

Andy Kirk

# Data Visualisation

A Handbook for Data  
Driven Design

3rd Edition

 **Sage**



1 Oliver's Yard  
55 City Road  
London EC1Y 1SP

2455 Teller Road  
Thousand Oaks  
California 91320

Unit No 323-333, Third Floor, F-Block  
International Trade Tower  
Nehru Place, New Delhi – 110 019

8 Marina View Suite 43-053  
Asia Square Tower 1  
Singapore 018960

---

Editor: Kate Keers  
Editorial assistant: Becky Oliver  
Production editor: Ian Antcliff  
Copyeditor: Richard Walshe  
Proofreader: Thea Watson  
Marketing manager: Ruslana Khatagova  
Cover design: Shaun Mercier  
Typeset by: C&M Digital (P) Ltd, Chennai, India  
Printed in the UK

© 2025 Andy Kirk

Apart from any fair dealing for the purposes of research, private study, or criticism or review, as permitted under the Copyright, Designs and Patents Act, 1988, this publication may not be reproduced, stored or transmitted in any form, or by any means, without the prior permission in writing of the publisher, or in the case of reprographic reproduction, in accordance with the terms of licences issued by the Copyright Licensing Agency. Enquiries concerning reproduction outside those terms should be sent to the publisher.

**Library of Congress Control Number: 2024939313**

**British Library Cataloguing in Publication data**

A catalogue record for this book is available from the British Library

ISBN 978-1-5296-2655-1  
ISBN 978-1-5296-2654-4 (pbk)

# Contents

Acknowledgements	vii
About the Author	ix
Discover Your Textbook's Online Resources	xi
Introduction	1
<b>Part A The Foundations</b>	<b>13</b>
1 What is Data Visualisation?	15
2 The Process and Principles of Visualising Data	29
<b>Part B The Hidden Thinking</b>	<b>61</b>
3 Formulating the Brief	63
4 Working with Data	95
5 Editorial Thinking	123
<b>Part C The Design Thinking</b>	<b>143</b>
6 Data Representation	145
7 Interaction and Dynamic Design	205
8 Annotation	231
9 Colour	253
10 Composition	273
References	293
Index	295



# Acknowledgements

Having been an enthusiastic participant in the data visualisation community since 2007, my growth and development in this field since then has been hugely influenced by the contributions of thousands of data viz designers and developers, visual journalists, academics, professionals, studios and freelancers. I want to acknowledge how much their inspirational works and smart discourse has influenced the thinking I've expressed in this book. Along the same timeline, my gratitude also goes to the individuals and organisations who have invested in my professional services, giving me the opportunity to do what I love for a career. Thank you, as always, to the many people at Sage, past and present, who have helped this book grow from a first proposal, through first, then second and now onto a third edition. To the readers, I don't wish to delay you getting into the contents, but thank you for choosing to invest in this book, through your outlay and your time working through it. I sincerely hope it elevates your capability, confidence and conviction in this subject, one which continues to stir my excitement and motivation to learn every day. My final special thanks, of course, go to my family for their love and support, and to Ellie, especially, for her understanding across the many occasions I've had cause to disappear in the evening and late into the night working on this book. I'll dedicate this edition, though, to the next generation: my nieces and nephews. The careers you end up in may not yet exist, but whatever path you decide to take in the future, I know you'll each make a huge splash. Keep being smart, keep working hard and keep being kind. You'll do just fine.



# About the Author

Andy Kirk is an experienced independent data visualisation expert, delivering professional services to clients worldwide. Andy works as a data visualisation design consultant, educator, author, researcher and speaker. He is the editor of *visualisingdata.com* and hosts the 'Explore Explain' video and podcast series.

After graduating from Lancaster University in 1999 with a BSc (hons) in Operational Research, Andy's working career commenced in a series of business analysis/information management roles with organisations including CIS Insurance, West Yorkshire Police and the University of Leeds.

He discovered data visualisation in 2007, during an era when it was yet to leap into the mainstream consciousness. This discovery was timed to perfection as it became the topic of the proposal for his Master's (MA) degree, through a self-directed research programme at the University of Leeds. Studying the subject over the subsequent two years unleashed and secured his passion for data visualisation, leading to the launching of his website in 2009 to continue the process of learning.

He became a freelance professional in 2010, since when Andy has had the privilege of working with a diverse set of clients, covering all shapes, sizes, domains, industries and locations across the world, helping them enhance their data visualisation capabilities. At the time of this edition's publication, he will have delivered over 400 public and private data viz training events in nearly 30 different countries, attended by over 8000 people.

Alongside his commercial activities, Andy has maintained academic duties in both research and teaching. In the mid-2010s he was an external consultant on a research project called 'Seeing Data', funded by the Arts & Humanities Research Council and hosted by the University of Sheffield. This study explored the issues of data visualisation literacy among the general public and, inter alia, helped to shape an understanding of the human factors that affect visualisation literacy and the effectiveness of design. During the 2010s he delivered postgraduate teaching at the Maryland Institute College of Art (MICA, US) and Imperial College London (UK) and, most recently, was an adjunct lecturer at University College London (UCL), co-teaching a data visualisation module on the Business Analytics Master's programme in the School of Management.



# Discover Your Textbook's Online Resources

Want more sources of reference and inspiration to further your understanding about data visualisation? Hosted and curated by the author, the supporting website for this book, at [book.visualisingdata.com](http://book.visualisingdata.com), has everything you need to continue exploring, practising and honing your data visualisation skills.

- **Explore the field:** Expand your knowledge and reinforce your learning with further reading, example references and learning tutorials organised by chapter.
- **Try for yourself:** Challenge, reflect and refine your skills by working on some practical exercises and activities linked to each chapter's topics.
- **Learn from the best:** Watch or listen to the episodes from the 'Explore Explain' podcast and video series, as Andy explores the visualisation design process followed by guest designers, developers and visual journalists. Hear from stellar talents from across the world as they explain the stories behind the inspirational works they've created. You'll get insight into how today's professionals navigate through all stages of contextual, analytical, editorial and creative activity outlined through this book.
- **Get technical:** See the latest catalogue of contemporary data visualisation tools, applications and libraries. 'Which tools make which charts?': through 'The Chartmaker Directory' you'll find a huge – and growing – collection of answers to that perennial question, with nearly 1500 examples and technical tutorials to enhance your chart-building capabilities.



# 3

## Formulating the Brief

In Chapter 2 we learned about the importance of adopting a design process to tackle data visualisation challenges through an organised sequence of activities. Supplemented by the guiding design principles, this offers a framework to help you make good decisions.

In this third chapter we commence the process with activities concerned with ‘formulating the brief’, defining the drivers behind and the factors shaping your work. Through this first stage we will seek to forge clarity about the context surrounding the project, the intended purpose it is serving, and we’ll conceive ideas about our initial creative vision.

No visualisation project is ever undertaken free of certain restrictions and frictions. You will spend time in this opening stage defining the contextual matters that give shape to much of the pragmatic thinking that runs through your work. You’ll think about the various people involved in this communication exchange and the characteristics of the deliverables expected and identify the constraints that will affect your choices.

Creating an actual, formal *brief* isn’t necessary, it’s the thinking this all entails that’s necessary. You can choose to be informal or formal, depending on what context dictates. For example, it can be useful when working with other people (clients, colleagues, collaborators) to document all the requirements and use it to inform the scope of a project. This can be shared, agreed upon and subsequently referred to. It will be in the interests of all parties to have such a document, especially for ‘mutually’ clarifying the expected deliverables. When working on more personal, perhaps self-initiated, projects you might choose to just keep basic notes that capture your thinking throughout the process, as much for your own record as anyone else’s.

A critical part of this opening stage is to establish clarity, as early as possible, about why you are producing this particular data visualisation. What is its *raison d’être*? As we’ve learned, the aim of a visualisation is to facilitate understanding. This process will therefore be driven by *what* specific understanding it is that you are seeking to facilitate for your audience. It’s not always known or knowable just yet, but you can initiate this thinking now, then define and refine later.

In defining your project’s purpose, you will also benefit from thinking already about how you might accomplish your aims: *how* will you facilitate understanding? We’re not yet concerned with thinking about design in any detail, but there is a need to identify your direction of travel in terms of communication strategy. Are you looking to convey and explain key messages, maybe via a visual story, showing and telling your audience what’s significant? Or maybe you are going to enable them to find and form their own discoveries? Is the tone of voice of your communication going to demonstrate more of an analytical emphasis, build around precise, quantitative perceptions, or will it need to focus more on heightening the gist of the data, amplifying the feeling of a subject?

We’ll close the chapter by allocating time to the *vision* you might have for your work. What early conceptual thoughts come to mind about what it is you think you might be creating? What

metaphors, colours and forms are evoked by your subject? What other sources of inspiration might you seek to emulate or draw ideas from? Capture these, sketch them, note down every idea. It will be helpful for later!

## 3.1 Defining Your Project's Context

Contextual thinking recognises the circumstances surrounding and affecting the visualisation project you are undertaking. These are the different points of friction or freedoms that will influence the boundaries of your creative ambitions. Some are imposed *on* you or determined *by* you.

At the start of any visualisation project, it will be common for not all the circumstances we're about to list to yet be definable or even be relevant to your work. Things change. That's why we need to

'Context is key. You'll hear that the most important quality of a visualisation is graphical honesty, or storytelling value, or facilitation of "insights". The truth is, all of these things (and others) are the most important quality, but in different times and places. There is no singular function of visualisation; what's important shifts with the constraints of your audience, goals, tools, expertise, and data and time available.' **Scott Murray, Author and Designer**

follow a process, to adapt elegantly to accommodate the impact of any new conditions emerging at any point in our workflow. Of course, the more things you *can* define now, the more the whole picture will become clearer, sooner. We want to eliminate as much of the unknown as we can, even if that is just knowing that a potentially influential factor doesn't have any fixed requirements. Sometimes you might see merit in imposing restrictions yourself, if none are imposed by others, to reduce the paradox of choice that can creep into creative processes. Restrictions are not necessarily a bad thing. Indeed, they often lead us to innovate, as discussed in the latter part of the previous chapter.

## What Contextual Factors will Shape Your Work?

There are several factors that we must consider relating to matters of people, of deliverables, and of constraints.

### People

**Client(s):** In those situations where you have been asked to develop a visualisation by somebody else, it is important to keep them in mind. For the context of this book, we'll consider clients to be anyone who has instigated or commissioned you to create a visualisation, which may include colleagues, managers or external customers. Inevitably, they will be invested in and have an influence on what you develop, but the extent of their stake will vary and your experience working under their influence will vary, accordingly. Some clients will seek constant involvement in the progress and direction of your work, steering you towards their needs. Managing such a relationship can be delicate, because there's a fine line between receiving useful guidance and experiencing distracting interference. The client may seek frequent updates on progress, have key decisions run past them for their approval and eventually determine the judgement about whether a solution is of sufficient quality to be considered finished. In my world as a freelance design consultant, my clients determine when (and if) I will get paid.

**Creator:** This is you, performing the role of the visualiser, leading this process towards a developed final solution. Data visualisation brings together a wide variety of capabilities and sensibilities that

fuel its richness, but can also make it seem overwhelming to master. The seven hats of data visualisation (Figure 3.1) represent an attempt to recognise the wide repertoire of specific skills, areas of knowledge, behaviours or attitudes that a visualiser may need to demonstrate across different stages of the process. Inspired by Edward de Bono's *Six Thinking Hats* (1985), it profiles the specification of a perfect 'unicorn' visualiser. The attributes listed under each of these hats can be viewed as a wish list of personal or team capabilities, depending on the context of your data visualisation work.



