisations as alternatives to Figure 1a.

2. A Reasoned Approach for Data Symbols

In order to develop a reasoned approach to the use of text in a plot, we need to elaborate on the reasoned approach that we use to choose the data symbols in a plot like Figure 1a. How we choose data symbols will form a basis for how we choose to use text.

The reasoned approach that leads to choosing bars as data symbols is based upon choosing how to encode data values as *visual features* (or visual channels, or aesthetics; Bertin 1983; Wilkinson 2005). For example, in Figure 1a, the number of pets is encoded as the *length* of a bar and the type of pet is encoded as the *colour* of a bar.

The effectiveness of a bar plot derives from the effectiveness of these visual features. For example, we know that encoding numeric data values as length leads to *accurate* comparisons (Cleveland and McGill 1984), and encoding categorical data values as colour hue is effective for *identifying* different categories (Munzner 2014).

Encoding numeric data values as *position* is another way to facilitate accurate comparisons, while encoding data values as angle, size, or colour luminance are less accurate. Encoding categorical data values as position or shape is also effective for identifying categories (Ware 2020). For example, Figure 2b shows an alternative to Figure 1a that encodes the number of pets as the horizontal *position* of data points, rather than as the lengths of bars, and encodes the type of pet as the *shape* of data points as well as the colour of the data points. The countries are still encoded as the vertical positions of the data points. This plot is also effective for comparing numbers of pets and identifying

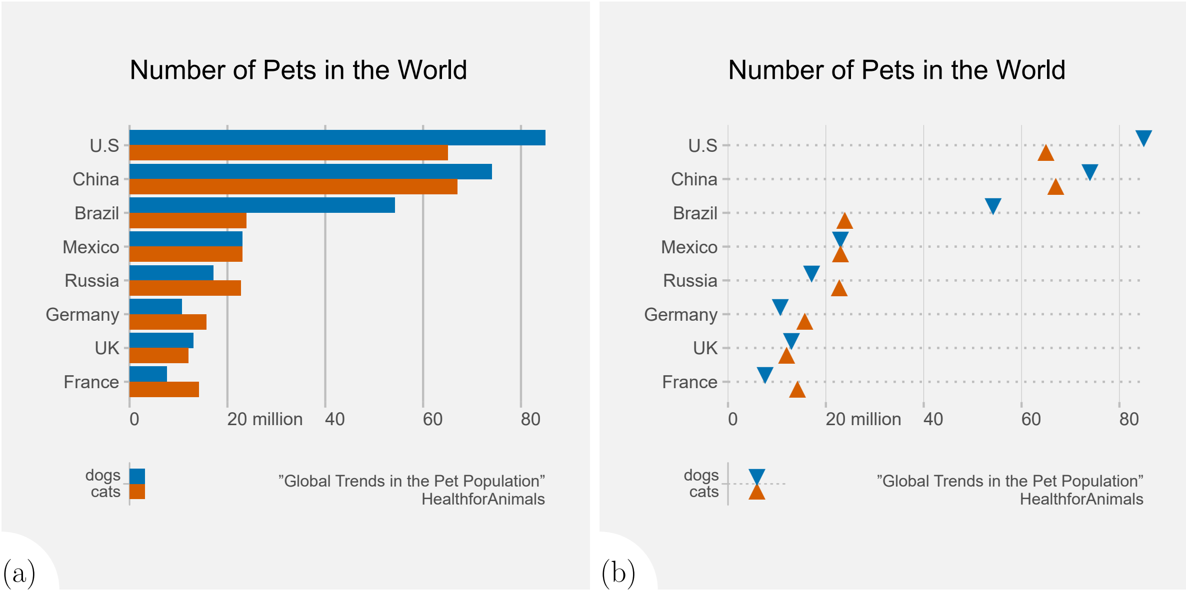

Figure 2: (a) The grouped bar plot from Figure 1a. (b) A dot plot using the same data as Figure 2a, but a different data symbol—data points instead of bars. The number of pets is encoded as the horizontal position of the data points, instead of being encoded as the lengths of bars.

different types of pets because it also uses effective encodings for those data values.

In addition to the encoding of data values as visual features, Figure 1a makes use of Gestalt Principles to form visual groups. For example, the *proximity* of each pair of bars signals that they both belong to the same country. The *similarity* of the bars—all bars for dogs are blue and all bars for cats are orange—signals that bars of the same colour correspond to the same type of pet (Wagemans et al. 2012).

A reasoned approach for choosing data symbols helps to explain the effectiveness of the bars in Figure 1a. The encoding of the number of pets as the lengths of the bars means that we can very easily see that the longest bar is more than four times the length of the shortest bar. In other words, we can make larger/smaller and ratio *comparisons* just based on the bar lengths.

However, there are other types of questions that we can answer with a plot like Figure 1a. A more expansive list would include:

(1) The largest count of pets is more than four times the smallest count of pets.

(2) There are more than 80 million dogs in the U.S. (plus similar information for each country and type of pet).

(3) The U.S. and China have more than double the number of pets compared to most other countries.

(4) Brazil has many more dogs than cats, while the counts for other countries are fairly similar. For Mexico, the counts are identical.

(5) There are more dogs than cats overall.

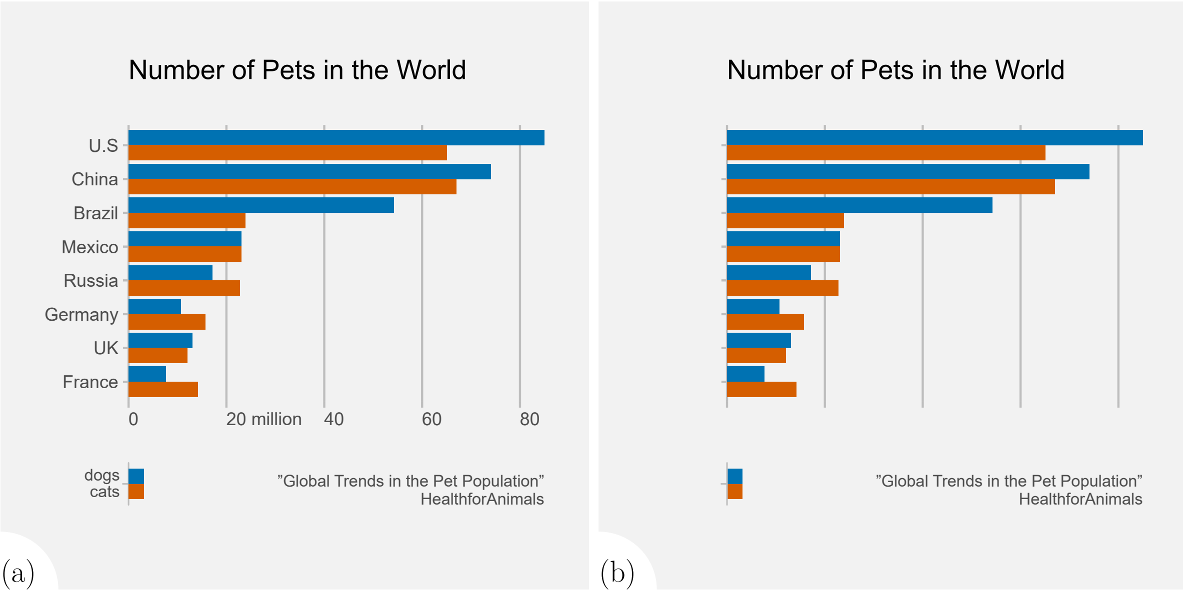

Figure 3: (a) The grouped bar plot from Figure 1a. (b) The same plot, but with all axis and legend labels removed. This demonstrates the limitations of geometric data symbols like bars—we are only able to answer a small set of questions just from the bars alone—and also demonstrates that text is playing an essential role in Figure 1a.

