CONTENTS

I

II

III

IV

Acknowledgements

This book would never have been written without Tapio Vapaasalo. *Neither would it ever have been finished without the hard work of our production editor,* Pia Alapeteri.

We are also deeply in debt to Alberto Cairo, *who has been amazingly supportive in our undertaking, as well as an inspiration.*

Annu Ahonen, Iida Turpeinen, Kimmo Vehkalahti, *and* Jussi Tuulensuu *have been immensely helpful in commenting the manuscript at its various stages.*

We want to thank everyone who agreed to be interviewed for the book: Fernando Goméz Baptista, Bonnie Berkowitz, Jonathon Berlin, Brian Boyer, Jen Christiansen, Ben Fry, Martine Hendriksen, Scott Klein, Miska Knapek, Vesa Kuusela, Charlie Loyd, Damon Burgett, Stefanie Posavec, Kim Rees, Brett Johnson, Jon Schwabish, Harri Siirtola, Robert Simmon, Yrjö Sucksdorff, Mikko Hynninen, Jan Willem Tulp, Ben Welsh, Sakke Yrjölä, *and* Javier Zarracina.

Likewise, we thank everyone who gave us the permission to feature their work in the book; especially Anssi Arte, Nadieh Bremer, Juan Colombato, Stewart Gray, Sami Heiskanen, Benjamin Hennig, Hannu Kyyriäinen, Viktor Landström, Johannes Nieminen, Tuomas Siitonen, Janne Pulkkinen, Pekka Veikkolainen, Hannes Vartiainen, Lauri Vanhala, *and* Aljaž Vindiš.

Paula Ahonen-Rainio, Keetos Balion, Otto Donner, Juuli Hurskainen, Tiina Koivusalo, Tommi Kovala, Esa Lehtinen, Eemeli Nieminen, Eeva Rautio, Ville Tikkanen, *and* Tuukka Ylä-Anttila *have helped us in other important ways.*

We wish to thank our friends and families for their patience and support.

FOREWORD

Every few days, we generate more data than had been produced since the dawn of history up to the year 2000. As Google's chief economist Hal Varian has stated, we "have essentially free and ubiquitous data … [now] the scarce factor is the ability to understand that data and extract value from it." In the end, data is valuable only to the extent that it can, in one way or another, be transformed into knowledge and wisdom.

Although artificial intelligence and other automated methods for analyzing data have seen impressive improvements during the past years, computers are not about to replace people as the final arbiter of meaning and relevance. For the foreseeable future, human beings—not machines—are going to be the end-users of data, and any insights that might arise from it. Consequently, the relevant question is: how can a person make sense of this deluge of data, and gain actionable insights from it?

However versatile and adaptable we are as a species, humans were not built for absorbing raw data. We have evolved to survive in an environment very different from the one most of us currently inhabit, engaged in tasks very different from combing through databases or spreadsheets for nuggets of information. Improvements in nutrition, health, and education notwithstanding, our brains and sensory systems have remained essentially unchanged for the 300,000 or more years our species has existed. This sets limits to our ability to absorb and process information.

Sensory stimulation is the only way for information from the outside world to enter our brain, and by far, vision is the most important of our senses. "The surface is the only way in," as professor emeritus Tapio Vapaasalo (our co-author in *Tieto näkyväksi*, the Finnish-language version of this book) has put it. Information design is the art and science of transforming data into visual structures, from which we can extract meaning—make comparisons, see trends, recognize patterns, spot exceptions and outliers, understand causal relationships, and much more.

As more and more activities become increasingly data-intensive, the ability to analyze data and communicate data-driven insights effectively become essential skills for people in a wide range of professions. Although graphic designers, communicators, and all those who are professionally involved in information dissemination have an important role to play, data-driven communication cannot remain the exclusive purview of such professionals. Instead, we believe it is highly useful—maybe even necessary—for all knowledge workers to understand the basics of information design, to be able to create clear, compelling, and insightful presentations of data—for their own use or for informing others.

Information graphics shape our understanding of reality. Charts are seen as authoritative, and a well-made chart can help convince people to change their views on important societal issues. Maps, which were once rare artifacts, mainly used by specialists, are now regularly perused by most smartphone users—a group of people numbering some 4 billion in 2019. Formerly seen mainly as as an implement of science and engineering, visualization is today applied to a myriad of different uses in fields as disparate as sports and agriculture.

