Dishonest Chart Techniques

Lukas Auer, David Heidinger, Nina Tschikof, and Christina Vogel

Group 2
DAT.C406UF Information Visualisation 3VU SS 2026
Graz University of Technology

11 May 2026

Abstract

Data visualisations are widely used to communicate complex information in areas such as science, journalism, business, and politics. However, visualisations can also be designed in ways that distort the interpretation of data and mislead viewers. These dishonest chart techniques exploit perceptual and cognitive biases through manipulations such as selective data presentation, misleading axis scaling, distorted visual encodings, or concealed uncertainty.

This paper presents a survey of common dishonest chart techniques organised into a taxonomy of nine categories. To further support understanding, interactive demonstrations were developed to illustrate how visual manipulations can alter the perception of otherwise unchanged data. By increasing awareness of dishonest chart techniques, this work aims to support more critical interpretation and more transparent communication of information visualisation.

Contents

Contents i

List of Figures iii

1 Introduction 1

2 Dishonest Chart Techniques 3
2.1 Subjectivity and Bias in Data Representation . . . 3
2.2 Characteristics of Dishonest Chart Techniques . . . 3
2.3 Cognitive and Perceptual Effects . . . 4
2.4 Impact on Interpretation and Decision-Making . . . 4

3 Taxonomy of Dishonest Chart Techniques 5
3.1 Cherry Picking . . . 5
3.2 Underbinning . . . 6
3.3 Axis Truncation . . . 7
3.4 Artificial Flattening or Steepening . . . 7
3.5 Time Gap . . . 7
3.6 Oversmoothing . . . 8
3.7 Manipulating the Visual Encoding . . . 9
3.8 Perspective Distortion . . . 9
3.9 Concealing Uncertainty . . . 10

4 Interactive Illustrations 13
4.1 Implementation . . . 13
4.2 Tools and Technologies . . . 13
4.3 Challenges . . . 14

5 Concluding Remarks 15

Bibliography 17

List of Figures

3.1 Cherry Picking: Global Mean Surface Temperature 6
3.2 Underbinning: Heights of NBA Players 6
3.3 Axis Truncation: Fox News Tax Rates 7
3.4 Artificial Flattening or Steepening: Line Chart 8
3.5 Time Gap: Missing Years 8
3.6 Oversmoothing: Peaks and Dips are Reduced 9
3.7 Manipulating Size Coding: The Shrinking Family Doctor. 9
3.8 Manipulating Colour Coding: Sequential and Rainbow Palettes 10
3.9 Perspective Distortion: 3D Pie Chart 10
3.10 Concealing Uncertainty: Projected Revenues 11

Chapter 1

Introduction

Data visualisations are widely used to communicate complex information in an intuitive and efficient way. They support decision-making across domains such as science, journalism, business, or politics. However, the same properties that make visualisations powerful also make them susceptible to misuse.

Dishonest chart techniques refer to design choices or data manipulations which are deliberately chosen to distort a viewer's perception of the underlying data. These techniques may sometimes occur unintentionally as a result of poor design choices, but regardless of whether intentional or not, they can significantly influence how patterns, trends and relationships are interpreted.

This paper provides a taxonomy of nine common dishonest chart techniques, each accompanied by one or more interactive illustrations. By categorising these techniques, it becomes easier to identify patterns, understand how different forms of distortion operate, and what to look out for. Through increasing awareness of dishonest practices, both creators and consumers of visualisations can contribute to more accurate and transparent communication of information.

Chapter 2

Dishonest Chart Techniques

Data visualisation has become an essential tool for exploring, analysing, and communicating complex datasets to people [Tufte 2001; Cairo 2019]. By transforming numerical or categorical data into visual representations, it enables users to identify patterns, trends and relationships that would be difficult to detect through raw data alone. The effectiveness of data visualisation is largely based on the ability of the human visual perception system. Humans are particularly skilled at detecting differences in position, length, and colour. This allows them to quickly interpret visual encodings. Compared to textual or numerical representations, visualisations can significantly reduce cognitive load and improve comprehension.

As a result, visualisations are widely used across various domains including scientific research, journalism, media, and politics. In many of these areas visualisations are not only used to present data, but also to support decision-making processes or to support desired outcomes. However, this widespread use also increases the impact of potential misrepresentations. Since visualisations are often perceived as intuitive and objective, viewers may trust them without critically questioning how they were constructed [Cairo 2019].

2.1 Subjectivity and Bias in Data Representation

Despite their apparent objectivity, data visualisations are inherently subjective. Every step in the visualisation process involves decisions that influence how the data is ultimately presented. These decisions include data selection, filtering, aggregation, scaling and the choice of visual encodings [Cairo 2019].