(6) The plot is about the number of pets in different countries.

(7) The information for the plot comes from the HealthforAnimals organisation.

A reasoned approach for choosing data symbols helps to explain why bars are effective for answering Question (1), but we will need a reasoned approach to text in order to make further progress on that list.

3. A Reasoned Approach for Text

Having established a reasoned approach to how data symbols like bars work in a plot, and what they are capable of, we can now begin to talk about a reasoned approach for the roles of the text labels in Figure 1a.

We will begin by observing some of the limitations of the data symbols in Figure 1a. In order to isolate the effectiveness of the bars in Figure 1a, Figure 3b shows a variation of the grouped bar plot with the text labels on the axes and legend removed.

Figure 3b reveals that we are unable to decode *absolute* numeric values just from the lengths of the bars. For example, we are unable to answer Question (2) because we cannot decode the number of dogs just from the length of a bar by itself.

Similarly, the position of the pairs of bars and the colours of the bars mean that we can very easily see that there are different groups of bars, but we are unable to decode the *absolute* identity of the groups just from the positions and colours of the bars. For example, we can see that one pair of bars in Figure 3b has identical lengths, but we do not know what country that corresponds to (or even that it corresponds to a country).

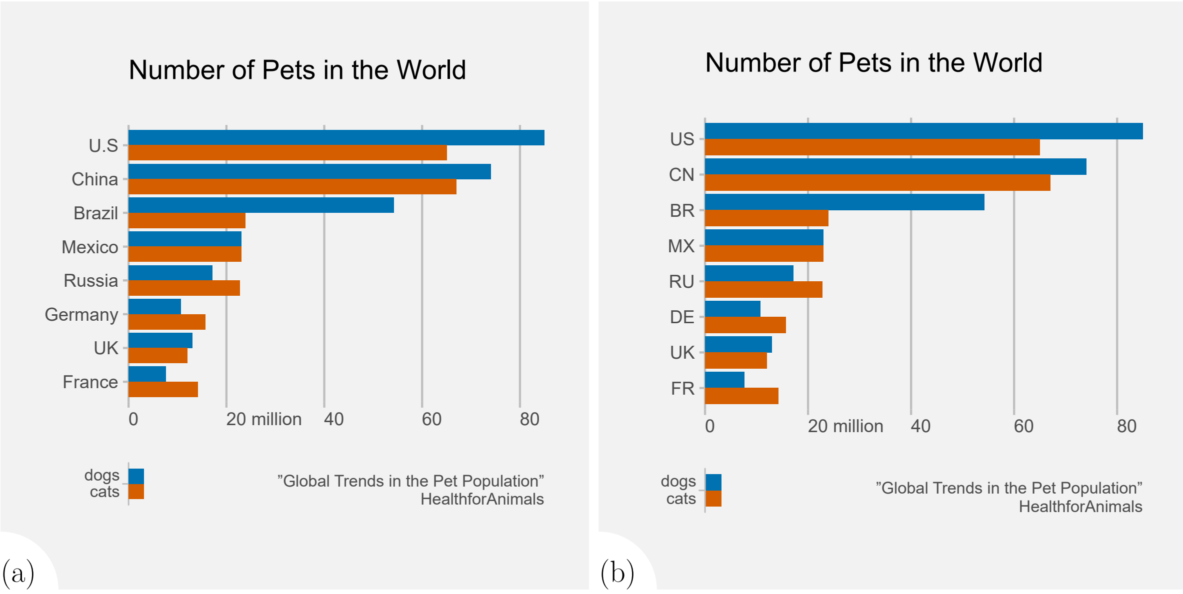

Figure 4: (a) The grouped bar plot from Figure 1a. (b) The same plot—in particular, the country names are encoded as text shapes on the y-axis—but the text shapes are two-letter abbreviations rather than the full country names. This makes it much harder to decode the country names from the text shapes.

It is clear that text is playing a fundamental role in answering even very simple questions using Figure 1a. In this section, we will build up a framework that allows us to explain the roles that text is playing.

The first thing to note is that the text labels on the axes and legend in Figure 1a encode data values just as the bars encode data values. This is most clearly demonstrated in the y-axis labels and the legend labels. For example, the data value `France` is encoded as a pair of bars *and* as the text "France" on the y-axis.

As we saw in the previous section with the bar data symbols, we will gain a clearer understanding if we consider that data values are being encoded, not just as text labels, but as specific *visual features* of the text labels. Furthermore, we must recognise that text labels possess many of the same visual features as data symbols like bars (Brath 2020). For example, just as the data value `France` is encoded as the vertical position of a pair of bars, the data value `France` is also encoded as the vertical *position* of the text "France". The *proximity* of the word "France" to the pair of bars means that we easily identify the text label and the bars as corresponding to the same data value.

In addition to encoding the data value `France` as the vertical position of the text "France", the data value `France` is also encoded as the *shape* of the text label—the letters that make up the text label—and this encoding is the key to the effectiveness of the text labels. A major problem with encoding data values as length, position, or colour is that we cannot decode *absolute* data values from length, position, or colour. However, we can decode absolute data values from the *shape* of text labels. For example, we can decode the data value `France` from the text shape "France".

The encoding of the data value `France` to the text shape "France" may seem rather obvious, but this is not the only possible encoding. For example, in Figure 1a the data

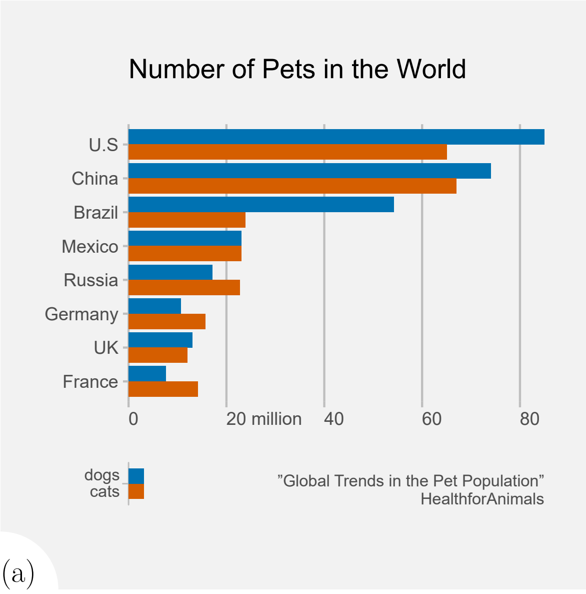

(a)

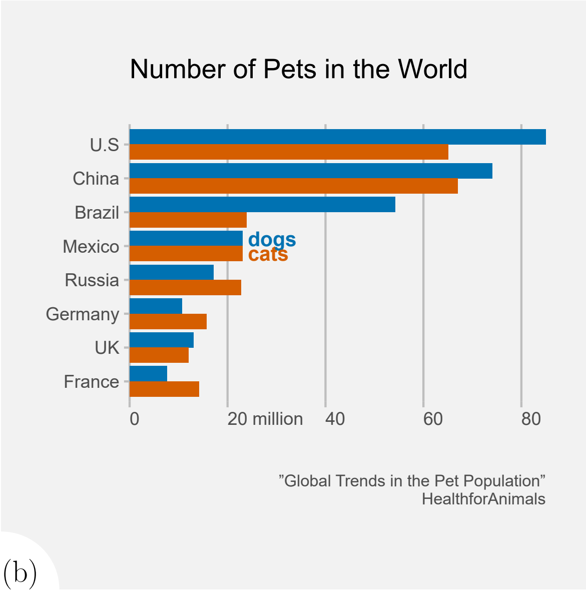

(b)

Figure 5: (a) The grouped bar plot from Figure 1a. (b) The same plot, but with direct labelling of pet type. There is no legend for decoding type of pet from colour in this plot. Instead, the type of pet is encoded as the text shape of labels placed directly on the plot beside the bars for Mexico.

value `United Kingdom` has been encoded as the text shape "UK". To emphasise the importance of the text shape encoding, Figure 4b shows a variation of Figure 1a with the countries all encoded as their two-letter abbreviations (ISO 3166-1 alpha-2). Figure 4b uses the same sort of encoding as Figure 1a—the country data values are encoded as the *shape* of text labels on the y-axis. However, it is much more difficult to decode the absolute country names from the two-letter abbreviations in Figure 4b. This shows that the choice of encoding from data values to text *shape* still needs to be carefully considered in order to ensure an effective decoding from the text shape.