**Figure 3.1** The Seven Hats of Visualisation Design

As you browse across these different capability profiles, which attributes do you possess or can you confidently demonstrate? Where are your gaps? Where are your 'good enoughs'? Which are the things you enjoy doing and the things you dislike? I know the things I am not skilled at (programming), the things I am good enough at (graphic design, where I rely more on instinct than skill gained from training), and the things I do not enjoy or find effortful

'There is not one project I have been involved in that I would execute exactly the same way second time around. I could conceivably pick any of them – and probably the thing they could all benefit most from? More inter-disciplinary expertise.'

**Alan Smith OBE, Head of Data and Visual Journalism at *Financial Times***

(proofreading, note-taking). Having that self-awareness at least means I know the current boundaries of my capabilities, where I need to improve and what I might need compromise on.

**Collaborator(s):** Being able to pool a diverse but balanced profile of different capabilities, to tackle shared challenges, is often the secret behind the best visualisations. Building on the examination of one's own skills, if you recognise gaps in your repertoire of capabilities, do you have the opportunity to work with others who bring complementary talents?

Other valuable collaborators will be subject matter experts, from whom you can obtain domain-specific guidance about matters that exceed your knowledge. They might be able to inform you about the most salient issues for topics you know little about, what certain terminologies mean or whether a pattern emerging in some analysis is a genuinely significant discovery or just noise. They might also offer technical support around specific application or technological issues.

'One thing I will stress is how important it is to actually talk to people on the phone, not just emailing people. I feel like when you talk to someone on the phone they will open up a little bit more, they'll give you actual insights that maybe you wouldn't have picked up on that can then enhance the project or give you an entirely new idea for a follow up piece.' **Simon Scarr, Deputy Head of Graphics, Thomson Reuters**

Identifying this extended cast of potential collaborators at the start of your process means you can plan for and anticipate the relationships you might need to manage: when is everyone going to be available, what personality traits might need handling, what obstacles might need navigating like work conflicts?

If the project is effectively a solo pursuit, you will enjoy greater autonomy, which can be liberating, but with this comes more responsibility on you to deliver the goods.

**Audience:** How to usefully classify the characteristics of your intended or potential audience?

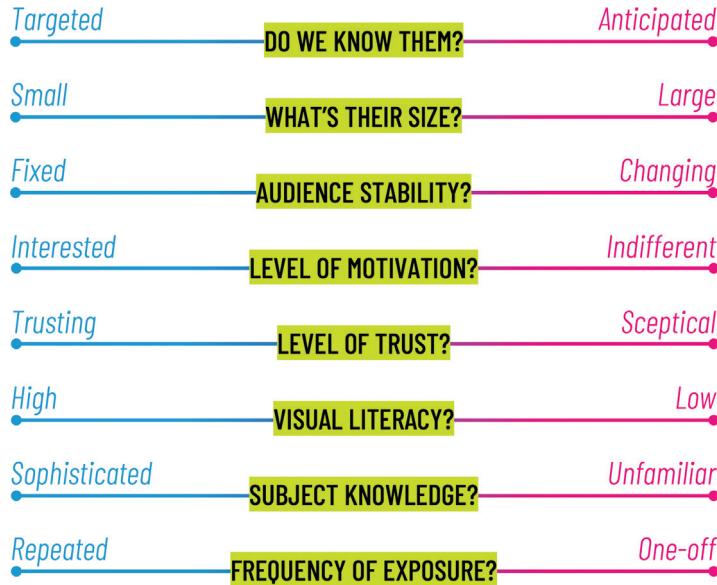
You are creating your work for other people so this is a question that will remain in your mind throughout this process. Answering this question, to the extent you can, needs tackling from the start.

The graphic in Figure 3.2 lists some of the most relevant audience classifications that will influence your thinking throughout this process and towards the eventual solution. The spectrum of attributes presented from left to right (blue to pink) displays the range of possible classifications for each characteristic. The blue classifications generally mean your visualisation process will have more certainty, more clarity and fewer obstacles to deal with. The colours are not tied together, though, in the sense that, for instance, if your relationship with the audience is more 'anticipated' rather than 'targeted', that does not imply all the other pink classifications are relevant.

Let me elaborate on a few of these characteristics. When you have direct knowledge of who your audience is – indeed, you may even have direct access to them – it is likely clearer what their needs are. Your choices can be moulded accordingly. If it is more unpredictable who will encounter your work, you're going to have to make a lot of assumptions to address that ambiguity.

If you have an especially large, diverse and changing profile of potential viewers, you are unlikely to be able to satisfy the needs of each cohort group within that mass; to cope you might need to formulate some imagined personas of the main two or three types of people you at least want this to work for, prioritising them over all other potential audience type variations. One size rarely fits all.

If you have a fixed group of viewers, let's say people who always browse through your weekly reports, you should be confident they will understand the technical context of the data you are presenting. You will therefore not need to include the type of detailed explanations that otherwise would be necessary for a less familiar audience.



**Figure 3.2** Classifying Audience Characteristics

In larger audience situations, it's reasonable to expect a varied range of motivation towards the subject you're presenting and the understanding you intend to facilitate. For those who are likely to be somewhat indifferent to your content, you may need to work harder to convince them about its potential relevance. Maybe an immediate headline to hook them in or a novel visual flourish to appeal to their tastes?

The visual literacy capabilities amongst your viewers will inform the need, or otherwise, for offering assistance when it comes to perceiving different chart types. Additionally, if there are any viewers with known or potential accessibility impairments, as covered in the previous chapter, their needs must be factored into your design thinking. There may also be language obstacles, potentially requiring you to create alternative translated or at least captioned versions of your eventual solution.

## Deliverables

**Format:** You will need clarity about the specific format of the deliverables you are intending or are required to produce. Is your work going to be published in print, for a digital platform, or even taking some alternative physical form? Will it be static, interactive or animated?

Maybe you will have to produce something to be published across multiple media. For example, newspapers typically publish their graphics in print, for a website, maybe for a mobile app, and also often share them across social media. While it may be essentially the same graphic repurposed to be published across each format or platform, there may still be subtle alterations required to optimise the presentation for each respective version. This increases your workload.

**Quantity:** As well as the different output formats, it's also important to get a sense of the different dimensions of your project's deliverables. How *many* charts are you making? What shape and what size? For example, are you producing 12 different graphics for a month-by-month slide deck or contributing to a large report that will need two charts developed for each of the

'I love, love, love print. I feel there is something so special about having the texture and weight of paper be the canvas of the visualisation. It's a privilege to be able to design for print these days, so take advantage of the strengths that paper offers – mainly, resolution and texture. Print has a lot more real estate than screen, allowing for very dense, information packed visualisations. I love to take this opportunity to build in multiple story strands, and let the reader explore on their own. The texture of paper can also play a role in enhancing the visualisation; consider how design and colour choices might be different on a glossy magazine page vs the rougher surface of a newspaper.' **Jane Pong, Graphics Editor at Bloomberg News**

'I like to imagine that I have a person sitting in front of me, and I need to explain something interesting or important about this data to them, and I've only got about 10 seconds to do it. What can I say, or show them, that will keep them from standing up and walking away?' **Bill Rapp, Art Director at Center for American Progress**

50 questions asked in a survey? Perhaps it's a web-based project with four distinct sections, each requiring at least one interactively-adjustable view, or it could be just a single chart to be emailed to your manager. It is not always possible to determine output dimensions this early in the process, but you should certainly maintain awareness of how realistic the expected deliverables are going to be, given your project's potential constraints, especially with respect to timescales.

**Delivery:** This extends your thinking about the characteristics of your audience in combination with your work's output format. It concerns the setting in which they will encounter and experience using your work: how will your visualisation be delivered to its audience? Is it going to be published for audiences to consume remotely – away from you – or will it be presented and delivered in-person? If you are not personally available to offer verbal explanations of key findings, descriptions of the data gathering process, assumptions or calculations, you may need to include these as annotated properties in your absence.

Do the needs of your audience, and the setting of their encounter with your work, require their understanding to be facilitated immediately, at pace, or is a more prolonged duration of understanding viable? Think about four distinct setting types requiring different delivery solutions:

- The *boardroom*: A setting characterised by limited time, limited patience and limited attention. Immediate insights and key messages need to be conveyed at a glance. This is the birthplace of the arbitrary 'one-pager'. There will likely be reduced appetite for engaging with anything that requires effort, such as an unfamiliar chart type or rich interactivity, unless someone is there to present it.
- The *coffee shop*: A more relaxed setting that might be compatible with a piece of work that involves more effort and requires more time to learn about the subject. Unfamiliar representations might be reasonable to use as long as sufficient assistance is provided about how to read the displays. Interactive features to enable interrogations may not pose the same obstacles to understanding in the way they would in other situations.
- The *cockpit*: A situation characterised by operational need, whereby the visualisation is offered as a tool or instrument, enabling monitoring and signalling of important matters.

Sufficient breadth and depth of content will be required to serve the multitude of different potential scenarios of understanding that might arise. Think of a wide-ranging organisational dashboard that serves multiple potential levels of enquiry, from high-level orientation to in-depth localised detail, to aid operational actions. Also consider cases like needing a reference device such as a map that provides answers to serve myriad different navigation needs.

- The *prop*: This is where a visualisation plays the role of a supporting visual device to accompany other forms of communicated understanding. This could be a presenter verbally highlighting key insights via a presentation alongside visual aides, or an article whereby the report author provides a written summary of the key findings that emerge from an accompanying chart or figure.

**Frequency:** If your visualisation work is going to be part of a frequently updated product, you might maintain the same fundamental design but the efficiency and reproducibility of your solution when updated with new data will be paramount. This could relate to the regular provision of an updated monthly sales report or a more irregular, but constant, need to refresh election polling forecasts as new data emerging from latest surveys get published. The upfront development of these solutions may be extensive, perhaps requiring the creation of adaptable templates or programmatically driven automation. Thereafter the updating of each release of data may only involve a limited amount of work thanks to the initial investment in time.

For visualisations intended to be one-off, standalone pieces you may have more scope to create a bespoke solution, not intended for ongoing reproducibility, so long as the efforts involved remain in your favour in terms of timescales (and maybe budget).

## Constraints

**Timescales:** The primary constraint is how much time there is to develop your solution. Most projects have a deadline attached to them, whether this is imposed by clients, mutually agreed upon, or set by yourself. Even if you do not need necessarily to adhere to a deadline, let's say for personal projects, it can still be useful to define a target date to help hasten your progress.

At the other end of the project's timeline, there is a start date. This may not be *now*. You may have to wait for certain conditions to be in place before you can even commence your work. If you are conducting analysis of some survey results, you will not have a complete, final dataset of responses to work with until the survey is closed. You know it's coming, but you can't start just yet.

During a project there may be milestone dates to aim for. These might be representative of occasions when you need to hold a meeting to show progress to a client. It might be a point in time when you know you'll need to commit to stop exploring your data and make some final editorial decisions ahead of the design stage.

Required task duration is the key time-related factor. There will be massive differences in the ambitions of a project that needs completing in two days compared to another that has two months. But if the two-day project only involves a small-sized task, one that might take only a few hours, that is going to be easily deliverable. A two-month deadline sounds great, on paper, but if the tasks involved in fulfilling that work are estimated to require three months' work, it will be a struggle to deliver in time.

Anticipating how long a task will take is a difficult thing to judge. You usually don't know until you've done it. And that's often too late to be useful! Even with experience working on a diverse range of projects, seemingly similar tasks can end up catching you out by taking quite different lengths of your time. Noting down how much time you spend on each major task, across your design process, helps you to at least be better informed about how you usually need to allocate your time. In turn, this gives a better chance of accurately estimating expected commitments on future projects.