Bias can be introduced at multiple stages of this process. During data collection, certain variables may be omitted or overrepresented. During preprocessing, the data may be cleaned or transformed in ways that affect its interpretation. Finally, during the design of the visualisation, certain choices such as axis scaling, colour mapping or chart type can further shape how the data is perceived by the audience. For example, selecting a specific time range may highlight or hide trends, while grouping data into categories may obscure underlying variability. Similarly, the use of certain colour schemes can emphasise or downplay specific values.

These examples illustrate that visualisations do not provide a neutral view of data. Instead, they represent a constructed perspective shaped by human decisions. Recognising this subjectivity is crucial for critically evaluating visualisations and avoiding misleading interpretations.

2.2 Characteristics of Dishonest Chart Techniques

Dishonest chart techniques can be defined as deliberate visualisation practices that lead viewers to a misinterpretation of the underlying data. These techniques do not necessarily involve incorrect data, but rather misleading representations of otherwise valid data [Cairo 2019; Tufte 2001]. Such techniques are sometimes used unintentionally, due to a lack of expertise in visualisation design, but here the focus is on their intentional use in order to influence the viewer's perception and support a particular narrative.

A key characteristic of dishonest chart techniques is that they exploit limitations in human perception and cognition. Viewers tend to assume that visual representations are proportional to the data they encode [Cleveland and McGill 1984]. At first glance, they may for example interpret the length of a bar or the area of a shape as directly corresponding to a numerical value, even when this relationship is manipulated. Common forms of distortion include exaggerating differences, compressing variability, hiding relevant context, or selectively presenting data. These distortions can significantly alter the perceived importance or meaning of the data. Importantly, dishonest techniques are often subtle and difficult to detect. This makes them particularly problematic as viewers may not be aware that their interpretation is being influenced.

2.3 Cognitive and Perceptual Effects

The effectiveness of dishonest chart techniques is closely related to how humans process visual information [Cleveland and McGill 1984]. Cognitive psychology has shown that individuals rely heavily on visual cues when interpreting data, often using heuristics rather than detailed analytical reasoning [Kahneman 2011]. For example, viewers may focus on prominent visual features such as large shapes, bright colours, or steep slopes, even if these features do not accurately reflect the underlying data. Similarly, they may overlook missing information, such as omitted baselines or hidden uncertainty.

Certain visual encodings are more prone to misinterpretation than others. While position and length are generally perceived more accurately, area and volume are often misjudged [Cleveland and McGill 1984]. This makes techniques that manipulate size or perspective particularly effective at misleading viewers. Additionally, cognitive biases such as confirmation bias can amplify the effects of dishonest visualisations. People may be more likely to accept misleading representations, if they align with their existing beliefs. These perceptual and cognitive factors explain why even small distortions in visualisation design can have a significant impact on interpretation.

2.4 Impact on Interpretation and Decision-Making

Misleading visualisations can have far-reaching consequences for interpretation and decision-making [Cairo 2019]. Since visualisations are frequently used to communicate important information, distortions in their design can influence how people understand and act on data.

In domains such as politics, finance, or healthcare, decisions are often based on visualised data. Misrepresentations in these contexts can lead to incorrect conclusions, poor decision-making, or the spread of misinformation. For example, exaggerated trends may create a false sense of urgency, while concealed variability may suggest a level of certainty that does not exist. Similarly, selective data presentation can support misleading narratives and influence public opinion.

The increasing use of data visualisations in digital media further amplifies these risks. Visualisations are often shared widely through media and consumed quickly, leaving little space for critical evaluation. Therefore, it is essential to improve awareness of how visualisations can mislead and to promote more responsible design practices.

Chapter 3

Taxonomy of Dishonest Chart Techniques

This chapter presents a taxonomy of dishonest chart techniques. It builds upon several pieces of previous work. Lan and Liu [2025] analysed 2227 flawed data visualisations from which they derived a taxonomy of 76 specific design flaws. These design flaws were categorised into three high-level categories and ten subcategories. The misinformation category includes flaws that distort or misrepresent information, such as truncated axes, misleading 3D effects, or incompatible chart types. Their taxonomy proposed by Lan and Liu is particularly relevant for this work, because it provides a broad and structured overview of misleading and flawed visualisation techniques.

The book by Cairo [2019] offers an in-depth analysis of how charts can be used for either revealing or concealing information. Multiple chapters explain the of different ways in which charts can be used to lie to the viewer. Cairo argues that charts are often perceived as objective and trustworthy by the general public, which makes misleading visualisations particularly influential in shaping opinions and public understanding. The early chapters helped identify common patterns of misleading chart design.

Yau [2025] presents interactive examples of misleading chart design and explains how visualisations can distort the interpretation of data. His use of interactive examples inspired the development of interactive charts to illustrate this survey.