By encoding countries as the vertical positions and shapes of y-axis labels, and encoding countries as the vertical positions of bars, we can answer Question (3) because we can decode from the text labels exactly which positions correspond to the U.S. and China, and we can decode a ratio exceeding two from the lengths of the corresponding bars compared to the bars of the other countries.

A similar case can be developed for the labels on the x-axis and the encoding of data values as the length, position, and colour of bars *combined with* the encoding of data values as the position and shape of text labels allows us to answer Questions (2), (3), and (4). The text labels are encoding information just as the bars encode information, but the *shapes* of the text labels are essential to decoding *absolute* numbers and countries and the *positions* of the text labels allow us to identify and compare numbers between countries.

The use of text labels on both the x-axis and the y-axis also demonstrates another strength of encoding data values as text shape. The y-axis labels encode qualitative data values (categories) while the x-axis labels encode quantitative data values (numbers). This demonstrates that we can use text shape to encode any sort of data value. In

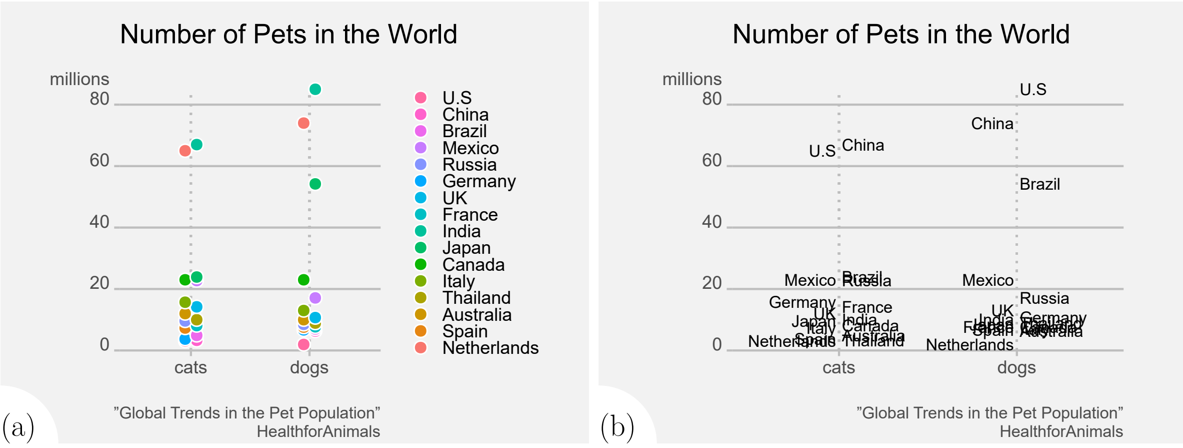

Figure 6: (a) A dot plot of the number of cats and dogs in different countries. This plot uses an expanded data set with more countries compared to Figure 1a. The different countries are encoded as the colours of the data points. Because of the limited capacity of colour hue, it is difficult to identify different countries because many hues are very similar to each other. (b) This plot uses text labels as data symbols instead of data points and encodes the country names as text shape instead of colour hue. This demonstrates that text shape has a much greater capacity than hue because it is much easier to identify different countries from the text shapes in this plot than it is to identify different coutries from different hues in Figure 6a.

other words, text shape is very *expressive* (Munzner 2014). This contrasts text shape with many other visual features, which tend to be less expressive. For example, the colour hue of a bar is effective for encoding different categories, like dog and cat, but is ineffective for encoding numbers, like the count of cats and dogs. The length of a bar is effective for encoding the count of cats and dogs, but is inappropriate for encoding different categories (Zhang 1996).

We can describe how the text in the plot legend works using reasoning similar to what we have just used for the axis labels. The encoding of pet type as the shape of the text labels in the legend allows us to decode the absolute pet types of `dog` and `cat` and the proximity of the bars and text labels in the legend allows us to identify which colour corresponds to each type of pet.

The legend involves an additional layer of indirection compared to the text labels on the y-axis because, having decoded pet type from the bars and text in the legend, the colours of the bars in the main plot have to be compared to the colours of the bars in the legend. Figure 5b demonstrates how *direct labelling* can be used to remove that layer of indirection. In Figure 5b, the vertical position of the text labels for dog and cat uses the same encoding of vertical position as the bars for Mexico. Furthermore, the colour of the text labels for dog and cat use the same encoding of colour as the bars for type of pet. This allows a more direct decoding of pet type.

Figure 6a shows a dot plot of an expanded data set on the number of pets in 16 different countries. In this plot, the type of pet is encoded as the horizontal position of data points, number of pets is encoded as the vertical position of the data points, and the country is encoded as the colour (hue) of the data points. The purpose of showing this

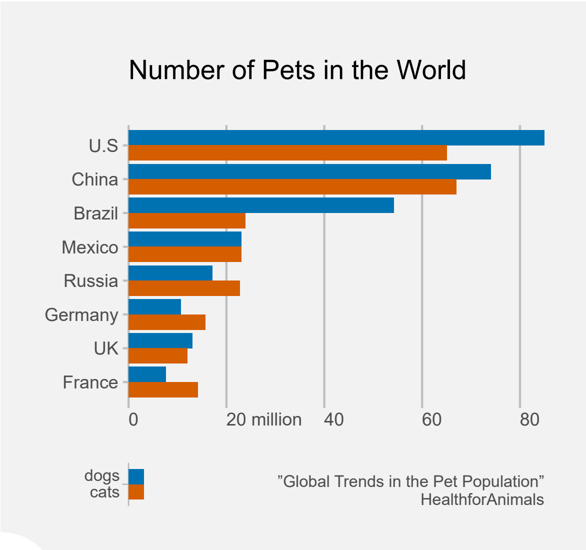

(a)

country	pet	count (millions)
U.S	dogs	85.00
	cats	65.00
China	dogs	74.00
	cats	67.00
Brazil	dogs	54.20
	cats	23.90
Mexico	dogs	23.00
	cats	23.00
Russia	dogs	17.10
	cats	22.75
Germany	dogs	10.70
	cats	15.70
UK	dogs	13.00
	cats	12.00
France	dogs	7.60
	cats	14.20

(b)

Figure 7: (a) The grouped bar plot from Figure 1a. (b) A table display of the same data. It requires more time and effort to compare between countries and to extract trends from this table compared to the grouped bar plot.

plot is to demonstrate the limited *capacity* of hue to represent many different categories (Healey 1996). With this many different hues, it is very difficult to decode which points correspond to which countries because many of the hues appear very similar to each other.

Figure 6b shows a variation on Figure 6a with text labels instead of data points. The number of pets is still encoded as the vertical position of the text, but the country is now encoded as the text shape. This demonstrates that text shape has a much higher capacity than hue. It is very easy to decode the different countries from the different text shapes even though there are many different countries.