'What is the *least* this can be? What is the minimum result that will 1) be factually accurate, 2) present the core concepts of this story in a way that a general audience will understand, and 3) be readable on a variety of screen sizes (desktop, mobile, etc.)? And then I judge what else can be done based on the time I have. Certainly, when we're down to the wire it's no time to introduce complex new features that require lots of testing and could potentially break other, working features.' **Alyson Hurt, Supervising Graphics Editor at NPR**

**Pressures:** These are forces adjacent to or external to a project that will impact in different ways. Financial constraints are one such pressure that may affect your project. Can you afford the costs of potentially necessary hardware, software licences, photographic or audio assets? Do you have budget available to employ outsourced services, like translating or illustration work? How much time you want to work on a visualisation project may not be viable if the costs involved are not aligned with the allocated budget.

Other pressures often emerge from the politics surrounding your subject, the data or the messages coming from this data. I have been involved in several projects where the charts produced showing data about cities or countries had to be sorted alphabetically, and not

by any other ranking measure, in order to preserve a certain diplomatic neutrality. Your work may be revealing of certain insights for which the implications are not going to be politically compatible at that moment in time. Situations such as these can be difficult to handle: you want to be respectful to the sensitivities involved, but you will also want to preserve the integrity of what you are communicating.

Cultural considerations may materialise when you are creating work concerning content or for people across different regions. Issues around the use of imagery, colour connotations or the symbology of certain forms may need to be delicately and sensitively handled. There may also be environmental considerations, particularly concerning the output of your work, that need to be observed specific to your context, as we touched on in the previous chapter.

**Design:** Restrictions around certain design choices are common, often informed by style guidelines that must be followed through the adopting of specific colours, typeface or fonts. Where appropriate, I will usually attempt to negotiate away from the absolute adherence of such guidelines as they can be unnecessarily restrictive to a visualisation design. Those negotiations are not always successful! You may dislike the colours, you may hate the typeface, you may be asked to include logos, which can take up valuable space and unbalance your composition, but rules are rules. We therefore need to know about these requirements now, not later.

Layout or size restrictions may also exist, dictating the space in which you have to work. For example, when producing graphics for journals or for digital outputs on a tablet or smartphone, you might have quite a small amount of space to utilise. Conversely, your output might need to

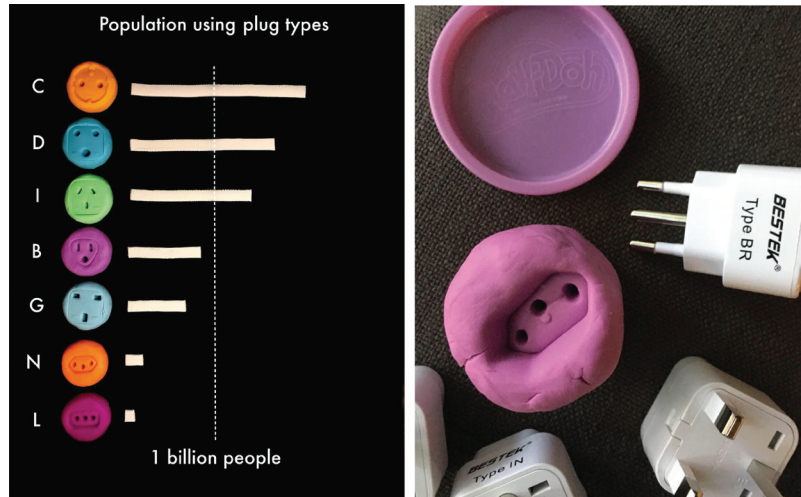
be very large, which can introduce completely different demands with the importance of legibility and resolution quality heightened.

Further creative pressure might come from what I would describe as market influences. The visualisation you develop may have to compete for attention alongside other work. For example, if you are creating a visualisation for a charitable organisation, how do you get a message across louder and more prominently than others competing for the same eyeballs? If you are working on an academic research project, how do you get your findings heard among all the other studies battling to create an impact? The creative desire to emulate or differentiate from other practices can prove to be a strong motive in your design thinking.

**Technological:** As I have mentioned in the book's introduction, there are myriad tools, applications and programming libraries in data visualisation, offering a varied landscape of capabilities. The technology you have access to will affect how digitally ambitious your work can be and/or how efficiently you will be able to make it. You can only achieve what your tools enable you to achieve. This influence will shape several stages of your process:

- *Working with data:* Technologies to help with gathering, examining, transforming and exploring data: What volume of data can your tools handle? How quickly do they perform actions such as searching, editing or re-calculating, especially with large datasets? Do they enable the automation of repeatable or time-consuming tasks, maybe through scripting? What range of statistical techniques are offered? How effective are the methods for quickly analysing and summarising the data? Do they enable means for exploring your data visually?
- *Data representation:* Technologies to help with making charts: What range of different charts do they offer? Is the process of constructing the charts straightforward or time-consuming? Do the tools facilitate customised solutions which might open up workarounds for expanding the options beyond standard capabilities? Can I export or publish any charts made external to the tool, such as an output for the web or as an image for a report?
- *Interaction and dynamic design:* Technologies to develop interactive features for exploration and control. What range of different interactive techniques do they offer? Is the process of incorporating these driven by a graphical user interface or does it require manual programming? Do the tools facilitate any degree of workarounds or means for expanding their potential capabilities in innovative ways? Can I embed dynamic assets such as videos or audio clips?
- *Data presentation:* Technologies to manage the inclusion of annotations, the use of colour and the composition of your work: What range of annotated features can you include? Are you able to control fully the appearance of those features? To what extent can you manage the colour properties applied to every visual element? Likewise, what degree of control do you have over the size, placement and ordering of all elements? Are you able to embed static assets such as photo-imagery or illustrations?
- *Publishing:* Technologies to disseminate the final output format of your work: Do the tools enable publishing as a presentation slide deck, printable or pdf report, published as an image for a website, compiled as a video, fabricated as a physical artefact?

In certain contexts, you may not need any technology. It may be possible to embrace more analogue or artisanal approaches, as demonstrated in the charming example shown in Figure 3.3, which is a visualisation (or more specifically a 'data physicalisation') created out of Play-Doh.



**Figure 3.3** 'Popularity of International Outlets', by Amy Cesal

The analysis concerns the various plug types used around the world. The imprints of the configuration of pins for each plug type are stamped into coloured chunks of Play-Doh, and then white lengths are measured out to represent the populations of people from the countries whose power systems use each type.

## 3.2 Establishing Your Project's Purpose

If the objective of a data visualisation is to facilitate understanding, as an outcome, it stands to reason that being clear about *what* understanding this will involve, and *how* you will accomplish its facilitation, will be a defining task of immediate priority.

### What Understanding Are You Facilitating?

For some people, the most natural mindset to occupy, when thinking about the purpose of their visualisation, is to work out what message they're seeking to communicate: what is this visualisation going to say to the viewer?

An alternative way to think about defining purpose, and the approach I find most useful, is to formulate questions.

In the subtitle of this book, I've used the words 'data driven' design, but you could make an argument for replacing these with 'curiosity driven' design. Defined as 'a desire to know or to learn something', curiosity is what ignites the purpose of a data visualisation, in my view. Without somebody having a desire to know or to learn something, data would never be captured, or stored, or gathered, or analysed, or visualised, or presented to others.

When defining *what* understanding your visualisation is going to facilitate, I find forming questions, rather than messages, to be the most useful way to shape my ideas. These will be questions that I know, I anticipate, I assume, or maybe just that I hope represent an expression of relevant understanding that will be useful to answer for my visualisation's audience.

I commence all my visualisation projects by writing down lots of different questions, covering a spectrum of potential enquiries. When the purpose of your work is not fixed or imposed by others, it can be useful to start off broadly and instinctively. Even if it's not feasible to pursue answers to all ideas you form, it at least stimulates your curiosity and may open up possible paths of enquiry to pursue.

An example is illustrated in Figure 3.4, based on an exercise I like to practise. Take any newspaper or news media headline and imagine as many different data questions as you can think of that would be interesting to ask – and get answers to – about a subject.



**Figure 3.4** Brainstorming Ideas for Potentially Relevant Questions about a News Topic

Take it a stage further. Assuming you were able to source the required data, imagine you were tasked with creating visualisations to accompany the story: answering *which* of these data questions would most enrich a viewer's understanding? This is an exercise I find really helps refine my instincts for using different language expressions – the who, why, when, where, what and how – that give analytical direction to answering data questions. As you'll see described further in Chapter 6, if charts communicate visual answers to data questions, we can't select the right chart until we know what is the right question. Adopting this mindset early helpfully aligns much of your subsequent decision-making.

The scope for early specificity in what understanding you are facilitating will vary depending on context. Sometimes its not yet possible to know what interesting data questions could be relevant to a subject. On some occasions I commence a project with very few ideas about potential questions coming to mind. This might be because I don't yet have enough knowledge about my subject matter

'If you want to work in data visualisation, you need to be relentlessly and systematically curious. You should try to get interested in anything and everything that comes your way. Also, you need to understand that curiosity is not just about your interests being triggered. Curiosity also involves pursuing those interests like a hound. Being truly curious involves a lot of hard work, devoting time and effort to learn as much as possible about various topics, and to make connections between them. Curiosity is not something that just comes naturally. It can be taught, and it can be learned.'

**Professor Alberto Cairo, Visualisation Specialist and Knight Chair in Visual Journalism, University of Miami**

to know what could be interesting. Similarly, it might be when I haven't yet become acquainted with the available data and I will need to explore that first before becoming familiar with what questions are even answerable.

Sometimes, it is the availability of a dataset, and the associated opportunity to explore its potentially interesting contents, that drives the motivation to initiate a visualisation project. Let's suppose you are a researcher specialising in the history of music and you come across a rich source of data detailing characteristics of some of the most successful pop songs of all time. You suspect there are *many* potentially fascinating things to discover in this data, but you do not yet know what *the* most interesting thing will be. Moreover, you've not yet even got an audience in mind, other than yourself. The process in those circumstances would be driven by a need to explore the data, pursue speculative exploration and thoroughly sift through this material until you unearth discoveries representing interesting insights you believe would be desirable for others to know or learn about.

Remember, this stage of activity represents the first but not the last occasion when there is the opportunity – and need – to do this thinking about purpose. You can define options now, refine them during the 'working with data' stage, and then finalise them when forming your ultimate editorial decisions by stage three.

There are other occasions when you will start a project with an already quite specific sense of what understanding you are facilitating. You already have clarity about the question your work will be seeking to answer. We talked about *clients* earlier, in the form of a manager, a colleague or a client, and it may be their own desire to know or learn something that's driving this process. They are going to simultaneously play both the role of client and target audience. Let's imagine a situation where they want to better understand sickness absence trends amongst staff in their department during the past year. They articulate this need to you, initiate a request for a data visualisation, and your job is to inherit this curiosity, use it to drive the purpose of your visualisation and work towards facilitating the understanding they seek.