Yet understanding such displays of data is not an innate skill, but one that has to be learned and honed. Visualization literacy should, in our opinion, not be seen as limited just to the ability to read charts and maps, but also to include the ability to make use of tools readily available for creating them. Taking control of chart- and map-making can be an agent of democracy for common citizens, especially those belonging to marginalized groups, to better participate in public discourse previously governed by an elite group of highly trained professionals.

This book is the culmination of a decade of teaching, writing about, and professionally pursuing information design. We hope it will provide a good overview of this varied and ever-evolving topic, and prove useful and enlightening for anyone wishing to communicate data—whether they are design professionals, analysts, journalists, scientists, civil servants, or practitioners of any other profession.

And above all, we hope you enjoy reading this book as much as we have enjoyed writing it.

In Helsinki,
January 6th 2019
The authors

PART

I

INTRODUCTION

A still from the animation *The next big spill – the Baltic Sea traffic visualized* (2013), created by Lauri Vanhala for HELCOM (Baltic Marine Environment Protection Comission).

INTRODUCTION

> ”The purpose of visualization is insight, not pictures”
>
> — Ben Shneiderman

Humankind has been creating visual documents of its surroundings for tens of thousands of years—long before the invention of writing. Alongside naturalistic imagery that attempts to capture what can be seen with our eyes, other visual representation methods, such as maps, were developed early in human history. Although maps resemble miniature images of terrain, the purpose is not to accurately reproduce a bird's-eye view of the landscape. They are conceptual representations based on established visualization conventions such as a scale factor. Developments in science, technology and society have brought about a need to present visually things that are not normally visible. Our eyes cannot directly see bodily functions or social and economic structures, but when presented in visual form, they are easier to understand.

According to the visualization researcher Colin Ware, most of our thinking happens as a kind of interaction with a variety of methods and tools that enhance our brains' data processing ability. Rapid technological progress constantly produces new, better tools, alongside traditional methods, such as drawing and writing—the most recent of which are computers and smart devices. In almost all fields of work, thinking is "carried out through distributed cognitive systems" in social networks, in which information is often processed by a large number of people.[1] Interaction is communication, communication is thinking, and visualization is in many cases the most effective way of communicating information.

---

[1] Ware 2004/2013, p. 2. Ware is here building on the work of anthropologist Edwin Hutchins, especially his book *Cognition in the wild* (1995).

Despite the development of tools, methods and processes, the fundamental laws governing the visual presentation of information have not changed. Our eyes and brain still function in the same way as they did for the first humans, hundreds of thousands of years ago. Understanding the capabilities and limitations of the visual system helps to identify a suitable visual presentation method for each communication need. The selection of the presentation method and a layout for the information, taking into account the capabilities and limitations of the visual system, is called **information design**.

Every knowledge worker's job description includes situations in which data has to be converted into visual form. Common tasks include giving presentations and creating presentation graphics and charts using spreadsheet software. In scientific publications, visualizing key research results in some form is a basic requirement. It is hard to imagine a modern school without educational materials, created by teachers or by textbook publishers, that use a variety of different graphic presentations.

The data visualization handbook is written for journalists, as well as graphic designers and other visual professionals, but also for people working in academia, education, public relations, government, and politics, who want to use information graphics and visualizations to support communication, analysis or decision-making. As the name suggests, the book can be used as a practical guide to creating visual presentations, but it is also suitable for use as a textbook.

The book introduces the reader to the general principles of information design, the limitations and strengths of human perception, and the requirements they place on visual presentation methods. The book explains the various genres of data visualization and includes a chapter on typographic issues related to visualization. The book ends with a discussion of workflows, ethics, and good professional practices in information design.

Vision is the strongest human sense

Humans and other apes have exceptionally well-developed eyesight. This distinguishes us from most other mammals, which rely much more on smell and hearing in perceiving the world around them. More than a quarter of the cells in our cerebral cortex are specialized for processing signals from the eyes.[2]

[2] Van Essen 2004

In fact, the human visual cortex is larger than all the other parts of the brain used for processing sensory information combined.