UKPLab [2025] maintain a curated collection of papers, datasets, and online resources related to misleading data visualisations and chart understanding. This resource served as a useful starting point for finding literature and datasets.

The taxonomy presented here comprises nine broad categories: cherry picking, underbinning, axis truncation, artificial flattening or steepening, time gap, oversmoothing, manipulating the visual encoding, perspective distortion, and concealing uncertainty. Each of them is illustrated with at least one interactive example in the accompanying talk slide deck.

3.1 Cherry Picking

Cherry picking is a form of selection bias, where only a subset of available data is presented, while other relevant portions are left out. This practice is particularly problematic in time-series data, where the choice of start and end points can significantly influence the perceived direction and magnitude of a trend. By selecting intervals that support a specific narrative, the visualisation may suggest patterns that are not really representative of the full dataset [Cairo 2019]. For example, Figure 3.1 shows a very large time time frame of global mean surface temperatures over millions of years, which makes it impossible to see the effect of global warming over the most recent 100 years.

From a cognitive perspective, viewers tend to interpret the presented data as complete and representative, especially when no indication of omitted information is provided. This can lead to too much trust in the observed chart and reduces the likelihood that viewers will question the underlying data selection

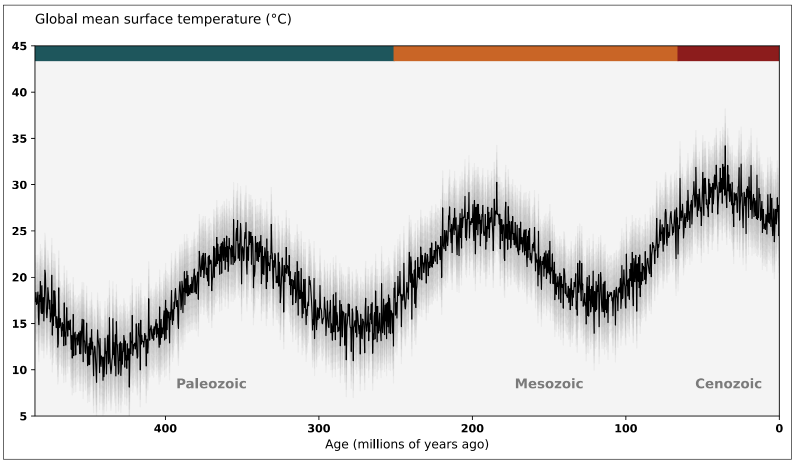

Figure 3.1: Cherry Picking: Global mean surface temperature over millions of years. It is hard to see the effect over global warming over the most recent 100 years. [Chart recreated in Python by the authors from the original by Judd et al. [2024].]

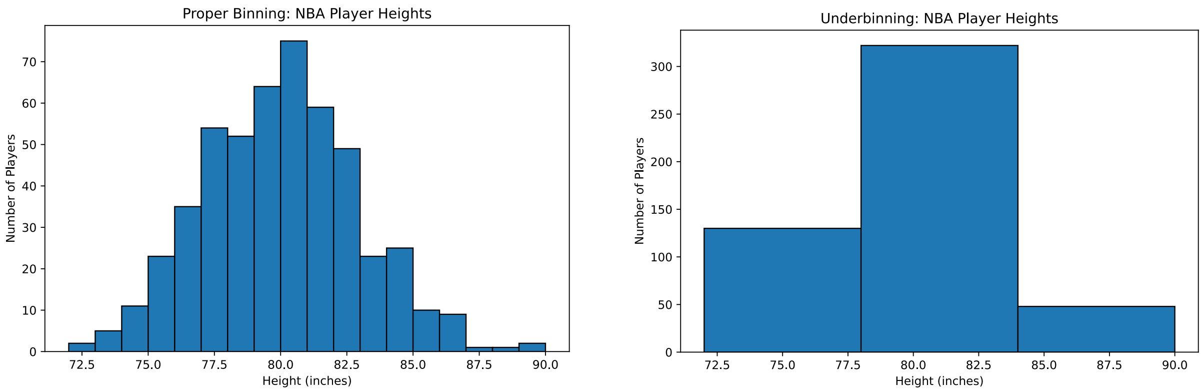

Figure 3.2: Underbinning: The heights of NBA players. On the left, 18 bins show a reasonable reflection of the underlying data. On the right, only 3 bins fail to show the nuances in the distribution. [Chart recreated in Python by the authors from the original by Yau [2025].]

process. As a result, cherry picking can effectively distort the interpretation without altering the actual data values, which makes this technique subtle but powerful.