One problem with both Figures 6a and 6b is that several of the data points and several of the text labels overlap with each other for lower counts of pets. Although we can decode that there are a larger number of countries with fewer pets, thanks to the density of the vertical positions of data points and text labels, it becomes impossible to identify individual countries. This emphasises that, if the text shape is not legible, its effectiveness for decoding absolute values disappears.

This problem of legibility leads to another important point. Although we have identified a number of ways in which it is very effective to encode data values as text shape—and in several cases, this is more effective than encoding data values as more common visual features such as colour and length—encoding data values as text shape is not in all ways superior to other ways of encoding information. In the next section, we consider some of the reasons why text shape is not always the best encoding.

4. Weaknesses of Text

In Figure 3b, we removed the axis and legend labels from the bar plot in Figure 1a in order to demonstrate the importance of the text labels. The table in Figure 7b goes in the opposite direction by removing the bars and producing an entirely text-based display of the data values. This will help to demonstrate some of the limitations of text labels.

As described in the previous section, encoding data values as the shape of text labels is extremely valuable because it allows us to decode *absolute* data values. It is also worth pointing out that the decoding of text is extremely *accurate*, even more accurate than the highly-rated length of bars (Shah and Hoeffner 2002). For example, we can decode the exact data value `17.10` from the text "17.10" in the table in Figure 7b and this is a much more accurate value than we are able to decode from the length of the blue bar for Russia in Figure 1a.

However, it is still much easier to answer questions like (2), (3), and (4), which involve simple comparisons, from Figure 1a than it is to answer those questions from the table in Figure 7b. One reason for that is because the lengths of bars are *congruent* with the number of pets, but the shape of text is not (Kosslyn 2006). A bar that is about twice as long as another bar naturally corresponds to a ratio of two and much more so than two text labels like "65.00" and "23.90".

A related issue is that the decoding of bar length is an immediate and unconscious process, but the decoding of text requires more processing and more deliberate congnitive effort (Sorapure 2019). For example, it is possible to calculate a ratio of approximately two from the text labels "65.00" and "23.90", but it requires deliberate mental effort to do so. Overall, two bar lengths have a stronger and more immediate visual impact than two text labels.

It is also easier to answer Question (5), which involves summaries of the data, from Figure 1a compared to the table in Figure 7b. The reason for this is that we can visually summarise a collection of bar lengths, or a collection of data point positions like in Figure 2b (Szafir et al. 2016; Whitney and Yamanashi Leib 2018), but we cannot easily do the same for a collection of text labels (Ware 2020). For example, we can easily see that the shorter bars in Figure 1a (Russia, Germany, UK, and France) average about 17 million. While it is possible to average the corresponding text labels in the table in Figure 7b, it is much slower and requires a much larger mental effort.

5. The Uniqueness of Text

We have so far ignored Questions (6) and (7), which involve information about the plot as a whole—metadata rather than data values. The reason for this is that the data symbols that we began with, the bars in Figure 1a, are of no use for answering these sorts of questions. Simple data symbols like bars, lines, and points are only useful for encoding simple data values. Text is the only option for encoding more complex information (Weber 2019).

We previously saw that the effectiveness of text labels is largely due to the ability to decode *absolute* data values, for example, from axis labels and legend labels. However,

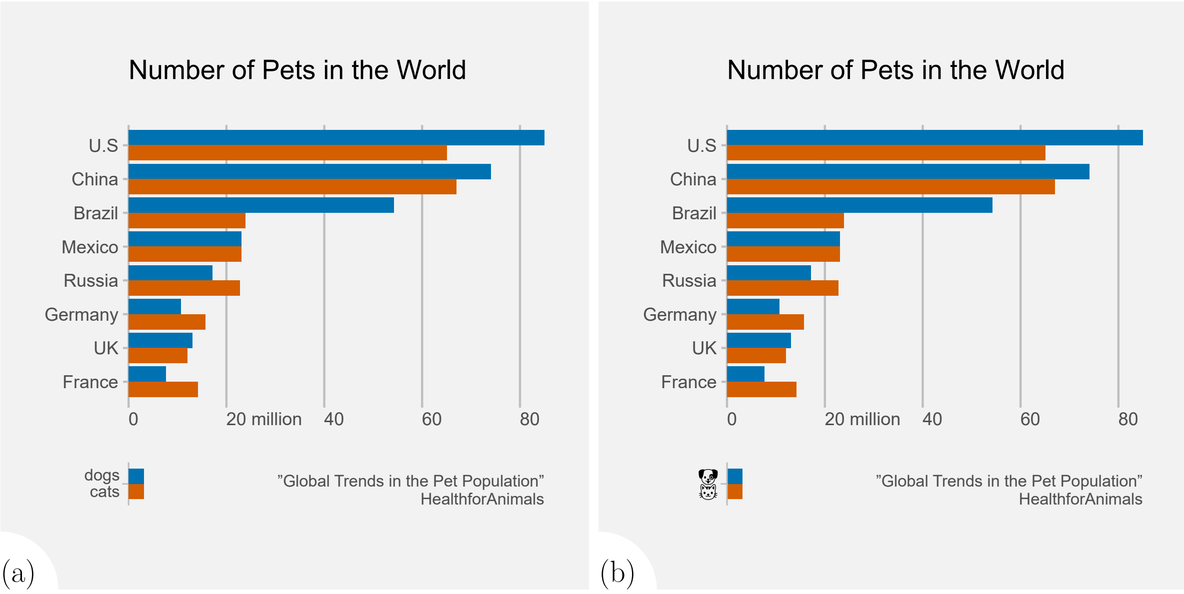

Figure 8: (a) The grouped bar plot from Figure 1a. (b) The same plot, but with cat and dog icons on the legend instead of text labels. This demonstrates that is is possible to decode absolute data values, like dog and cat, from data symbols other than text labels.

text labels are not the *only* way to encode absolute values. For example, Figure 8b shows a variation of Figure 1a with *icons* used instead of text labels on the legend. We can decode the data values `dog` and `cat` from these small images of a dog face and a cat face, just as we can decode those values from the text labels "dog" and "cat".

Nevertheless, text labels are much more flexible, less ambiguous, and more composable than icons (Ware 2020). Text has the unique capability to encode arbitrarily complex information. This makes it uniquely valuable for answering questions like (6) and (7). Rather than dogmatically stipulating that a plot must have a title and a caption, we can reason that text, in the form of a title and caption, is the only option for encoding complex information like the overall context and message of the plot and the source of the data set.

This may seem like another obvious point, but we can still gain understanding from viewing titles and captions as an encoding of information using the shape of text. In particular, as we saw in Figure 4b, the choice of text shape is important. If we choose to encode information as a poor text shape, then it will not be possible to decode information from the text label. Put in more familiar terms, it is important to choose our words carefully.

Figure 9b shows another way in which the choice of text shape in larger text elements matters. The choice of words in a plot title can have a large influence on the message that is taken away from a plot (Kong et al. 2019). We need to be careful about the information that we choose to encode as the shape of text for a plot title.