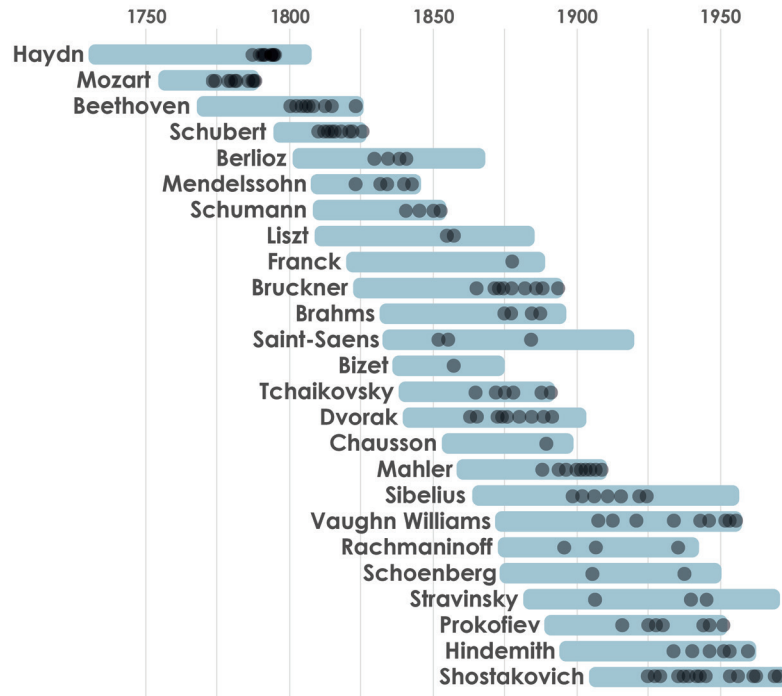
A different scenario would be your client expressing a request about what others might desire to know or learn: 'Can you produce some analysis showing trends of sickness absence amongst staff this year to share with all heads of departments?' might be the initiating request. They've commissioned you to undertake a visualisation project, but the purpose is to facilitate understanding for a different audience, not for the client and not for you. Whether this curiosity is a known desire held by the intended audience ('they need this'), an anticipated one ('they will find this interesting'), or an assumed one ('someone might find this interesting') will determine what chance the visualisation has of accomplishing its potential aims.

Being driven by the need to serve an anticipated or assumed interest of your audience is perhaps the most complex challenge in visualisation, in terms of having least control over and certainty of how to accomplish a successful outcome. If you know your audience well enough and are able to identify their needs, your purpose is shaped around what you *know* they want answering. Otherwise, you will need to make reasonable judgements to anticipate what *most* people will likely find *most* interesting.

Sometimes you are facilitating your own understanding, not anyone else's. Figure 3.5 shows an example of a visualisation produced by Chris Ingraham, a reporter covering data driven stories for the *Minnesota Reformer*. As he states in his post that accompanied his published visual, 'It's been a while since I made a chart just for me and, last night, I was wondering how the major symphony composers' timelines overlapped, so I put this together, with an assist from Carnegie Hall's fantastic data lab to come up with a metric for "major symphony".'

## The symphonic canon: a timeline

Major symphonic composers' lifespans and the dates of their symphonies



Notes: "Major" defined as composer of at least one of Carnegie Hall's 100 most-performed symphonies. For Haydn and Mozart, only 10 most-performed symphonies are displayed.

Sources: Carnegie Hall Data Lab, Wikipedia

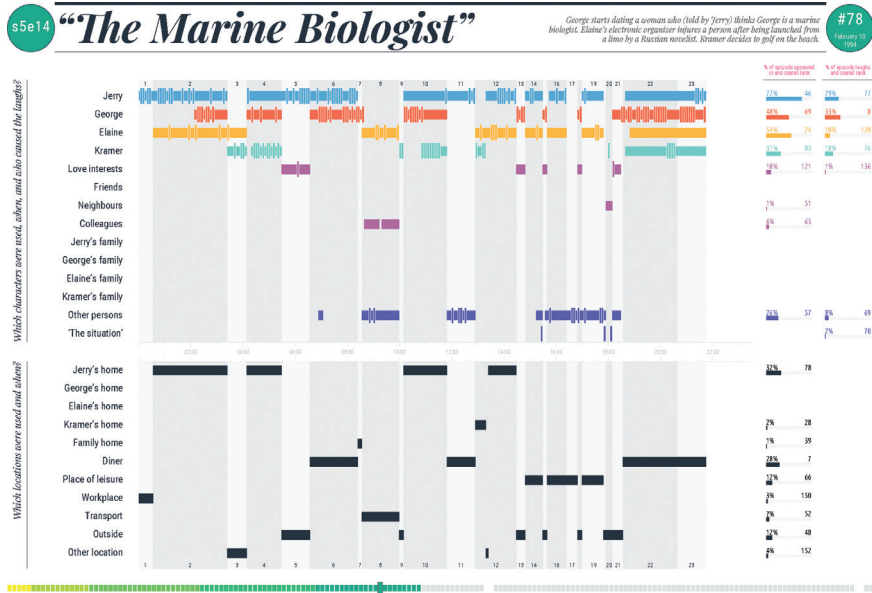
Christopher Ingraham

**Figure 3.5** 'The Symphonic Canon: A Timeline', by Chris Ingraham

Visualisation pieces like this represent a type of creative process often characterised as being a 'pet' or 'passion' project. They are entirely self-initiated, with no client involved. You have freedom to follow your own enquiry, shaped only by the limitations of your imagination and interests. With no fixed audience requirements, in a sense, you do it for you or you do it for other people like you who may share a similar curiosity.

Figure 3.6 shows an example of a visualisation I produced that entirely matches this motivation. This was triggered by my own specific curiosity. Titled 'The Seinfeld Chronicles' it remains the most ambitious data visualisation piece I've ever produced, driven by my pursuit of a curiosity I had about the *rhythm and texture* of the US sitcom Seinfeld.

Continuing the theme of music, it is said the 'musicality of language' is something that defines the signature comedic talent of Jerry Seinfeld and the show's co-creator, Larry David. Like music, situation comedy has texture and it has rhythm. The texture is formed by arrangements of different characters and different locations, which create the situations. The rhythm is shaped by the timing and pacing of these arrangements, which fuels the comedy. My work was an exploration of the show's musicality, investigating how Seinfeld and David orchestrated their resources – the people, places and the dialogue – to translate their creativity from the page and onto our screens.



**Figure 3.6** Excerpt from 'The Seinfeld Chronicles', by Andy Kirk

This project was initiated by me with the desire to facilitate my own understanding. I was not explicitly serving the expressed curiosity of any other client or defined audience, though I was anticipating the portfolio of visuals produced would be interesting for people who share my Seinfeld fanaticism. With early clarity about my purpose, this in turn clarified all my subsequent decisions about what data to gather, what analysis I could conduct. There were many things I *could* have explored about Seinfeld, about US sitcoms in general, or about TV shows more broadly, but my particular enquiry was the only purpose I was interested in serving.

Sometimes the understanding you are facilitating, though established, is still too enigmatic in its definition. In my Seinfeld case, 'What was the texture and rhythm of Seinfeld episodes?' is not a directly answerable question. To facilitate understanding through a visualisation I would need to form several specific sub-questions, that were more directly analysable and visualisable, which would collectively answer my overall enquiry.

- Structure: How was each episode organised into discrete scenes?
- Characters: When and for how long did each character appear on screen?
- Laughs: When did each character cause laughter amongst the studio audience?
- Locations: When and for how long were different locations used?

Ultimately, deciding that these are going to be the questions to which your contents will provide answers is the duty of your editorial thinking, when we reach stage three. That is when a commitment about what understanding you are facilitating will be necessary, ahead of the design thinking that follows in stage four.

Take another example. If you are a runner with a fitness-tracking device or app, you might finish a run and wonder, 'How good was that run?' This is a broad enquiry. To form an answer requires the synthesis of several distinct pieces of information ('How far? What time? What route? What achievements? What previous times?') that collectively provide understanding about how good the

run was. By contrast, if you just want to know, ‘In what time did I complete the run?’, this is a specific curiosity that can be effectively answered by a single piece of information.

On many occasions I have embarked on a visualisation project with a clear purpose in mind but then, having become better acquainted with the subject through its data, different enquiries have emerged as being more relevant. When working on your own pet projects, shifting your purpose, according to where you are seeing opportunities for discovery, is easy to justify. When you are working on a project tasked by a client, you might find there is less room for manoeuvre away from the initial pursuit you were tasked with. That said, you should still seek constant dialogue if you strongly feel another route might be more interesting. After all, it serves nobody’s purpose if you remain anchored to an enquiry that no longer reflects the most relevant aspects of a subject or if you are drifting away from the initial scope of the work.

## How Will You Facilitate Understanding?

Articulating your thoughts about the project’s purpose is not limited to what understanding will be facilitated. It extends to how you facilitate understanding most effectively. What type of visualisation may you need to create in order to accomplish your aims?

As with any other form of communication, there are different tactics required to achieve different outcomes. Think of the varied ambitions one might have with any given communication encounter: to educate, explain, inform, instruct, persuade, motivate, surprise, shock, scare. Any form of communication that has a purpose ‘to inform’ is going to require different content and delivery choices than a desire ‘to shock’. If you are seeking to have a persuasive impact when presenting to parents trends around childhood obesity, that would probably require a different communication strategy than if you were informing health professionals about the trends in their local area. There are different emotions involved, different stakes in the information presented, different forms of agency in what consequential actions are possible to take.

In the first chapter, I described how viewers go through three phases of understanding: *perceiving*, *interpreting* and *comprehending*. I explained how, as visualisers, we have no control over the outcome of comprehending, for which the onus is more on the viewer and determined by their attitude and connection to the subject matter. However, visualisers do have control and influence over the other phases, perceiving and interpreting. This materialises in the type of *tone* and the type of *experience* through which we choose to facilitate understanding.

## Tone

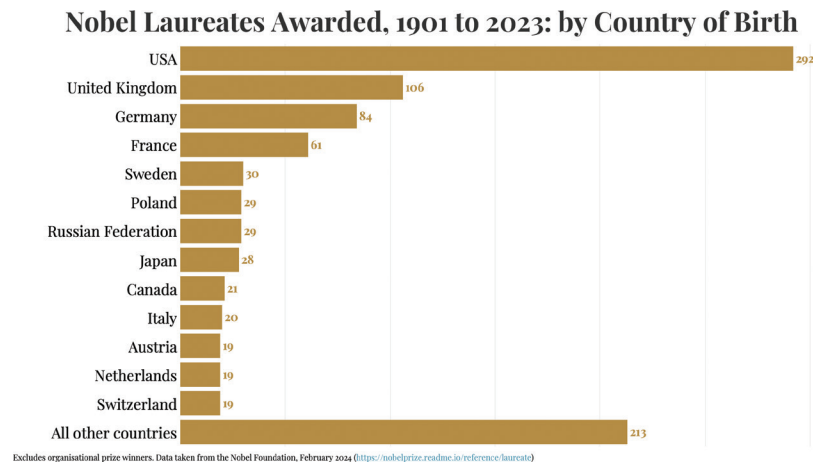
The tone conveyed by a visualisation has influence on the *perceiving* phase of understanding. In judging the most suitable tone for your work, you are deciding whether to place more emphasis on the viewer being able to *read* data or *feel* data (Figure 3.7).



**Figure 3.7** The Spectrum of ‘Tone’

**Reading tone:** A visualisation that conveys a *reading* tone places emphasis on optimising quantitative perception. It is concerned with being detail orientated, enabling precision and efficiency in judging the represented data. It is about heightening the analytical qualities of the content presented and not the emotional qualities. This is a visual style embodied by adjectives such as pragmatic, authoritative, analytical, conservative, utilitarian and even boring.

The design style employed, especially around chart and annotation choices, will typically seek to make it easier for a viewer to determine the magnitude of and the relationships between displayed values. Representation methods like bar charts, as shown in Figure 3.8, are the quintessential method for demonstrating this tone of voice.