3.2 Underbinning

Underbinning occurs when a chart, such as a histogram or a heatmap, uses too few bins to group the data. By dividing the dataset into overly broad intervals, the true shape of the distribution is hidden from the viewer. This reduction in detail obscures important features of the data, such as local clusters, gaps, and outliers. For example, if a dataset contains significant variations, but is grouped into only a small number of broad bins, the viewer will only see a very simplified picture. Figure 3.2 shows the heights of NBA players. Using only three bins fails to show the nuances of the distribution.

When viewing an underbinned chart, people often assume that the displayed blocks accurately reflect how the data is actually distributed. Since the visual representation looks complete and valid, viewers might not realise that crucial details have been smoothed over. This can lead to a misunderstanding of the actual situation, as the simpler chart might suggest a uniform trend where there is actually a lot of variation. Ultimately, while the raw data itself remains correct, the way it is grouped misleads the viewer by taking away the necessary context to understand the full dataset.

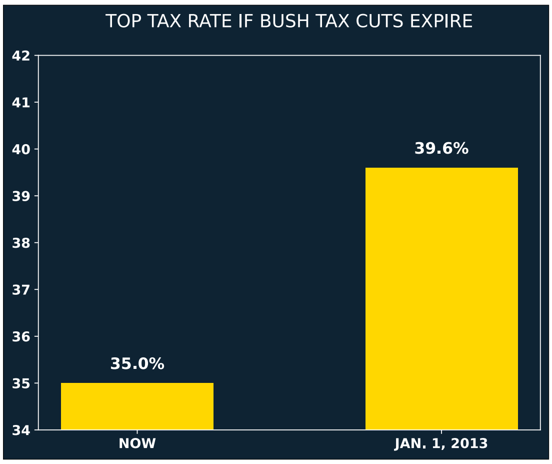

Figure 3.3: Axis Truncation: Starting the y-axis at 34 % rather than 0 artifically exaggerates the difference between the two values. [Chart recreated in Python by the authors from the original by [Fox 2012].]

3.3 Axis Truncation

Axis truncation is a technique in which the baseline of a chart, most commonly the y-axis, does not start at zero. Instead, the axis is cut off to begin at a value much closer to the actual data points. Since the bottom portion of the chart is removed, the remaining visual space is stretched out to fill the area. The primary effect of this manipulation is that it makes small differences between values appear much larger than they really are. Figure 3.3 shows two values of 35.0 % and 39.6 %, but the y-axis only starts at 34 %, artifically exaggerating the difference between the two values.

When people look at visualisations like bar charts, they naturally compare the physical heights of the bars to understand the data. If the zero baseline is missing, a minor increase might make one bar look twice as tall as another, even if the actual difference is very small. Although the numbers written on the axis are technically correct, the visual representation greatly exaggerates the situation and easily misleads the audience.

3.4 Artificial Flattening or Steepening

Artificial flattening or steepening happens when the visual dimensions of a chart are stretched or compressed to change how the data looks. This is often done by using extreme aspect ratios or by setting an overly broad y-axis. For example, if the y-axis covers a much larger range of numbers than the actual data points need, a moving trend line will look completely flat. On the other hand, stretching the chart vertically can make small, slow changes look like very steep and sudden spikes or drops. Figure 3.4 shows the same dataset with two different scales on the y-axis.

When people look at line charts, they naturally judge the importance of a trend by the steepness of the line. If a chart is artificially flattened, viewers might think a significant trend is just a minor detail and nothing to worry about. If the chart is artificially steepened, a normal change can quickly look like a dramatic event. The actual data points remain exactly the same in both cases, but manipulating the visual boundaries of the chart completely alters the message the audience takes away from it.

3.5 Time Gap

The time gap technique occurs when irregular or uneven time intervals are placed on a chart, but are laid out as if they were evenly spaced. For example, a chart might show data points for the years 2000, 2019, and 2020 right next to each other, with the exact same amount of physical space between them, as shown in Figure 3.5.