This also relates to another weakness of text. The choice of an effective encoding from data values to text shape depends on educational and cultural backgrounds. Although it may seem obvious to an English-speaking audience to encode the data value `Germany`

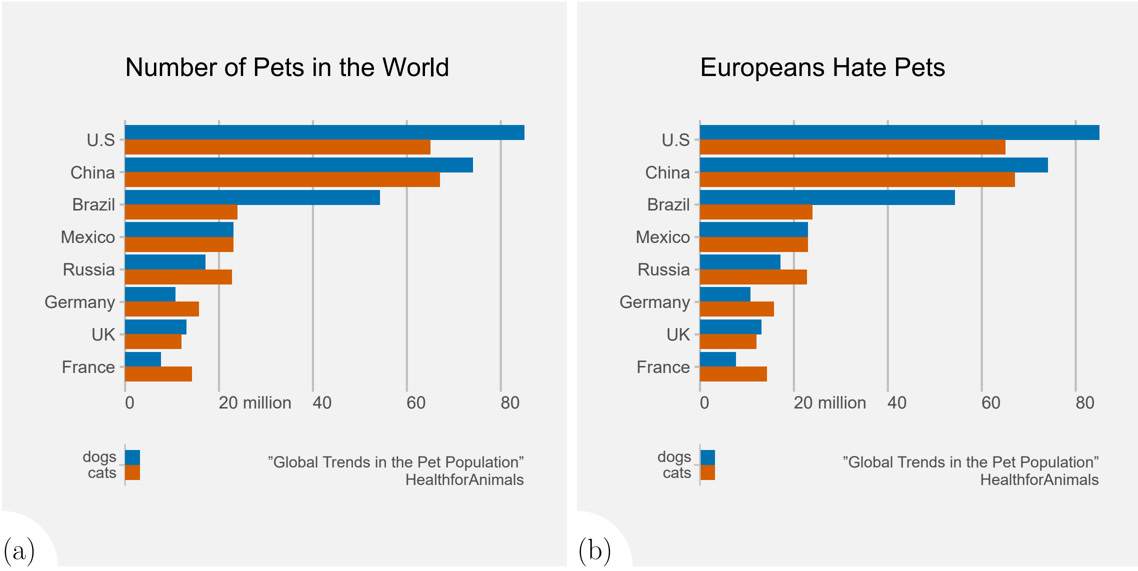

Figure 9: (a) The grouped bar plot from Figure 1a. (b) The same plot, but with a different title. The title provides a simple demonstration that the choice of text shape also matters when encoding complex information. The title on a plot can have a large influence on the message that is taken away from a plot.

as the text shape "Germany", a German audience might gravitate more towards the text shape "Deutschland". Similarly, the encoding of Germany as the two-letter code "DE" (Figure 4b) will be more or less difficult to decode depending on the viewers' educational and cultural background. By comparison, encoding data values as visual features like position or length are much less vulnerable to this sort of ambiguity (Willey and Liu 2022).

It is also important to note that larger text elements, like titles and captions, can also encode information as visual features other than text shape. For example, the text in the title at the top of Figure 1a differs from the text in the caption at the bottom of Figure 1a in terms of *size* as well as in terms of shape. The size of the title versus the caption encodes the relative importance of these different elements. The more important title is much larger and the less important detail of the data source is much smaller (Brath and Banissi 2016).

Figure 10b shows another example of encoding information in a plot title using visual features other than shape. In this case, the type of pet is encoded as the *colour* of specific words in the title. This is an interesting encoding because we have data values encoded as the colour of a *subset* of the overall text shape that makes up the title. This is something that is possible because text labels are typically made up of many pieces (letters), so we can encode information as subsets of the text (Brath and Banissi 2017). An analogy can be made to the encoding of different data values as different facial features in a Chernoff faces plot (Chernoff 1973). As a result of the additional encoding of pet type as the colour of parts of the title in Figure 10b, we have no need for a legend because the type of pet can be decoded from the shape of the coloured words and the similarity of the coloured words to the bar colours creates visual groups that allow different types of pet to be associated with different bars.

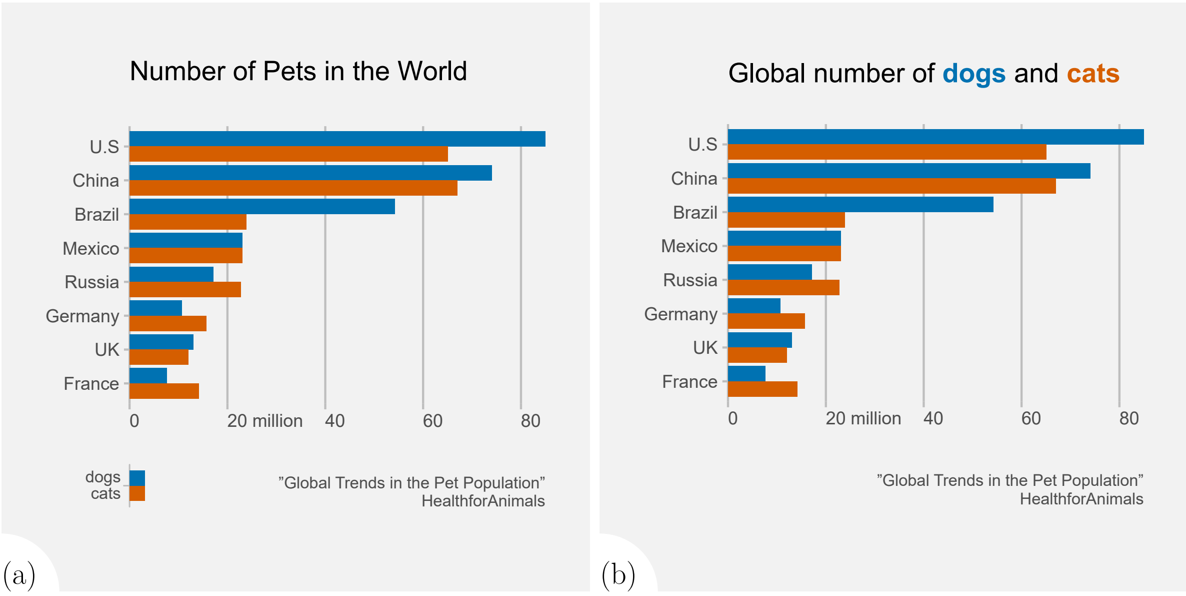

Figure 10: (a) The grouped bar plot from Figure 1a. The same plot, but the title contains the words "dog" and "cat" with the type of pet encoded as the colours of these words. As a result, there is no need for a legend.

6. Applications of a Reasoned Approach

We have so far established a framework that allows us to understand how very familiar uses of text—titles, captions, axes, and legends—work in a very familiar data visualisation—a grouped bar plot. The framework can also be used to understand how text works in other sorts of familiar plots, like word clouds and stem-and-leaf plots (Murrell 2026, Chapter 13). However, in this section, we will instead use the framework to gain insights on some less familiar uses of text in less familiar data visualisations.

We begin with the table in Figure 7b. Although this table consists almost entirely of text, our framework suggests that we can still analyse it in the same way that we analyse a data visualisation like Figure 1a. In particular, we can analyse the table in Figure 7b in terms of how data values are encoded as the *visual features* of text.

It is obvious that the table in Figure 7b encodes many data values as the *shape* of text elements. These are the words and numbers in each cell of the table. However, our framework allows us to see that data values are also encoded as other visual features of the text in the table. For example, country and pet type are encoded as the vertical positions of the text in the table as well as the shape of the text in each table cell. The arrangement of the data values into rows and columns in Figure 7b reflects a deliberate encoding of data values as *positions*. On the other hand, the number of pets is *only* encoded as the *shape* of the text in the third column of the table in Figure 7b.