We are able to visually perceive our environment, and changes in it considerably more quickly and precisely than using other senses. It is estimated that, each moment, our visual system sends our brains around eight times more information than all the other senses combined.[3]

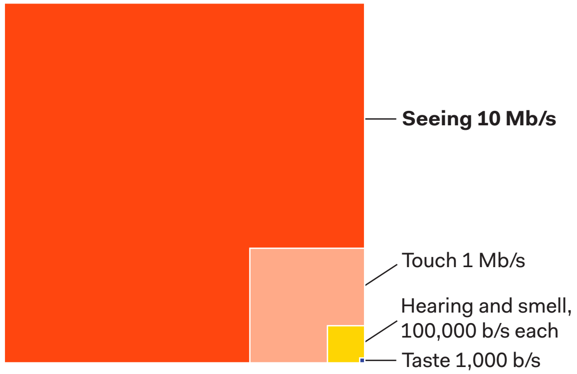

The amount of information processed by the sensory system cannot be measured directly, so such an estimate is inevitably imprecise. It is also worth pointing out that only a tiny fraction of the information transmitted by the senses reaches the conscious mind. If we look only at what is processed consciously, sight is not quite so overwhelmingly superior compared to the other senses—especially hearing. It is nevertheless obvious that vision is the strongest of human senses, and new information is usually adopted most readily when presented in visual form.

The supremacy of sight over the other senses is easy to see in many everyday situations. When we walk through a door, we can instantly form an overview of the space we have entered: how large it is, where the doors, windows and furniture are located, what surface materials have been used, whether there are other people or animals in the room, and so on. Forming a similar overview relying on the other senses would be much slower. It is no coincidence that the languages we speak include many idioms and expressions in which knowledge formation is described in terms related to seeing: *I see, an overview, a vision, to get the picture*.

The statistician John Tukey has said that "[t]he greatest value of a picture is when it forces us to notice what we never expected to see."[4] When shown in visual form, the data often reveals features that would remain hidden in a text or a table.

[3] Zimmermann 1989

[4] Tukey 1977, p. vi. Emphasis removed.

Vote shares in European Parliament election 2014 and referendum results by precinct in Denmark

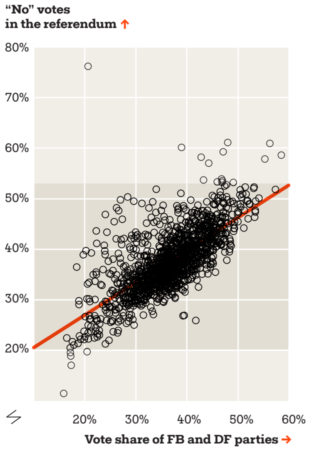

Source: Open Knowledge Danmark 2014. *Valgdata workshop*

In 2014, Danish voters cast their votes both in an election for the European Parliament, as well as in a referendum on whether the country should join the pan-European Unified Patent Court. Immediately after election day, activists belonging to the non-profit Open Knowledge Denmark analyzed the vote. In the adjacent scatterplot based on their work, each dot shows one of the 1,396 voting precincts in Denmark. Their position along the horizontal axis shows the percentage of votes received by the two parties, People’s Movement against the EU (Folkebevægelsen mod EU, FB) and the Danish People’s Party (Dansk Folkeparti, DF), which advocated a “no” vote in the referendum, and the vertical axis shows the share of “no” votes cast.

Look at the figure for a moment. Do you notice anything odd?

The points cluster quite densely around the diagonal, which means that there is a strong correlation between the two variables. This is not surprising: in precincts where a high number of votes were cast for the parties that opposed joining the Court, a high number of “no” votes were cast, and vice versa. The figure, however, reveals one precinct that is completely different from the others. The Taarbæk precinct shown in the upper left corner stands out dramatically from all the other dots in the figure.

When the issue was investigated, it was found that election officials in Taarbæk had, when reporting the result to the central electoral board, inadvertently reported the “yes” and “no” votes cast in the referendum the wrong way around. Despite multiple checks, no-one noticed the error, because the results were processed in the form of numbers in a table. However, when the data is converted into visual form, the outlying data point **cannot *not* be seen**. We can instantly see that one data point stands out in some way from others.