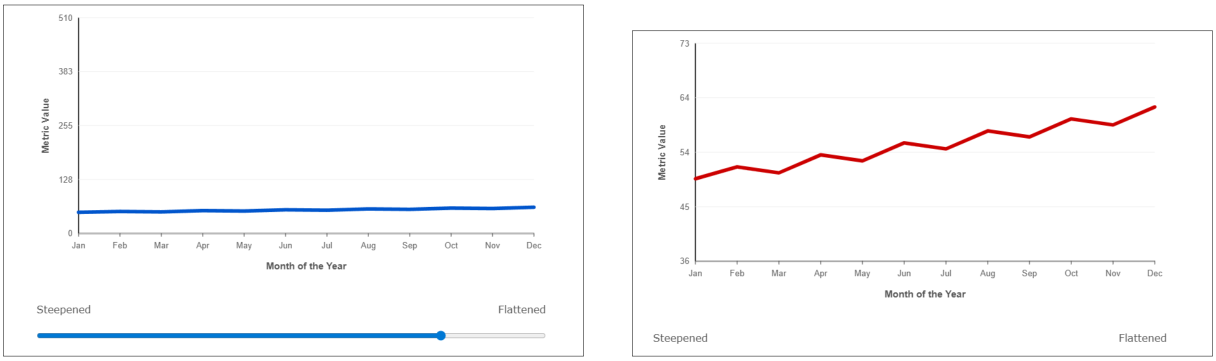

Figure 3.4: Artificial Flattening or Steepening: The line chart on the left has an overly large y-axis extent and looks artificially flat. The line chart on the right has a much narrower y-axis extent and looks steeper. [Image created by the authors from the accompanying interactive illustration.]

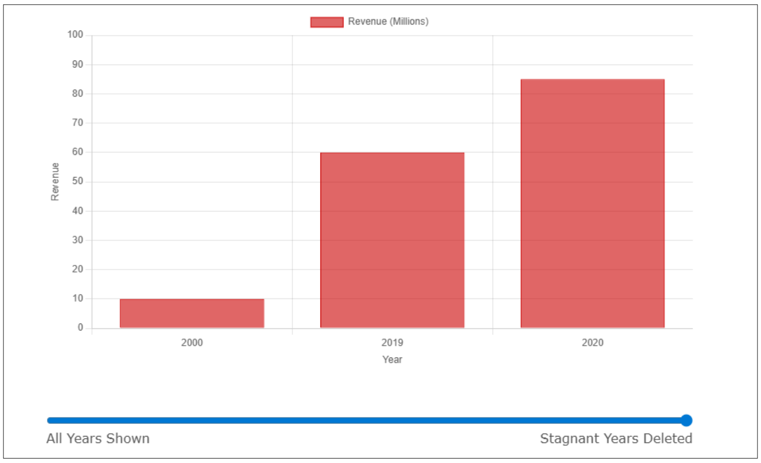

Figure 3.5: Time Gap: There are eight years missing between 2000 and 2019, but viewers may not notice this. [Image created by the authors from the accompanying interactive illustration.]

This manipulation effectively hides specific periods of time from the viewer. Usually, the years that are deleted are times when the data was stagnant or moving in an unwanted direction. When people look at a timeline or a trend chart, they naturally assume that time moves forward at a constant rate. Since the shown data points are spaced evenly, the trend looks very smooth and continuous. The viewer is misled into seeing a steady development that did not actually happen in reality, simply because the missing time gaps are not shown.

3.6 Oversmoothing

Oversmoothing occurs when aggressive smoothing algorithms or over-fitted curves are applied to a set of raw data points. Instead of showing the actual data with all its natural ups and downs, the chart draws a very smooth curve. This technique hides critical short-term fluctuations, volatility, and local extremes that are actually present in the dataset. Figure 3.6 shows an example of this.

When people look at an oversmoothed chart, they see a clear and stable trend. Since any sharp spikes and drops almost completely disappear, the viewer might think the situation is much more stable than it really is. While some smoothing can be a helpful tool to understand long-term developments, using it too heavily takes away important local patterns. This easily misleads the audience by making a highly variable and unpredictable dataset look perfectly calm and steady.

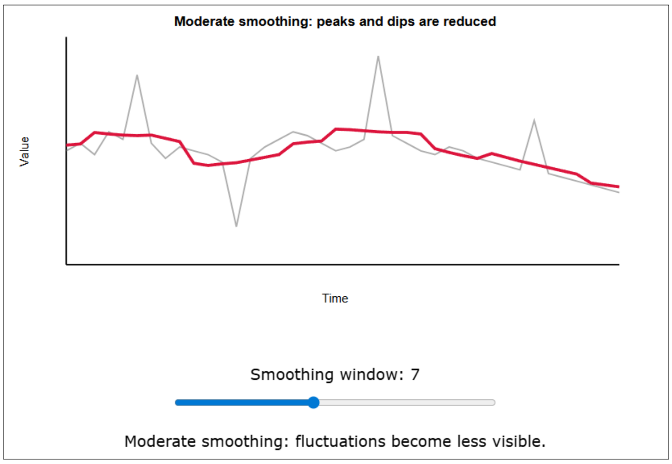

Figure 3.6: Oversmoothing: When smooting is over-applied, fluctuations become less visible. [Image created by the authors from the accompanying interactive illustration.]