The encoding of country and pet type data values as vertical and horizontal positions—as rows and columns—means that we can easily identify different groups of data values. Encoding data values as text shape means that we can easily decode the exact combination of country and pet type for each row and the exact count of pets for that combination. This allows us to answer questions like (2) because that just requires decoding a small number of specific data values: a specific country, a specific type of

country	pet	count (millions)
U.S	dogs	85.00
	cats	65.00
China	dogs	74.00
	cats	67.00
Brazil	dogs	54.20
	cats	23.90
Mexico	dogs	23.00
	cats	23.00
Russia	dogs	17.10
	cats	22.75
Germany	dogs	10.70
	cats	15.70
UK	dogs	13.00
	cats	12.00
France	dogs	7.60
	cats	14.20

(a)

country	pet	count (millions)
U.S	dogs	85.00
	cats	65.00
China	dogs	74.00
	cats	67.00
Brazil	dogs	54.20
	cats	23.90
Mexico	dogs	23.00
	cats	23.00
Russia	dogs	17.10
	cats	22.75
Germany	dogs	10.70
	cats	15.70
UK	dogs	13.00
	cats	12.00
France	dogs	7.60
	cats	14.20

(b)

Figure 11: (a) The table from Figure 7b. (b) The same table, but the rows for dogs have all been coloured the same blue and the rows for cats have all been coloured the same orange. In other words, the type of pet has been encoded as the colour of the text in the table.

pet, and a single number.

However, as we saw in Section 4, the table in Figure 7b is ineffective for answering questions like (5), which involve decoding summaries of multiple data values, because we cannot easily summarise even a quite small number of text shapes.

In summary, tables of data values are useful for looking up individual data values, but poor for decoding trends across multiple values. This is familiar knowledge, but we have used the framework to provide a reasoned understanding of that knowledge. We can also go further and use the framework to think about what might work better.

One important element of the framework is the idea that we can encode data values as visual features of text labels *other than* the text shape. For example, the table in Figure 11b shows that we can add the colour encoding of the type of pet to the table in Figure 7b. This is a demonstration that we can encode data values not just as the shape of text, but also as the colour of text. The result is that we can much more easily identify the rows for cats and dogs compared to the table in Figure 7b because the type of pet is encoded as both the vertical position of the text labels *and* as the colour of the text labels. The vertical positions are useful for grouping a particular combination of country, count, and pet type and the colours are useful for grouping all data values that correspond to the same type of pet.

Another way that we could improve the table in Figure 7b is to allow the decoding of data summaries from the counts in column three of the table in Figure 7b. Figure 1a and Figure 2b demonstrated that encoding data values as the lengths of bars or as the positions of data points does allow us to decode data summaries, like the average number of pets for European countries. This suggests that we could try encoding the count data values as the *position* of the text in the table in Figure 11b.

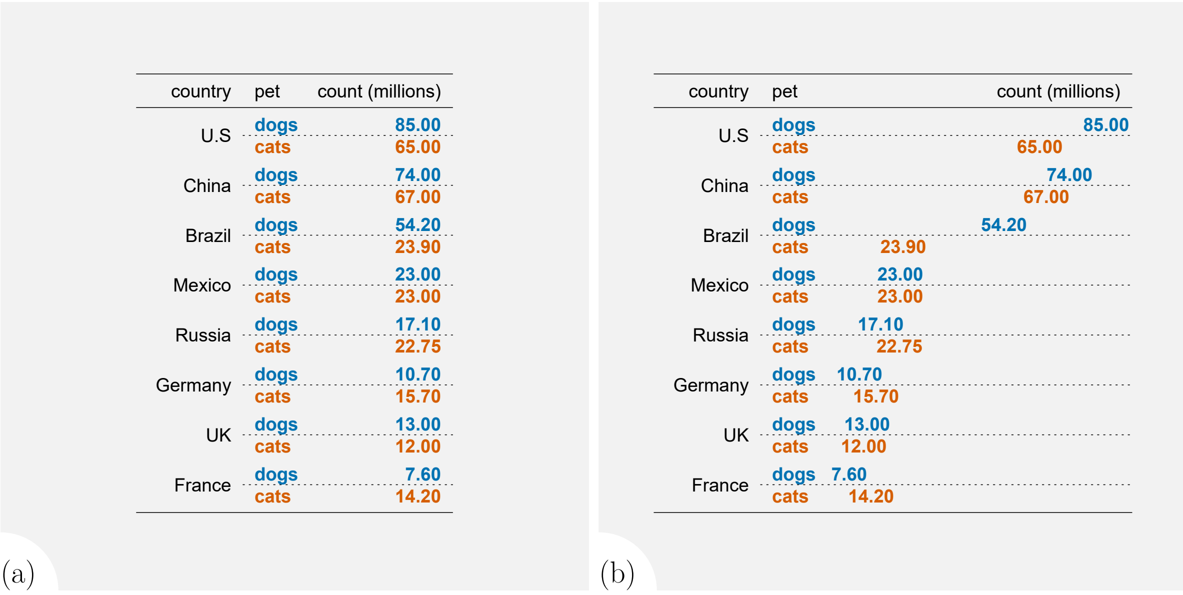

(a)

country	pet	count (millions)
U.S	dogs	85.00
	cats	65.00
China	dogs	74.00
	cats	67.00
Brazil	dogs	54.20
	cats	23.90
Mexico	dogs	23.00
	cats	23.00
Russia	dogs	17.10
	cats	22.75
Germany	dogs	10.70
	cats	15.70
UK	dogs	13.00
	cats	12.00
France	dogs	7.60
	cats	14.20

(b)

country	pet	count (millions)
U.S	dogs	85.00
	cats	65.00
China	dogs	74.00
	cats	67.00
Brazil	dogs	54.20
	cats	23.90
Mexico	dogs	23.00
	cats	23.00
Russia	dogs	17.10
	cats	22.75
Germany	dogs	10.70
	cats	15.70
UK	dogs	13.00
	cats	12.00
France	dogs	7.60
	cats	14.20

Figure 12: (a) The table from Figure 11b. (b) The same table, but the horizontal positions of the text labels in the third column reflect the magnitude of the number of pets. In other words, the number of pets is encoded as the horizontal position of the text in the third column of the table.

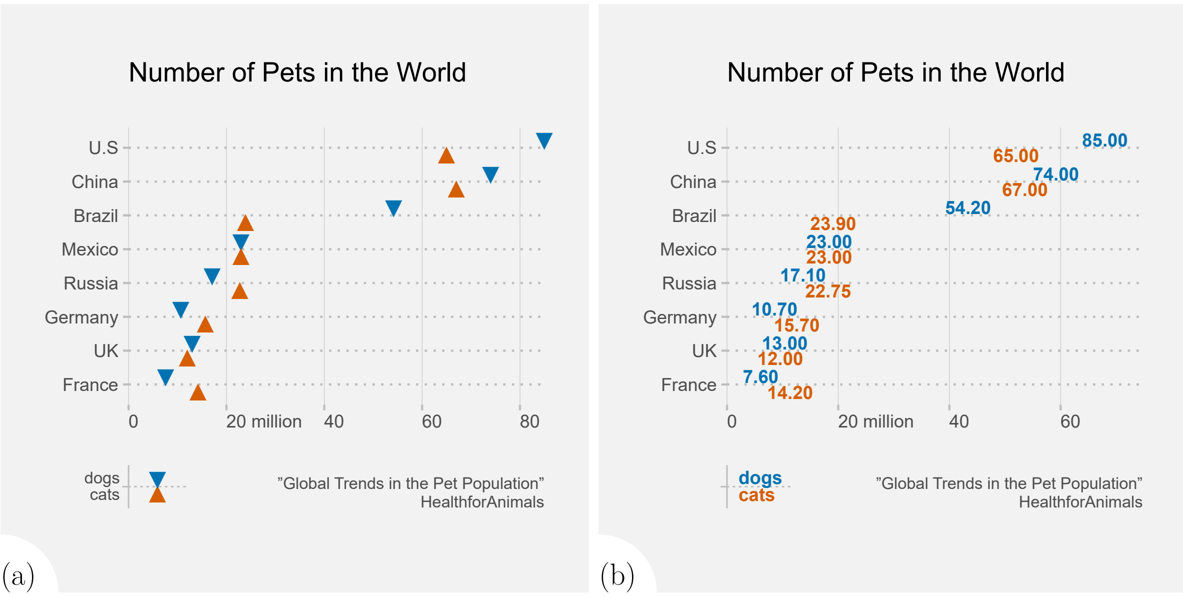

Figure 13: (a) The dot plot from Figure 2b. (b) A plot with text labels as data symbols instead of data points and the number of pets encoded as the shape of the text labels as well as the horizontal position of the labels. This means that we can decode absolute numbers from the text shape as well as decoding summaries of pet numbers from the text positions.