**Figure 3.8** ‘Nobel Laureates Awarded, 1901–2023: by Country of Birth’

By representing the size of a quantitative value using the proportional size of a line mark (even though we call them bars, they are technically lines), bar charts facilitate both general sense-making and precise point-reading, thus heightening the perceptual accuracy for the viewer. The inclusion of direct labels or axis scales is typically how bar charts are presented. But even if we didn’t have the value labels against each bar, as shown in this example, you would quickly ascertain that the USA value is around three times the size of the next largest value, displayed via the UK bar. Although similar in magnitude, you can see that the value for Canada is slightly larger than the one for Italy, which in turn is slightly larger than the bars for Austria, Netherlands and Switzerland. Once again, even without labels, you’d probably estimate France’s number to be around 60 and Germany’s close to around 85. For exactness in value-reading the direct labels serve that purpose, but without their inclusion, the judging and comparing of quantitative size using a bar chart is sufficiently accurate.

Bar charts are so ubiquitous and so necessary because, as shown, they make it efficient to accurately answer quite reasonable questions like ‘What is the size of value X?’ and ‘By how much is the size of value Y larger than value Z?’ Many visualisations you will produce will most likely lean towards offering this kind of *reading* tone.

Indeed, you might reasonably ask why would you ever *not* seek to optimise the accuracy and efficiency of value judgements? Surely anything that compromises on this is undermining the

accessibility of your design and maybe even jeopardising its trustworthiness? Well, this is why it is so crucial to carefully think about what it is you are seeking to accomplish. There are other paths to consider, as typified by the *feeling* end of this tonal continuum.

**Feeling tone:** In contrast to reading values, sometimes we might justifiably decide to place more emphasis on the *feeling* of data. A feeling tone might be suitable in circumstances where there is less need for precision and more need for big-picture sense-making. It might be necessary to portray your subject with an injection of emotional stimulation, to give the understanding being facilitated greater potency. In contrast to the reading of data, a tone of voice aligned to feeling data is embodied by adjectives such as impactful, emotive, human, figurative, visceral and novel.

With a feeling tone, the perceptual judgements are more qualitative in nature, aligned with the notion of ‘getting the gist’. This means the viewer can ascertain general observations, at-a-glance, of the hierarchy of large, medium and small properties of your data. The viewer will gain a sense of major patterns that reveal clusters or trends suggesting which things are generally going up and/or going down. Just not by exactly what amount. They will be able to notice things that are bigger-than or smaller-than, but again without judging precisely by how much.

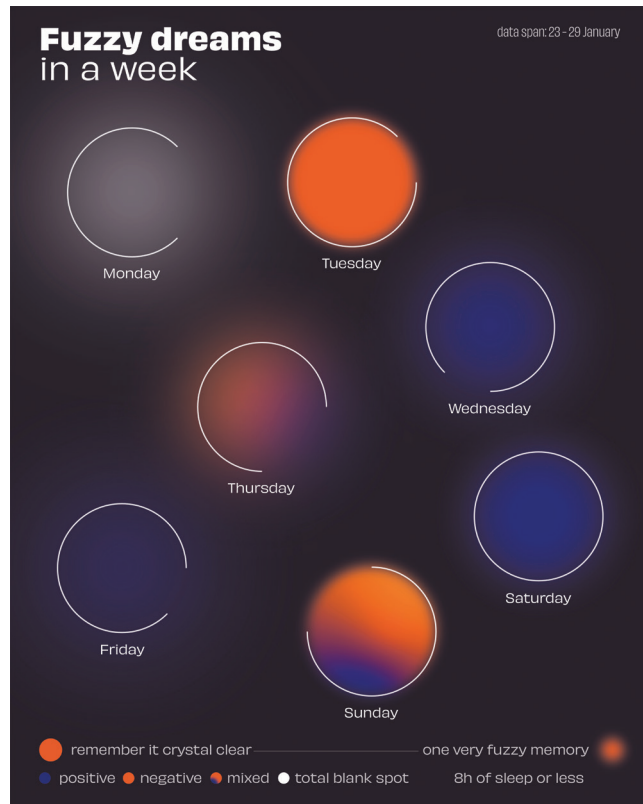
As described at the start of this book, the benefit of visually representing data is that it offers something different – and often something better – than the function afforded by a table of data values. It enables a viewer to see the data. Sometimes you want the viewer to see the data *and* the subject.

In Figure 3.9 we see a graphic titled ‘Fuzzy Dreams in a Week’ summarising dreams experienced over 7 days. Each day is shown as a separate circle with a ring surrounding it, the length of which is the duration of each night’s sleep relative to a maximum filled ring of 8 hours. The dreams’ representation comes through the colouring – orange is a negative dream, blue positive, with both colours used for emotional mixtures – and the sharpness or blurriness for the clarity of memory of the dreams. The term ‘fuzzy’ is aligned perfectly with the use of the visual attribute of blur. Some dreams remain vivid when you awake, other dreams can be immediately forgotten in their detail but maybe not in their emotional legacy. You’re not invited to read the charts, rather ‘get the gist’.

In the project illustrated in Figure 3.10, you see analysis from 2016 of the families who, back then, had the most financial clout when it comes to funding presidential candidates. The data quantities are portrayed using stacks of ‘Monopoly’ pieces piled up on the White House lawn. The red houses represent the small number of distinct families who have contributed nearly half of the initial campaign funding, the green pieces are representative of the total households in the USA. Like Figure 3.9, here you cannot perceive the data in a quantitative sense – through any notion of counting the number of piled pieces – nor are you supposed to. But you can visually approximate the remarkably disproportionate balance and power of wealth in a qualitative sense, through the juxtaposition of *many* green pieces compared to the *few* reds.

Thinking about tone is to recognise semantically what your subject is about: what activity, instance or phenomenon does it represent? Is it about people and places, or products and transactions? Is it about the real or the abstract? Data is more than just values formed of numbers or

‘There’s a strand of the data viz world that argues everything could be a bar chart. That’s possibly true but also possibly a world without joy.’ **Amanda Cox**, Executive Editor of Data Journalism at **Bloomberg News**



**Figure 3.9** 'Fuzzy Dreams in a Week', by Kosara Keskinova



**Figure 3.10** 'Buying Power: The Families Funding the 2016 Presidential Election' (*New York Times*)

text. Considering the underlying phenomena helps you feel its spirit more closely. It prepares you for the responsibility you will face when curating understanding about *this* subject matter. As we saw in Chapter 2, with the case of the 'Gun Deaths in Florida' graphic (Figure 2.10), some subjects are inherently more emotive than others. You might choose to amplify or suppress the

emotion of the subject, and you need clear conviction in deciding how to find the most suitable tone of voice.

For subjects that carry the weight of strong emotion, there might be good reason to exploit the inherent feelings. Encapsulating emotional sensations like fear, shock or fun through your design choices might accelerate the meaning of the subject and potentially affect that elusive ‘comprehending’ phase of understanding.

Provoking feelings through data could be seen as being somewhat manipulative. To a certain degree it possibly is and there are risks associated with misjudging handling emotional attributes. A playful approach to portraying data about a serious topic will demonstrate insensitivity and probably undermine the trustworthiness of your work, even if you have created an accessible and elegant solution. As long as you are faithful to the underlying data and the subject’s visual embodiment is not superficial or deceptive, I believe it is an entirely appropriate option to consider following when the circumstances suit.

‘Find loveliness in the unlovely. That is my guiding principle. Often, topics are disturbing or difficult; inherently ugly. But if they are illustrated elegantly there is a special sort of beauty in the truthful communication of something. Secondly, Kirk Goldsberry stresses that data visualization should ultimately be true to a phenomenon, rather than a technique or the format of data. This has had a huge impact on how I think about the creative process and its results.’

**John Nelson, Cartographer at Esri**

It is important to note that any visualisation work that leans more towards ‘feeling’ is typically the exception and will be relevant to only the minority of situations. But it’s worth being aware of this end of the tonal spectrum, even if only to give context to why the majority of your work appears to favour the bar chart. A skilled visualiser needs an adaptive view. Experience imbues an ability to judge the appropriate occasions whereby the purpose of visualisation will require different levers being pulled.

When it comes to defining the best-fit choice of tone, it is often possible to think of a blend of options in combination. There will be projects you work on that involve multiple chart assets, multiple interactions, different pages and deeper layers. The mantra proposed in 1996 by Ben Shneiderman, one of the most esteemed academics in this field, namely ‘Overview first, details on demand’, informs the idea of thinking about different layers of readability and depth in your visualisation work accessed through interactivity. Some of the chart types that you will meet in Chapter 6 can only ever hope to deliver a *gist* of the general magnitude of values (the big, the small and the medium) and not their precise details. A treemap, for example, is never going to facilitate the detailed perceiving of values. In the example shown in Figure 3.11, showing Standard & Poor’s 500 index stocks categorised by sectors and industries, the area of each rectangular shape represents the relative market capital for each company included. The colours indicate changes in valuation over the past 24 hours.

Our perceptual system is relatively poor at estimating scales of areas. If you wish to compare the size of one stock (e.g. Microsoft, MSFT in the top left) with another (e.g. Google, GOOG in the top middle) it will not be easy to make such a judgement accurately. However, you can get the sense that they are both relatively large. A chart like this offers a single, at-a-glance view of the hierarchy of values (big, medium, small) as well as prominent observations of colour (vivid red vs vivid green, colour blindness notwithstanding). In this case, features of interactivity exist allowing the user to hover over a given shape to reveal a tooltip containing precise details as value labels. With those details available, a viewer’s perception shifts from a big-picture qualitative feel to more precise quantitative reading: overview first, then details on demand.

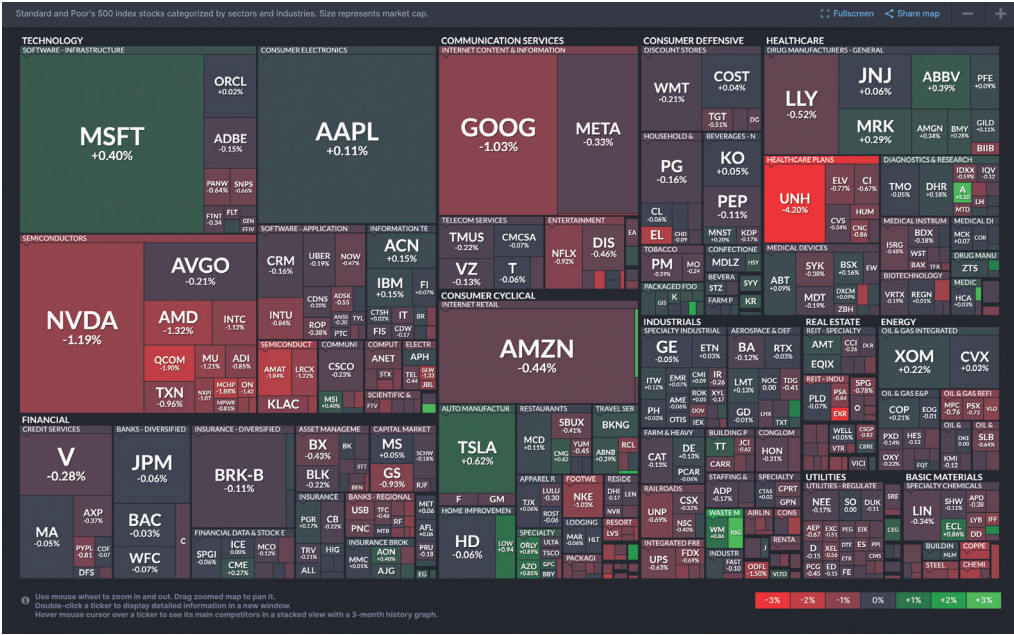


Figure 3.11 'Finviz: Standard & Poor's 500 Index Stocks' (www.finviz.com)

## Experience

The experience offered by a visualisation has an influence on the *interpreting* phase of understanding. Specifically, in judging what is the most suitable experience through which to facilitate understanding you are deciding where the onus lies for interpreting. Whereas tone embodies a continuum, the judgement of the most suitable experience is more distinct: *explanatory*, *exhibitory* or *exploratory* (Figure 3.12).