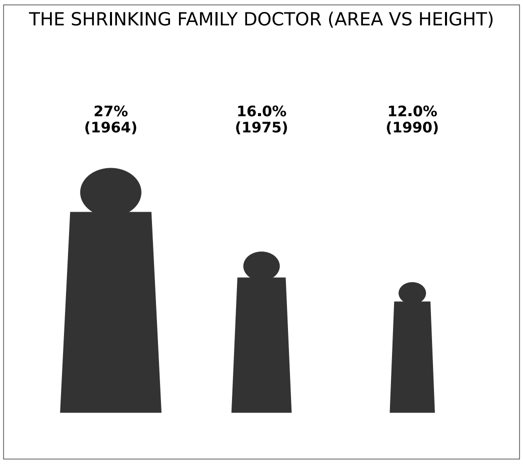

Figure 3.7: Manipulating Size Coding: The data value is actually mapped to the height of the icon, but the human eye sees the difference in area, making 27 % seem many times larger than 12 %. [Chart recreated in Python by the authors based on the original LA Times graphic [LA Times 1979].]

3.7 Manipulating the Visual Encoding

Manipulation of visual encoding occurs when the choice of how to show data, such as using size, shape, or colour, is used to mislead the viewer. Instead of presenting the numbers clearly, the design uses confusing or exaggerated visual cues. For example, a chart might map a simple one-dimensional number to the two-dimensional area of an icon or glyph. If the number drops by half, the overall area of the drawing is reduced to a quarter of its original size, leaving the viewer to assume that the drop is much more dramatic than it actually is. This can be seen in Figure 3.7

Similarly, using a rainbow colour palette for simple continuous data can create false boundaries and patterns that do not exist in reality, as can be seen in Figure 3.8. This easily changes how the audience interprets the data without altering the actual values.

3.8 Perspective Distortion

Perspective distortion occurs when three-dimensional effects are improperly used in a two-dimensional visualisation.

For example, by tilting a pie chart backwards, the parts of the chart located in the front appear larger

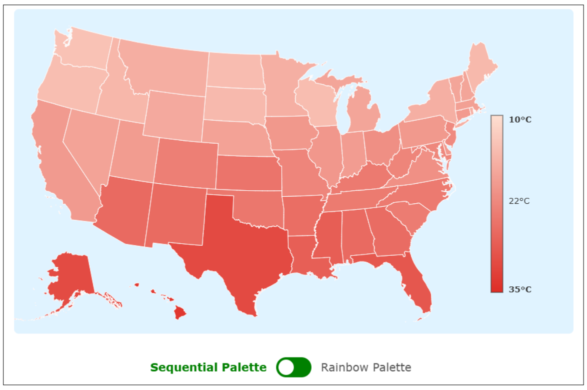

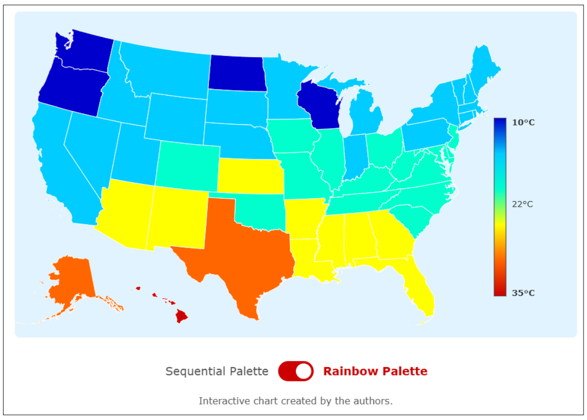

Figure 3.8: Manipulating Colour Coding: Using a rainbow palette on the right for continuous data can create false boundaries and patterns. [Image created by the authors from the accompanying interactive illustration.]

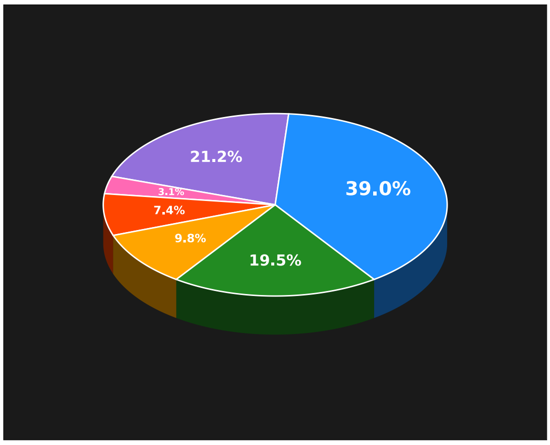

Figure 3.9: Perspective Distortion: The 19.5 % slice at the bottom of the 3d pie chart looks larger than the 21.2 % slice at the top. [Chart recreated in Python by the authors, based on Steve Jobs' original talk slide from his 2008 Apple Macworld Keynote [Paragraft 2008].]

to the viewer, whereas the parts in the back seem smaller. The tilt creates a false sense of size and importance. For example, in Figure 3.9, the 19.5 % slice placed at bottom of a tilted pie chart looks larger than the 21.2 % slice at the top, simply because of 3d perspective. When people look at a 3D chart, it becomes very difficult for them to compare the true angles and areas accurately. Ultimately, adding an unnecessary third dimension does not provide any new information, but does heavily distort the true proportions and misleads the audience.