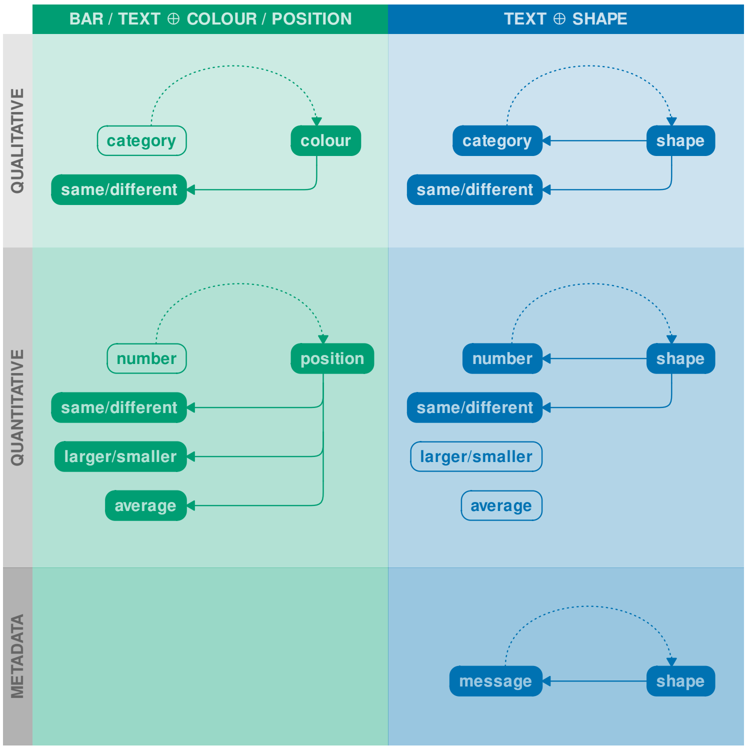

Figure 14: A summary of the framework. The green column shows the qualities that text shares with traditional data symbols like bars or points. If we encode categories as colour (dotted line), we can decode different groups (solid line). If we encode numbers as position, we can also decode the size of differences and even average multiple positions. The blue column shows the unique advantages and disadvantages of encoding data values as the *shape* of text. If we encode categories as text shape, we can decode exact categories. If we encode numbers as text shape, we can decode exact numbers. However, we cannot (easily) decode the size of differences from text shape, nor can we decode averages. Finally, text shape is almost unique in its ability to encode and decode arbitrarily complex and abstract information.

The table in Figure 12b encodes counts as the horizontal position of text labels in the third column, as well as the shape of the text labels. This is not a typical format for a table, but we are more able to decode visual summaries, like the fact that three of the five largest counts are for dogs, while all other counts are fairly similar. In other words, there are more dogs than cats overall.

There are still clear rows of data values for each combination of country and pet type and we are still able to decode exact data values. However, now that the horizontal positions of the text labels for counts encode the counts, we are also able to decode data summaries from those positions, like the average number of pets for European countries—not by summarising the text *shapes*, but by summarising the text *positions*.

The table in Figure 12b is effectively a version of Figure 2b with text labels as data symbols instead of data points. Figure 13b makes the comparison even more explicit. Just as the table in Figure 12b is an unusual table format, Figure 13b is an unusual plot, but the framework that we have developed allows us to provide a reasoned justification for both of these visualisations.

7. Summary

We have developed a framework that treats text labels in a data visualisation in a similar way to the data symbols in a data visualisation. The framework is summarised in the bullet points below and in Figure 14.

- We can think of text as just another visual element that we can use to encode information, just like a line, or a bar, or a data point.
- We can encode information as the visual features of text, just as we can for other visual elements. For example, we can encode categorical data values as the colour of text or as the size of text.
- Visual features that are effective for other visual elements are also effective for text. For example, colour is effective for encoding a small number of categories and size is also effective for encoding quantitative data values.
- The most important visual feature for text is the text shape—the letters, words, and sentences that make up the text.
- Text shape is important because we can decode absolute data values from text shape, rather than just same/different, like colour, or larger/smaller, like size. In addition, the decoding of text shape is very accurate, more accurate than size, and it is possible to identify a very large number of different categories from different text shapes, many more than from different colours.
- Basic Gestalt principles of similarity and proximity allow information that is decoded from text labels to be associated with other data symbols like bars and points.
- Text shape is unique in being capable of encoding arbitrarily complex or abstract information, as well as very simple individual data values.
- On the other hand, it is very much harder to decode summaries of data values from text shape, something that is possible with, for example, the position of many data points.
- Furthermore, text shape is not congruent with data values. For example, a larger number, encoded as a text shape, does not naturally convey a larger value in the same way that a larger bar does.

This framework provides a reasoned justification for familiar uses of text in familiar plots: labels on axes and legends are the only way to convey absolute data values; titles and captions are the only way to encode complex metadata. In addition, the framework provides a basis for exploring new uses of text in less familiar plots, in particular, encoding simple data values as simple visual features of text *other than* shape.

8. Discussion

The framework for text that is outlined in this article is a descendant of the frameworks of Bertin (1983), Wilkinson (2005), and Munzner (2014). A data visualisation is considered to be an encoding of information as the visual features of graphical elements. Furthermore, there is a focus on what is the most effective encoding in terms of accurate and efficient decoding of information, following, for example, Cleveland (1994) and Ware (2020).

However, where text is included in those ancestral frameworks, it is treated as a separate case from other graphical elements, though with an acknowledged role as a tool for communicating complex and contextual information in, for example, titles and captions (Wanzer et al. 2021).

The main contribution of this framework is to observe that text has other roles to play within a data visualisation, including fundamental roles in the communication of much simpler information.

In this framework, we point out that text has many things in common with other graphical elements, like bars and lines and data points. We recognise that information can be encoded as visual features of text, such as colour, position, and size, just like for other graphical elements (Brath 2020).

In addition, encoding information as text *shape* provides the essential capability to decode exact data values. “La première propriété de la lettre c’est d’être non-ambiguë” (The first property of a letter is that it is unambiguous, Bertin 1967, Appendix A).

One limitation of this framework is that it is a simplification. While there are some established results regarding the effectiveness of decoding simple information from the visual features of text, more research is required (Brath and Banissi 2016). Furthermore, there are clearly some visual features that have a more complicated decoding, for example, the interaction between text font size versus text length (Alexander et al. 2018) and between text angle and text shape (Tinker 1956).

Another limitation of this framework is its focus on efficient decoding. We do not address the role of text on other measures of data visualisation effectiveness, such as memorability (Borkin et al. 2016).

Nevertheless, the hope is that this framework provides a deeper understanding that will lead to more deliberate, justifiable choices about the use of text in data visualisations. An empirical study to assess whether the framework does indeed have a positive impact is a possible avenue for future research.

Acknowledgments

Many thanks to the anonymous reviewers who provided very helpful suggestions for improvements to this article.

All data visualisations were generated with R using the following packages: **grid**, **ggplot2**, **xdvir**, **gridGeometry**, **ggtext**, **ggokabeito**, and **gggrid**.

The cat and dog face emojis used in Figure 8b are from the Noto Emoji Google Font.