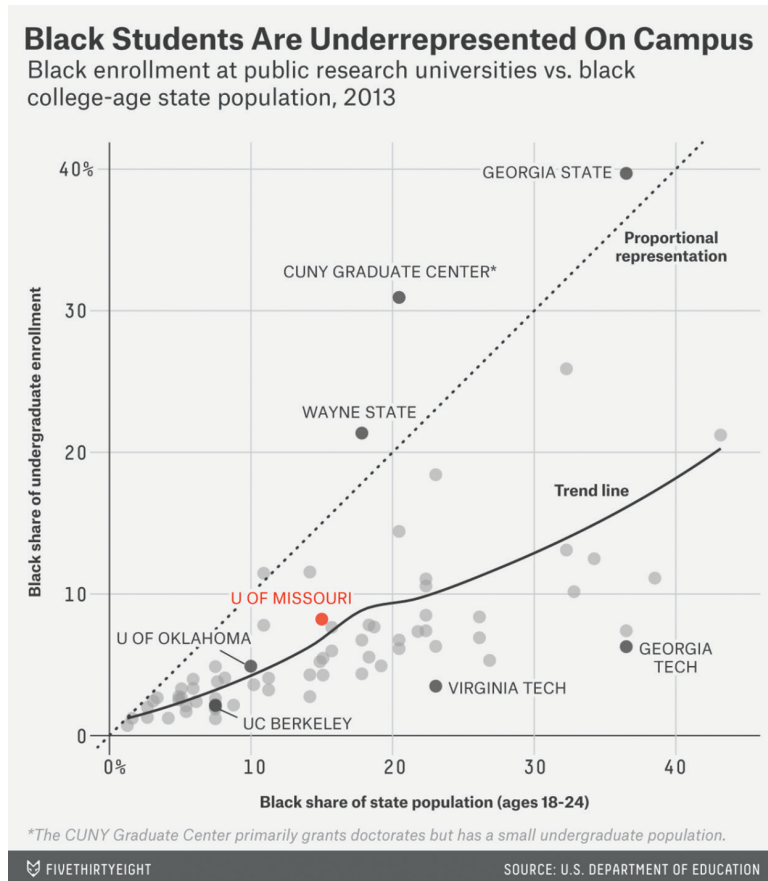
Figure 3.12 The Classifications of 'Experience'

**Explanatory** visualisations are characterised by the visualiser taking responsibility to elevate important observations and explain them to the viewer so they can assimilate the meaning of what is presented. I often think of *quotation marks* as being emblematic of explanatory visualisations as they convey a visualiser saying something.

Explanatory visualisations will include simple visual cues that direct the eye's attention towards key features of a chart's display. This might involve the use of contrasting colour properties to

establish a hierarchy of prominent marks in a chart. It might be the use of annotated captions that provide textual commentary about key interpretations.

An example of an explanatory experience is shown in Figure 3.13, featuring a chart which was featured in an article (published November 2015) reporting on protests across US schools regarding the under-representation of black students.



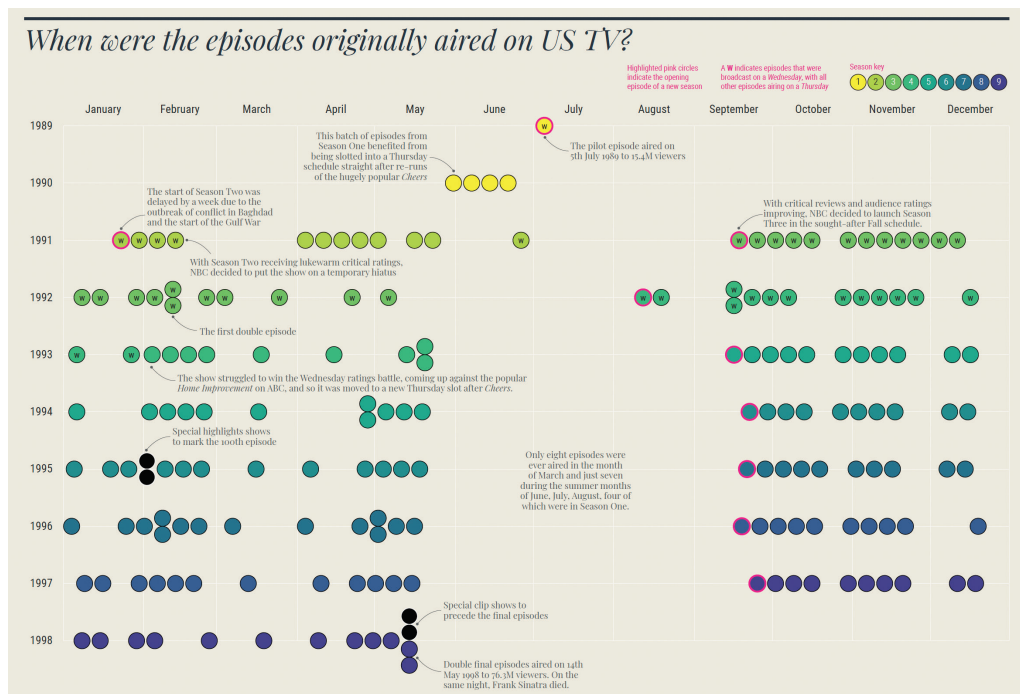
**Figure 3.13** Taken from 'Mizzou's Racial Gap is Typical on College Campuses' by FiveThirtyEight

Here you see a scatter plot comparing the share of enrolled black students for different public research universities (along the vertical y-axis) with the share of the college-age black populations in the respective states (along the horizontal x-axis). With the protests beginning at the University of Missouri, the chart uses red to highlight this data point as the primary item of interest. Other notable colleges, as mentioned in the article, are emphasised using darker dots and labels to illustrate useful comparisons. With additional visual overlays like the trend line and dotted line marking where proportional representation lies, the viewer's attention is drawn to the implication of what it means to be positioned in different regions of this chart – is it good or bad, typical or atypical? It's not just *where* the dots are, but what the *where means*.

A useful way to consider the role of an explanatory visualisation is to think how you would verbally explain key insights from a chart if you were there, with your audience, in-person. Think about what features you would point out as being most interesting? Which values would you draw attention towards, and which would you ignore? The traits of a good explanatory visualisation are that it can effectively stand alone without the need for in-person explanations, beckoning the viewer towards the important interpretations.

Explanatory visualisations need the visualiser to possess (or acquire) sufficient knowledge about the topic in order to identify the most relevant insights worthy of elevating for the viewer. If you have something worthy of saying, say it with an explanatory visualisation.

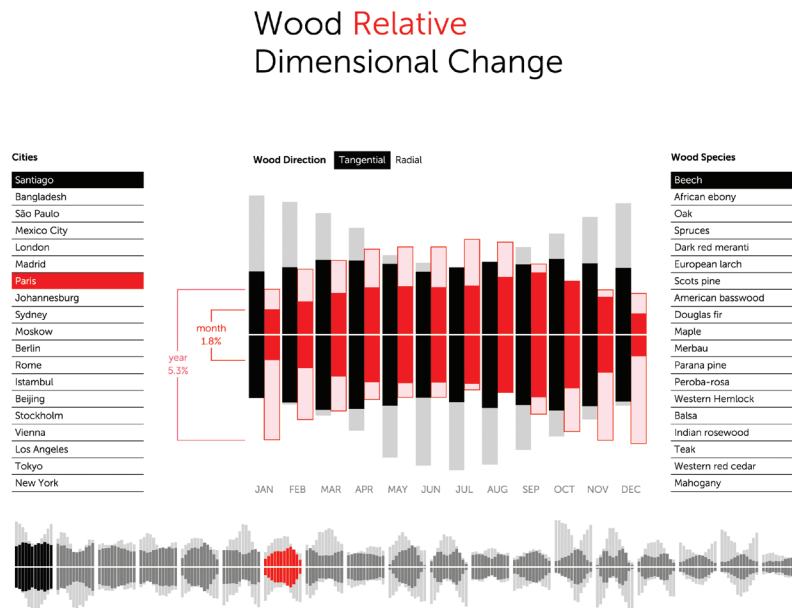
An example demonstrating this is shown in the chart in Figure 3.14, which is another piece of analysis taken from my Seinfeld project. This is a timeline graphic showing when each episode, across the nine seasons, was originally broadcast on US TV. You will see scattered across the display several captions that describe interesting (at least, they were to me!) observations about some of the scheduling patterns and/or gaps. These are things I would be pointing to and verbally explaining if I had the chart on a screen and an audience in front of me. In the absence of that setting, I can recreate my explanatory commentaries through the inclusion of simple text boxes and arrows.



**Figure 3.14** Excerpt from 'The Seinfeld Chronicles', by Andy Kirk

**Exploratory** visualisations differ considerably from explanatory visualisations in that they are designed to enable the viewers or – more specifically in this case – the *users* to interact, discover, and form their own interpretations. The *question mark* symbol is an appropriate icon for exploratory visualisations, given they exist to help a user answer their own questions.

The basic level of exploration will come through simple interrogation of the data, through features like filtering. This may enable a user to choose to display only certain categories of interest or a selective range of values. An example of this type of visualisation is shown in an interactive project in Figure 3.15. It was developed to allow users to explore different measures concerning the dimensional changes of different wood species, over time, across selected cities of the world. Unlike with explanatory visualisations, here you will find no captions or emphasised observations. There are no indications from the visualiser about what contents are significant or insignificant, what is good or bad: no assistance to help the user interpret the meaning of this data. Projects like this exist to enable users to form their own interpretations.

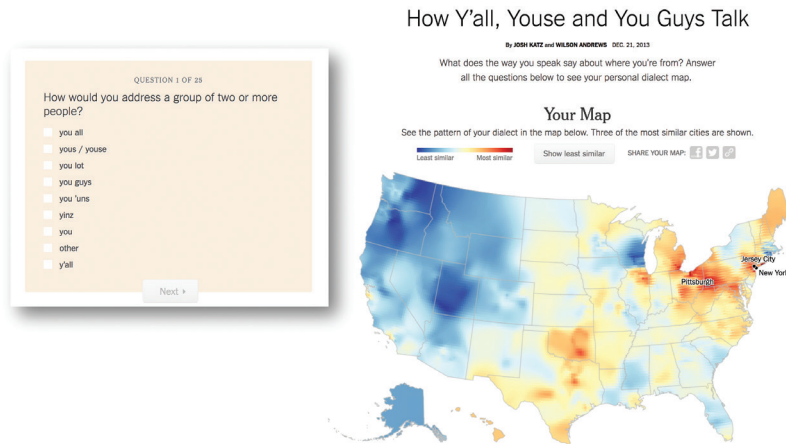


**Figure 3.15** 'Dimensional Changes in Wood', by Luis Carli (luiscarli.com)

Facing the responsibility for translating 'what it means', this kind of experience will be most suitable for audiences that have the requisite knowledge and/or motivation to seek such interpretations by themselves. Indeed, it would be a reasonable assumption that the users will be better equipped to do this than to receive interpretations from the visualiser.

Some exploratory visualisations will expand the interactivity offered towards two-way participatory experiences, whereby a user is able to contribute their own data and receive some form of feedback. As users we are naturally drawn to experiences like quizzes or projects that allow us to make sense of our place in the world (e.g. how does my salary compare with others? How well do I know the area where I live?). Figure 3.16 shows the *New York Times*' so-called 'Dialect quiz map', an especially successful exploratory visualisation employing this participatory approach to great effect.