3.9 Concealing Uncertainty

Concealing uncertainty happens when a chart does not include important visual indicators of variance, such as error bars or confidence intervals. Instead of showing the possible range of a value, the chart only displays a single, exact-looking line or data point. For example, Figure 3.10 shows a line chart of projected revenues with a 90 % confidence interval to explicitly indicate the increasing uncertainty in the projections.

This technique suggests absolute precision in data that is actually just an estimate or a forecast. When people look at a projected trend without any signs of uncertainty, they naturally assume the future outcome is completely guaranteed. Since the potential error or variance is hidden from the screen, the viewer is

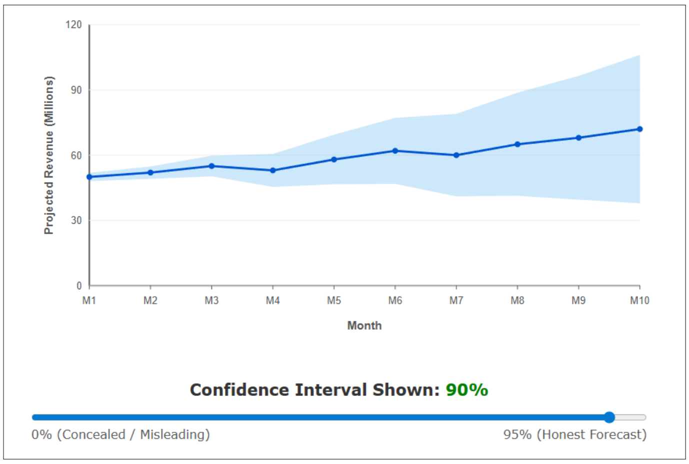

Figure 3.10: Concealing Uncertainty: Uncertainty in data points should be clearly indicated. Here, a confidence interval of 90 % shows the increasing uncertainty in projected revenues. [Image created by the authors from the accompanying interactive illustration.]

misled into trusting the projected numbers too much. In reality, no forecast is perfect, but removing the visual boundaries of doubt creates a false sense of security and accuracy in the data.

Chapter 4

Interactive Illustrations

Interactive illustrations were created in JavaScript for each of the nine demonstrationsdishonest chart techniques. The primary design goal for the illustrations was to make the abstract concepts of chart manipulation tangible and easy to understand. Instead of just showing static "before and after" pictures, the illustrations were designed to let users experience the distortion firsthand.

By using interactive elements like sliders and toggle switches, the viewer can actively change the parameters of a chart, such as the number of bins, the y-axis starting point, or the scale of the axes. This hands-on approach allows the audience to see in real-time how a perfectly normal dataset can slowly transform into a highly misleading visualisation.

4.1 Implementation

From a technical perspective, all the interactive illustrations follow a similar internal structure. Each demo consists of a visual display area and a set of user controls. When a user interacts with a control, such as dragging a slider, an event listener immediately triggers a JavaScript function.

This function recalculates the necessary variables and updates the visual elements on the screen dynamically. For example, in the axis truncation demo, moving the slider recalculates the visible height of the bars relative to the new y-axis minimum, and updates the SVG rectangle heights instantly. In the cherry-picking demo, the slider calculates a smooth interpolation between the full dataset view and a zoomed-in timeframe, updating an SVG clipping mask to hide the rest of the timeline. This immediate feedback loop is what makes the illustrations feel responsive and fluid.

4.2 Tools and Technologies

To build the interactive illustrations, a combination of standard web technologies and specific data visualisation libraries were used. The foundation of the project is written in standard HTML, CSS, and JavaScript. This allowed us to easily integrate the interactive illustrations directly into our slide deck without needing complex software. For the presentation itself, we used the Rslidy framework [Andrews et al. 2026], which is a lightweight tool for creating HTML-based presentations.

The charts themselves were built with Scalable Vector Graphics (SVG). By manipulating SVG elements with pure JavaScript, completely custom animations and interactive features could be created. For some of the more complex visualisations, external JavaScript libraries were used to save time and ensure accuracy. Chart.js was used to build the bar chart for the time gap demo, because it provides smooth, built-in animations that made it easy to hide and show specific years. For the colour mapping demo, D3.js aand TopoJSON were used to render an interactive map of the United States. This allowed dynamic switching between a sequential colour palette and a rainbow palette.

4.3 Challenges

Developing the interactive illustrations presented several technical and design challenges. One of the primary technical hurdles was the mathematical complexity involved in rendering custom SVGs from scratch. For instance, the perspective distortion demo required complex trigonometry to calculate the correct arcs, ellipses, and depth lines to simulate a tilted three-dimensional pie chart without relying on a dedicated 3D library.