References

Alexander, E. C., Chang, C.-C., Shimabukuro, M., Franconeri, S., Collins, C., and Gleicher, M. (2018). Perceptual biases in font size as a data encoding . *IEEE Transactions on Visualization & Computer Graphics*, 24(08):2397–2410, ISSN: 1941-0506, DOI: 10.1109/TVCG.2017.2723397.

Bertin, J. (1967). *Sémiologie graphique: Les diagrammes, les réseaux, les cartes.* Gauthier-Villars, Paris.

Bertin, J. (1983). *Semiology of Graphics: Diagrams, Networks, Maps.* University of Wisconsin Press, Madison, Wisconsin, ISBN: 9780299090609. Translated by William J. Berg from the 1967 French edition *Sémiologie Graphique.*

Borkin, M. A., Bylinskii, Z., Kim, N. W., Bainbridge, C. M., Yeh, C. S., Borkin, D., Pfister, H., and Oliva, A. (2016). Beyond memorability: Visualization recognition and recall. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):519–528, DOI: 10.1109/TVCG.2015.2467732.

Brath, R. (2020). *Visualizing with Text.* AK Peters Visualization Series. CRC Press, Boca Raton, Florida, ISBN: 9780367259266.

Brath, R. and Banissi, E. (2016). Using typography to expand the design space of data visualization. *She Ji: The Journal of Design, Economics, and Innovation*, 2(1):59–87, ISSN: 2405-8726, DOI: 10.1016/j.sheji.2016.05.003.

Brath, R. and Banissi, E. (2017). Font attributes enrich knowledge maps and information retrieval. *International Journal on Digital Libraries*, 18(1):5–24, ISSN: 1432-5012, DOI: 10.1007/s00799-016-0168-4.

Cairo, A. (2019). *How Charts Lie: Getting Smarter about Visual Information.* W. W. Norton & Company, New York, ISBN: 9781324001560.

Chernoff, H. (1973). The use of faces to represent points in k-dimensional space graphically. *Journal of the American Statistical Association*, 68(342):361–368. DOI: 10.2307/2284077.

Cleveland, W. S. (1994). *The Elements of Graphing Data.* Hobart Press, Summit, New Jersey, DOI: 10.2307/1422498.

Cleveland, W. S. and McGill, R. (1984). Graphical perception: Theory, experimentation, and application to the development of graphical methods. *Journal of the American Statistical Association*, 79(387):531–554, DOI: 10.2307/2288400.

Healey, C. (1996). Choosing effective colours for data visualization. In *Proceedings of 17th Annual IEEE Visualization '96*, pages 263–270. DOI: 10.1109/VISUAL.1996.568118.

Hearst, M. A. (2023). Show it or tell it? Text, visualization, and their combination. *Commun. ACM*, 66(10):6875, ISSN: 0001-0782, DOI: 10.1145/3593580.

Kong, H.-K., Liu, Z., and Karahalios, K. (2019). Trust and recall of information across varying degrees of title-visualization misalignment. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, CHI '19, page 1–13, New York, NY, USA. Association for Computing Machinery, ISBN: 9781450359702, DOI: 10.1145/3290605.3300576.

Kosslyn, S. M. (2006). *Graph Design for the Eye and Mind*. Oxford University Press, ISBN: 9780195311846, DOI: 10.1093/acprof:oso/9780195311846.001.0001.

Munzner, T. (2014). *Visualization Analysis and Design*. AK Peters Visualization Series. CRC Press, Boca Raton, FL, ISBN: 9781466508910, DOI: 10.1201/b17511.

Murrell, P. (2026). *How Data Visualisation Works*. The University of Auckland. Version 0.9 (build 2026-04-01).

Robbins, N. B. (2005). *Creating More Effective Graphs*. Wiley-Interscience, Hoboken, NJ, ISBN: 0-471-27402-X.

Schwabish, J. (2021). *Better Data Visualizations: A Guide for Scholars, Researchers, and Wonks*. Columbia University Press, New York, ISBN: 9780231193115.

Shah, P. and Hoeffner, J. (2002). Review of graph comprehension research: Implications for instruction. *Educational Psychology Review*, 14(1):47–69, ISSN: 1573-336X, DOI: 10.1023/A:1013180410169.

Sorapure, M. (2019). Text, image, data, interaction: Understanding information visualization. *Computers and Composition*, 54:102519, ISSN: 8755-4615, DOI: 10.1016/j.compcom.2019.102519.

Stokes, C. and Hearst, M. (2022). Why more text is (often) better: Themes from reader preferences for integration of charts and text. https://arxiv.org/abs/2209.10789.

Szafir, D. A., Haroz, S., Gleicher, M., and Franconeri, S. (2016). Four types of ensemble coding in data visualizations. *Journal of Vision*, 16(5):11–11, ISSN: 1534-7362, DOI: 10.1167/16.5.11.

Tinker, M. A. (1956). Effects of angular alignment upon readability of print. *Journal of Educational Psychology*, 47(6):358–363, ISSN: 0022-0663, DOI: 10.1037/h0044504.

Wagemans, J., Elder, J. H., Kubovy, M., Palmer, S. E., Peterson, M. A., Singh, M., and von der Heydt, R. (2012). A century of gestalt psychology in visual perception: I. Perceptual grouping and figure–ground organization. *Psychological Bulletin*, 138(6):1172–1217, DOI: 10.1037/a0029333.

Wanzer, D. L., Azzam, T., Jones, N. D., and Skousen, D. (2021). The role of titles in enhancing data visualization. *Evaluation and Program Planning*, 84:101896, ISSN: 0149-7189, DOI: 10.1016/j.evalprogplan.2020.101896.

Ware, C. (2020). *Information Visualization: Perception for Design*. Morgan Kaufmann, San Francisco, 4th edition, ISBN: 9780128128756.

Weber, W. (2019). Towards a semiotics of data visualization – an inventory of graphic resources. In *2019 23rd International Conference Information Visualisation (IV)*, pages 323–328. DOI: 10.1109/IV.2019.00061.

Whitney, D. and Yamanashi Leib, A. (2018). Ensemble perception. *Annual Review of Psychology*, 69:105–129, ISSN: 1545-2085, DOI: 10.1146/annurev-psych-010416-044232.

Wickham, H., Çetinkaya Rundel, M., and Grolemund, G. (2023). *R for Data Science*. O'Reilly Media, 2nd edition, https://r4ds.hadley.nz/.

Wilke, C. O. (2019). *Fundamentals of Data Visualization*. O'Reilly Media, https://clauswilke.com/dataviz/.

Wilkinson, L. (2005). *The Grammar of Graphics*. Statistics and Computing. Springer, New York, 2nd edition, ISBN: 9780387245447, DOI: 10.1007/0-387-28695-0.

Willey, C. R. and Liu, Z. (2022). Re-assessing the role of culture on the visual orientation perception of the rod and frame test. *PLOS ONE*, 17(10):1–15, DOI: 10.1371/journal.pone.0276393.

Zhang, J. (1996). A representational analysis of relational information displays. *International Journal of Human-Computer Studies*, 45(1):59–74, DOI: 10.1006/ijhc.1996.0042.

Affiliation:

Paul Murrell
The University of Auckland
Private Bag 92019
Auckland 1010
New Zealand
E-mail: paul@stat.auckland.ac.nz
URL: https://www.stat.auckland.ac.nz/~paul/

Journal of Data Science, Statistics, and Visualisation https://jdssv.org/
published by the International Association for Statistical Computing
http://iasc-isi.org/

ISSN 2773-0689
May 2026, Volume VI, Issue III *Submitted:* 2026-02-26
doi:10.52933/jdssv.v6i3.172 *Accepted:* 2026-04-29