In this case users are invited to complete 25 questions about their use of language terms in different scenarios. Based upon their responses and the others gathered in the associated (and ever-growing) study, a map of the US is graphically revealed showing the similarity or otherwise of their apparent dialect. This map is customised entirely by the contributions of the participating user: it shows *them*



**Figure 3.16** ‘How Y’all, Youse and You Guys Talk’, by Josh Katz and Wilson Andrews (*New York Times*)

who *they* are. There are thousands, maybe even millions, of different combinations of possible outcomes, but the end result encountered by the user is unique to their participation.

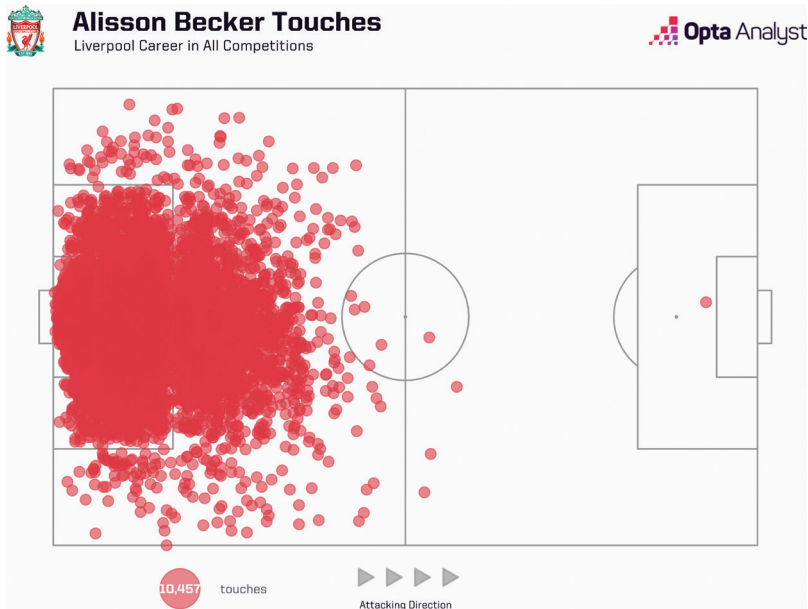
We discussed the idea of ‘unimpeded access to understanding’ as being a key attribute of the accessible design principle. This can be a common failure of exploratory visualisations when a user is left with a feeling of ‘So what?’ or ‘I don’t know what you want me to do with this project?’ When this happens there is a clear disconnect between the intentions of your project and the expectations of those using it. You might have offered an exploratory tool when the needs of the audience would have been better served by an explanatory offering.

**Exhibitory** visualisations are characterised by being neither explicitly explanatory nor functionally exploratory. These visualisations are simply visual displays of data asking the viewer to do the work to establish their meaning. I find the *ellipsis* an appropriate emblem for exhibitory visualisations, representing how the visualiser passes the onus on to the viewer to interpret.

One way to think about this type of experience is in relation to the exhibit of an artwork: it takes the interpretative capacity of the viewer to be able to understand the *content* of a display as well as the *context* of a display. Similar to exploratory experiences, exhibitory projects rely entirely on the audience having the motivation and capacity to interpret. Unlike exploratory visualisation, exhibitory works do not provide means for interactive adjustments and interrogations, so interpretation is formed only through the act of seeing and thinking.

It’s maybe reasonable to ask, therefore, what the value is of an exhibitory visualisation? If your work is intended for a specific, targeted audience whom you know to be sufficiently knowledgeable about the subject matter and the analysis, it might not be necessary to offer an explanatory experience. Moreover, there may be no scope for enhancing the experience with any features of interrogating interactivity.

An example of a perfectly suitable exhibitory visualisation is shown in Figure 3.17. This chart maps the locations of each of the 10,457 instances (accurate to February 2024) when Liverpool FC’s goalkeeper, Alisson Becker, has touched the ball during his career at the club. It was published on the Opta Analyst website, a place visited by people with a passion and knowledge of football, and an interest in football analytics.



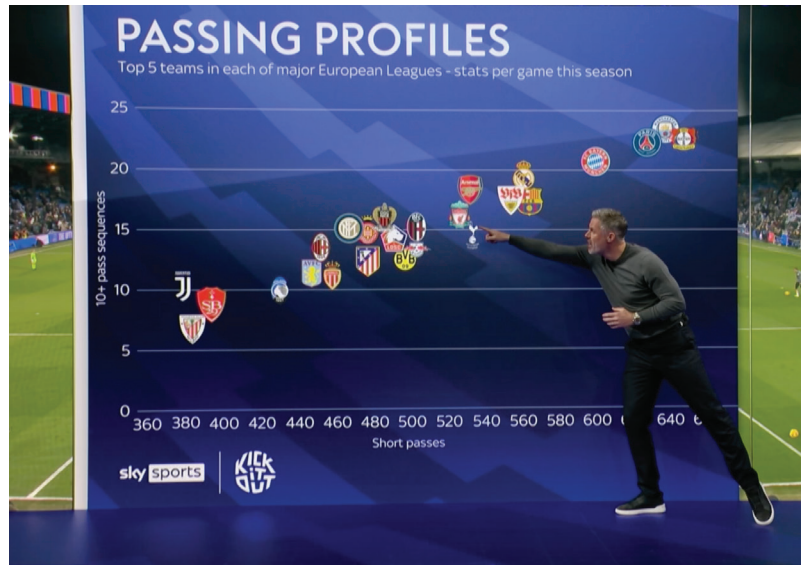
**Figure 3.17** 'Alisson Becker Touches', by Opta Analyst

In this chart we see lots of red dots filling most of the left-hand side of the pitch map, a handful of dots near the halfway line, and one dot far over on the right. This single dot clearly stands as an outlier, but what's the story of that data point? That mark represents the instance of a touch made by Alisson at 94 minutes 18 seconds during the final moments of the West Bromwich Albion vs Liverpool Premier League match on 16 May 2021. It was more than just a touch, though, it was header from a corner that flew into the back of the net. A winning goal in a crucial match, scored by a goalkeeper. Since the Premier League was formed in 1992 this was only the sixth time a goalkeeper has scored, the first to do so to win a game, and the first goalkeeper to score in a competitive game for Liverpool since they were founded in 1892. During this unsettling, artificial period of pandemic-era football, played in enforced empty stadiums, the impact of this goal for Liverpool fans was huge. For the player, having lost his father just three months earlier, it was a hugely celebrated but deeply profound moment of emotion.

For certain audience demographics, this chart alone won't mean a thing. They won't have any interest in football and they won't have reason to care about the analysis. For others, they might understand the subject, but may not necessarily grasp the significance of the analysis. Some cohorts know its significance, they just won't wish to confront it (Hello, Evertonians). For others – people like me, lifelong Liverpool fans, who might consider this amongst the greatest charts of all time – it speaks for itself. Because I understand (and care about) the subject's context, I am able to draw interpretations without the need for explanatory assistance.

Sometimes a visualisation isn't intended to speak for itself, but a presenter can speak for it. I described a few pages back the imagined scenario of presenting a chart to an audience to help you to imagine what explanatory features you might include if you weren't there. Well, what if you are there? If you are delivering a presentation with visual aides, as demonstrated in Figure 3.18, you would use exhibitory-style visualisations as your visual prop and deliver the explanatory

observations verbally as you point out key features using physical gestures. I would define this situation as involving an *exhibitory* visualisation but with the understanding facilitated through *explanatory* delivery.



**Figure 3.18** Clip Taken from Sky Sports Monday Night Football, Featuring Presenter Jamie Carragher's Analysis of Passing Profiles

A further final case for creating exhibitory visualisations will be if they are being included as figures within written articles or reports. In and of itself the visualisation does not need to be explanatory in its presentation, rather it exists as a graphic image to refer to from within the adjacent passages of text. That written communication is therefore where the explanatory narrative manifests.

### 3.3 Conceiving a Creative Vision

The definition of *vision* is 'the ability to think about or plan the future with imagination or wisdom'. Though we are far away from getting deep into design thinking yet, at this early stage of the process there is an opportunity – and probably an urge – to harness initial ideas about how you might imagine your work. These are the early seeds of thoughts about what the eventual solution you are working towards might look like.

In *Thinking Fast and Slow*, author Daniel Kahneman describes two models of thought that control our thinking activities. He calls these System 1 and System 2 thinking: the former is responsible for our instinctive, intuitive and metaphorical thoughts; the latter is much slower and more ponderous, by contrast, requiring greater mental effort when being called upon. System 1 thinking is what you want to harness at this first stage: what are the mental impressions that form quickly and automatically in your mind when you first think about the design opportunity you are facing?

You cannot switch off System 1 thoughts. Mental visualisations are what we instinctively ‘see’ in our mind’s eye when we consider the subject or nature of a task. You will not be able to stop them happening so, rather than stifling your natural mental habits, this is a good opportunity to allow yourself space to begin imagining.

When you think about your visualisation project, and about the subject matter it pertains to, what feelings does it evoke? What colours do you see? What shapes and patterns strike you as being semantically aligned with the topic? Sometimes instinctive ideas are reflections of our culture or society, especially the connotations of colour usage. This can be useful thinking to glimpse into the similarly impulsive way the minds of your audience might connect with a subject when they encounter your eventual solution.

Think back to the example shown in Figure 3.10 about political ‘buying power’. As a commonly recognisable metaphor of wealth, using Monopoly pieces was an entirely reasonable way to represent the data as being symbolically congruent with the subject involved. In Figure 3.19 we see a visualisation about the wine industry, showing the top grape varieties grown. In the upper part of the graphic, the size of production for each grape variety is shown using a bubble chart, which creates a metaphorical representation of a bunch of grapes.

You can clearly see how this design might have been conceived from early ideas formed before the data was even collected and analysed. Not only is the representation consistent with the subject, but it also offers an immediately recognisable metaphor. Any viewer will make a seamless connection between subject and form.

To unlock your creative imagination it is helpful to allow yourself to be inspired by the creative and created world that surrounds you. Exposing your senses to different sources of influence can broaden the range of creative solutions you might conceive.

By researching the techniques that are being used across the data visualisation field, and looking through the huge library of books that now exist, you’ll maybe encounter examples that demonstrate how others might have tackled visualising similar subjects or portraying similar types of analyses.

Beyond the visualisation sphere, consider other forms of creative design or imagery, from everyday encounters, whose aesthetic qualities you appreciate. Start a scrapbook or creative mood board that collects the colours, patterns, shapes and metaphors you’ve connected with. These references may spark ideas for the style, tone or essence of your upcoming projects. They might not have immediate value for all work but may prove useful for project contexts where there is an extra emphasis on creative expression.

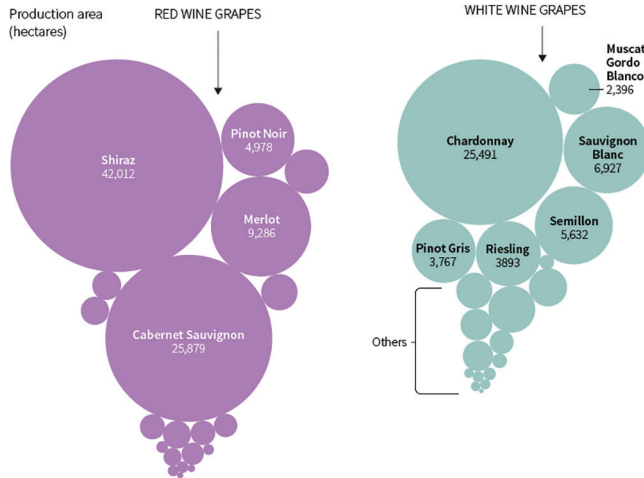
It is important to acknowledge some ethical matters surrounding all this. Influence and inspiration are healthy: the desire to emulate what others have done is understandable. Plagiarism,

‘I truly feel that experimentation (even for the sake of experimentation) is important, and I would strongly encourage it. There are infinite possibilities in diagramming and visual communication, so we have much to explore yet. I think a good rule of thumb is to never allow your design or implementation to obscure the reader understanding the central point of your piece. However, I’d even be willing to forsake this, at times, to allow for innovation and experimentation. It ends up moving us all forward, in some way or another.’ **Kennedy Elliott,**  
**Story Editor at the New York Times**

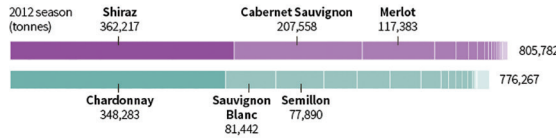
### The wine industry

Shiraz is the number one grape variety in Australia, covering 42,000 hectares of vineyard. Along with Chardonnay, the two varieties make up 45% of the total grape production.