Another significant challenge was designing smooth transitions between the “honest” and “dishonest” states. In the cherry-picking and artificial steepening demos, the transition needed to be fluid, so the viewer could visually track exactly how the data representation was being altered. Calculating the intermediate steps for the axis scaling, bounding boxes, and data point positions required careful tuning of the JavaScript logic to ensure the animations felt natural.

Finally, ensuring that the visualisations remained responsive and legible across different screen sizes required extensive CSS styling and the careful use of the SVG `viewBox` attribute. Balancing the interactive controls, such as the sliders and toggle switches, with the visual output to create an intuitive and educational user experience took several iterations to get right.

Chapter 5

Concluding Remarks

Data visualisations are powerful tools for communicating complex information, but they can also be used in misleading ways. This paper examined a range of dishonest chart techniques that distort the interpretation of data. Although these techniques often preserve the underlying data values, they can significantly alter how viewers perceive patterns, trends, and relationships.

To better understand these forms of distortion, this paper proposed a taxonomy consisting of nine common dishonest chart techniques: cherry picking, underbinning, axis truncation, artificial flattening or steepening, time gap, oversmoothing, manipulating the visual encoding, perspective distortion, and concealing uncertainty. To support practical understanding, interactive illustrations were developed that allow users to directly experience how visual manipulations influence perception.

Overall, this work aims to contribute to a more critical understanding of data visualisations and to promote more transparent and responsible visual communication. As visualisations continue to shape public opinion and decision-making, the ability to recognise dishonest chart techniques is becoming increasingly important for both creators and consumers of visual information.

Bibliography

Andrews, Keith, Fabian Platzer and Patrick Hipp [2026]. *Rslidy: Lightweight, Accessible, and Responsive HTML5 Slide Decks*. 07 May 2026. https://github.com/tugraz-isds/rslidy (cited on page 13).

Cairo, Alberto [2019]. *How Charts Lie: Getting Smarter about Visual Information*. W. W. Norton, 15 Oct 2019. 256 pages. ISBN 1324001569 (cited on pages 3–5).

Cleveland, William S. and Robert McGill [1984]. *Graphical Perception: Theory, Experimentation, and Application to the Development of Graphical Methods*. Journal of the American Statistical Association 79.387 (Sep 1984), pages 531–554. doi:10.1080/01621459.1984.10478080. https://jstor.org/stable/2288400 (cited on page 4).

Fox [2012]. *If Bush Tax Cuts Expire*. Fox News, Aug 2012. https://peoplesgdarchive.org/item/10481/if-bush-tax-cuts-expire-fox-news-infographic (cited on page 7).

Judd, Emily J., Jessica E. Tierney, Daniel J. Lunt, Isabel P. Montañez, Brian T. Huber, Scott L. Wing and Paul J. Valdes [2024]. *A 485-Million-Year History of Earth's Surface Temperature*. Science 385.6715 (20 Sep 2024). doi:10.1126/science.adk3705. https://henry.pha.jhu.edu/SCIENCE.pdf (cited on page 6).

Kahneman, Daniel [2011]. *Thinking, Fast and Slow*. Farrar, Straus and Giroux, 25 Oct 2011. 256 pages. ISBN 0374275637 (cited on page 4).

LA Times [1979]. *The Shrinking Family Doctor*. 05 Aug 1979. https://infovis-wiki.net/wiki/Lie_Factor (cited on page 9).

Lan, Xingyu and Yu Liu [2025]. *"I Came Across a Junk": Understanding Design Flaws of Data Visualization from the Public's Perspective*. IEEE Transactions on Visualization and Computer Graphics 31.1 (Jan 2025), pages 393–403. doi:10.1109/TVCG.2024.3456341. https://arxiv.org/abs/2407.11497 (cited on page 5).

Paragraft [2008]. *The Chart-Junk of Steve Jobs*. 03 Jun 2008. https://paragraft.wordpress.com/2008/06/03/the-chart-junk-of-steve-jobs/ (cited on page 10).

Tufte, Edward R. [2001]. *The Visual Display of Quantitative Information*. 2nd Edition. Graphics Press, Jan 2001. 200 pages. ISBN 0961392142 (cited on page 3).

UKPLab [2025]. *Awesome Misleading Visualizations*. Ubiquitous Knowledge Processing Lab, 01 Sep 2025. https://github.com/UKPLab/awesome-misleading-visualizations (cited on page 5).

Yau, Nathan [2025]. *Defense Against Dishonest Charts*. 13 Feb 2025. https://flowingdata.com/projects/dishonest-charts/ (cited on pages 5–6).