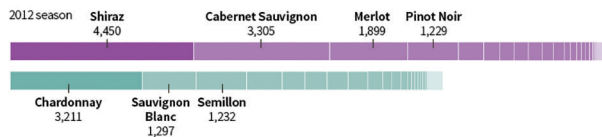
#### Top grape varieties grown



#### Production for wine-making



#### Businesses producing grapes - may produce more than one variety\*\*



**Figure 3.19** 'Grape Expectations', by S. Scarr, C. Chan and F. Foo (*Reuters Graphics*)

'I focus on structural exploration on one side and on the reality and the landscape of opportunities in the other... I try not to impose any early ideas of what the result will look like because that will emerge from the process. In a nutshell I first activate data curiosity, client curiosity, and then visual imagination in parallel with experimentation.' **Santiago Ortiz, Founder and Chief Data Officer at DrumWave**

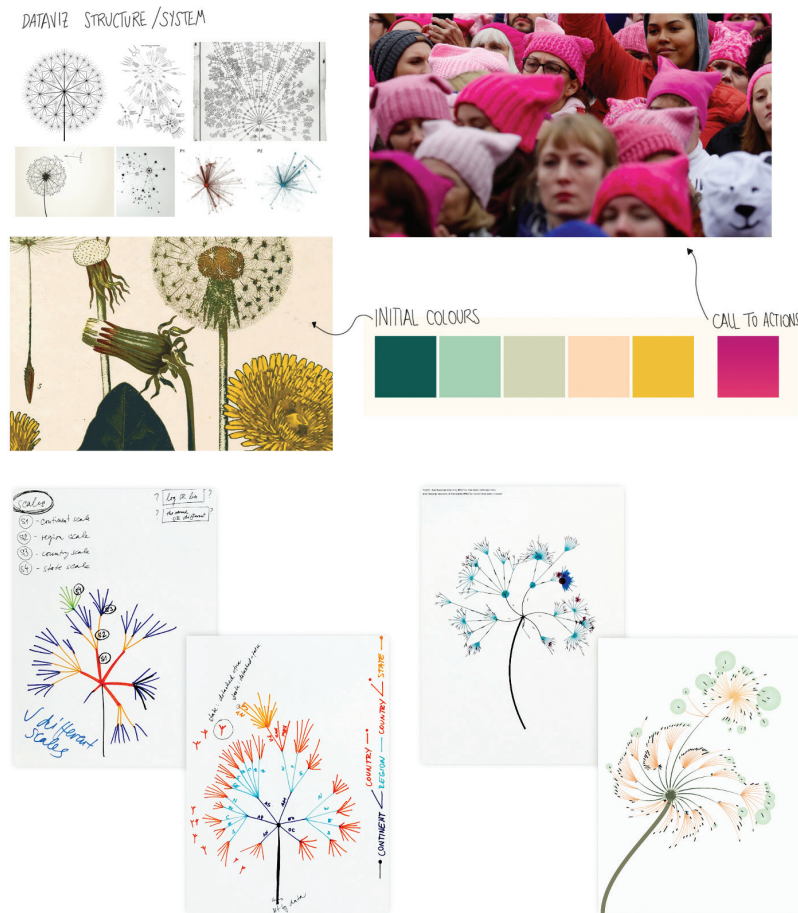
copying and stealing uncredited ideas is just wrong. There are ambiguities in any creative discipline about the boundaries between influence and plagiarism, and the world of visualisation design is not alien to that challenge.

Being influenced by the great work you see around the field is not stealing, but if you do incorporate explicit ideas in your work that have been influenced by others, at the very least you should do the right thing and credit the originators who will be unquestionably deserving of such attribution.

I've already explained the importance of getting into the habit of sketching and how it's not about possessing strong artistic talent. It helps extract ideas out of your mind and quickly capture them in visual form. Figure 3.20 shows a montage of various sketches and sources of inspiration that influenced the design concept of a project visualising the spread of the #MeToo movement.

For some people, the most fluent and efficient way to 'sketch' is through their digital application of preference, rather than on paper. Regardless of the medium you use, sketching is useful when you are working with collaborators as a means of discussing, generating

'Look at how other designers solve visual problems (but don't copy the look of their solutions). Look at art to see how great painters use space, and organise the elements of their pictures. Look back at the history of infographics. It's all been done before, and usually by hand! Draw something with a pencil (or pen... but NOT a computer!). Sketch often: The cat asleep. The view from the bus. The bus. Personally, I listen to music – mostly jazz – a lot.' **Nigel Holmes, Explanation Graphic Designer**



**Figure 3.20** 'MeTooMentum', by Valentina D'Elfilippo (design) and Lucia Kocincova (development)

‘I draw to freely explore possibilities. I draw to visually understand what I am thinking. I draw to evaluate my ideas and intuitions by seeing them coming to life on paper. I draw to help my mind think without limitations, without boundaries. The act of drawing, and the very fact we choose to stop and draw, demands focus and attention. I use drawing as my primary expression, as a sort of functional tool for capturing and exploring thoughts.’ **Giorgia Lupi, Information Designer and Partner at Pentagram**

‘Often when any kind of client or a potential client comes in with a preconceived notion or a solution, we try to challenge that... Sometimes it feels like it’s the natural fit, but sometimes we – annoyingly, probably for some clients – go, “Do you really need a map?” And they’re like, “Actually, no.” Or, “Actually, what are you talking about? Yes, of course we need a map kind of thing!”’ **Andrea Lau, Data Visualisation Specialist, Managing Director & Co-Founder at Small Multiples**

and critiquing ideas. I find it particularly helpful when trying to conceive innovative solutions to unusual or particularly complex challenges. It may be that my eventual solution looks nothing like my rudimentary sketches, but it gives me a way to cycle rapidly through iterations of concepts that may be worth exploring later on.

There are limits to the value of ideas and the role they should be allowed to play. Early sparks of inspiration in your thinking should be embraced, but do not be precious or stubborn. Always maintain an open mind and recognise that ideas have a limited role. This is why *harnessing* is the appropriate term used to describe this activity. Data will be your critical material. It may be that your ideas are ultimately incompatible with the properties and qualities of the data you are working with, in which case you should just let go and move on.

Finally, it is worth noting the diplomatic prospect of taking on board other people’s ideas. A recurring challenge of commercial work comes through working with clients who are unequivocally and emphatically clear about what they think a solution should look like from the very start.

Often your involvement in a project may arrive after these ideas have already been formed, during which time they already made it on to the brief issued to you by the

client (‘Can you make something like *this*, please?’). This is where your tactful ‘communicator’ hat comes to the fore. The ideas presented to you may be reasonable and well-intended, but it is your responsibility to lead the creation process. You can welcome input in the form of proposed concepts but, as with the limitations of your own ideas, there will be other factors with a greater influence: the nature of the data, the type of curiosities you are pursuing, the essence of the subject matter, and the nature of the audience, among many other things. These will be the factors that ultimately dictate whether any early vision of potential ideas ends up being of value.

## Summary: Formulating the Brief

This chapter commenced the opening stage of the design process concerned with initiating, defining and planning the requirements of your work.

## Defining Your Project's Context

The first section looked at context and identifying all the circumstances that will shape your project. These included factors such as the following.

### People

- Client(s): Who is the ultimate customer? Who are the influencers, interferers?
- Creator: What skills/knowledge are possessed by the visualiser?
- Collaborator(s): Working in a team, what blend of skills exist? Do you need to consult subject matter experts?
- Audience: How can you usefully classify the characteristics of the most relevant viewers/users?

### Deliverables

- Format: What is the communication medium: Print, digital, video, physical, presentation – all of the above?
- Quantity: How many and what range of different output assets will need producing?
- Delivery: Consumed remotely away from the creator or presented by a communicator? Rapid or prolonged duration?
- Frequency: One-off project? Will it become a regular/repeated task? Does it already exist or brand new?

### Constraints

- Timescales: When is it due? When can it commence? Any milestones? Duration of availability to commit?
- Pressures: Any factors or sensitivities around matters of financial, political, cultural, legal, environmental concern?
- Design: Any style restrictions (colour, type, logo)? Any size restrictions? Design techniques to avoid/emulate?
- Technological: What tools are available? Any hardware, infrastructure considerations? Platform compatibility?

## Establishing Your Project's Purpose

The second section considered the purpose of your work and the need to establish clarity about why you are doing it and what you are trying to accomplish.

Firstly, we looked at *what* understanding are you facilitating? This was framed around defining the motivating curiosity driving your work – to what data question(s) will your audience be offered visual answers?

Secondly, we looked at *how* will you facilitate understanding? Depending on your work's purpose – what will success look like? – you will need to find the most suitable balance in tone and the experience through which understanding will be facilitated:

- **Tone:** The difference between visualisations that place an emphasis on 'reading' vs 'feeling' data.
- **Experience:** The difference between 'explanatory', 'exhibitory' and 'exploratory' visualisations.

## Conceiving A Creative Vision

Finally, you learned about the value and limitations of harnessing ideas. What mental images, shapes, forms and keywords instinctively come to mind when thinking about the subject matter of this challenge? What influence and inspiration can you source from elsewhere that might start to shape your thinking?

## General Tips and Tactics

- Not all contextual factors can be defined this early, nor will they be stable throughout a process. Circumstances can change: certain undefined things at the start will emerge with a clear definition later; some things that start as being clearly defined will need reconsidering.
- Another reminder for the value of note-taking. They are so important to keep about any thoughts you have had relating to ideas of your curiosity, articulation of purpose, any assumptions, things you know and do not know, where you might need to get data from, who the experts are, questions, things to do, issues/problems, wish lists, etc.
- Keep a ‘scrapbook’ (digital bookmarks, print clippings) of anything and everything that inspire and influence you – not just data visualisations. Log your ideas and inspire yourself.
- This stage is strongly characterised by ambition management and it will be to your benefit if you treat it with the thoroughness it needs. The impact of any corners being cut here will be amplified later on. It can be tempting to race away with immersing yourself in your data, or your design thinking, but show patience: remember to attend to the right things at the right time.

### What Now? Visit [book.visualisingdata.com](http://book.visualisingdata.com)

**EXPLORE THE FIELD!** Expand your knowledge and reinforce your learning with further reading, example references and learning tutorials associated with this chapter’s content.

**TRY FOR YOURSELF!** Challenge, reflect and refine your skills by working on some practical exercises and activities linked to this chapter.

**LEARN FROM THE BEST!** Watch or listen to the episodes from the ‘Explore Explain’ podcast and video series, as Andy explores the visualisation design process followed by guest designers, developers and visual journalists. Hear from stellar talents from across the world as they explain the stories behind the inspirational works they’ve created. You’ll get insight into how today’s professionals navigate through all stages of contextual, analytical, editorial and creative activity outlined through this book.