It is worth remembering, however, that a visualization can only show us certain features in the data, therefore obtaining a more detailed explanation for the deviation requires familiarity with the original tables and texts. Image and text are subsequently mutually supportive, not substitutive, means of communication.

In his book *Thinking, fast and slow*, psychologist Daniel Kahneman, a Nobel laureate in economics, describes two systems that govern our thinking. System 1 is fast and intuitive, while system 2 is slow and analytical. Visual communication particularly supports the fast system 1 thinking, while language and other conceptual structures, such as mathematics, support the slow system 2 thinking. Both systems have their strengths and weaknesses, and often the best result is obtained by using both in conjunction.

Visualizations have been described[5] as a mixture of image and text. Indeed, a well-designed visualization supports both fast and slow thinking. Visual elements help viewers to understand the structure of data and to form a quick, intuitive overall picture of a phenomenon, even a complex one. The strength of text on the other hand, lies in conveying either precise, abstract, and analytical information. It focuses on details and enables a more thorough interpretation.

Using different cognitive systems in parallel helps in processing, internalizing, and memorizing information. Although images alone are recalled better than text, what is recalled best are compositions that include the same information in both visual and verbal form.[6] Accoring to the **dual coding theory**, developed by the psychologist Allan Paivio, this is due to the fact that visual information is stored as images in one part of the brain, the non-verbal part, and verbal information as concepts in another, the verbal part. A message is dual coded when it is perceived as both words and images, that is, processed in both systems. The resulting memory trace becomes stronger than when the message is perceived in just one of the two ways.

What is information design?

Information design is about presenting information in the clearest way possible. According to the definition suggested by the researcher Robert E. Horn,[7] information design is “the art and science of preparing information so that it can be used by human beings with efficiency and effectiveness.” Clarifying the structure of data presented in written form or in a table also falls under this definition, but the concept refers above all to creating visual displays of information. Information design consists of selecting, organizing and presenting information, taking into account the needs and characteristics of the selected target audience, and the context of use.

Is graphic design information design? This is a legitimate question as graphic design—or visual communication design, as it is also known—is in essence about designing the **visual presentation of information**. When the graphic design firm

[5] E.g. Huovila 1996

[6] E.g. Paivio 1991, Levie & Lentz 1982, Atkinson et al. 1999

[7] Horn 2000

Pentagram first coined the term *information design* in the 1970s to describe its work, the intention was to highlight the fundamental purpose of graphic design as supporting information and communication. With the establishment of *Information Design Journal* (1979) however, information design has been established as its own area of expertise, separate from graphic design.

An example of the divergent aims and methods of graphic design and information design, which is worth mentioning here, are the problems related to pharmaceutical packaging design. According to a survey by a Finnish pharmacists' trade magazine, patient safety is compromised –or at risk of being compromised– on a weekly basis in 30% of pharmacies, due to packaging for different medicines being too similar.[8] One of the reasons for this concerning situation is probably that the graphic or packaging designer responsible for the look of the package usually focuses primarily on the consistent application of the manufacturer's brand identity, instead of first considering the visual distinctiveness of packages in different use scenarios. The perspective of information design is different: brand and esthetics are always subordinate to the communicative function of packaging or graphics.

Terminology and typology

The origins of information design and data visualization lie in several disparate fields. The methods and practices discussed in this book originate in a number of differing disciplines: graphic design and illustration, journalism, cartography, computer science, business administration, and statistics, among others.

A side effect of this diverse heritage is that the terminology used by different visualization practitioners, authors, and researchers varies—at times widely. Different terms are sometimes used to describe the same thing, or the same word to describe different things.

Because of this, we have paid special attention to terminology in this book. We have tried to be unambiguous in our choice of words and give our definitions of specialist terms when they first appear in the text. Some of the principal terminology is outlined below.

+++

In his book, *Good charts*,[9] Scott Berinato presents a typology of charts based on two fundamental questions about each chart:

1. Is the information **conceptual** or **data-driven**?
2. Is the purpose to **declare** (or explain) or to **explore** the information?

Although Berinato mostly discusses the typology in terms of business graphics, we believe it can be applied to any type of visual presentation of information. We call the two dimensions **conceptual–measurable** and **explanatory–exploratory** (which we discuss below).

The structure of a supercell cloud

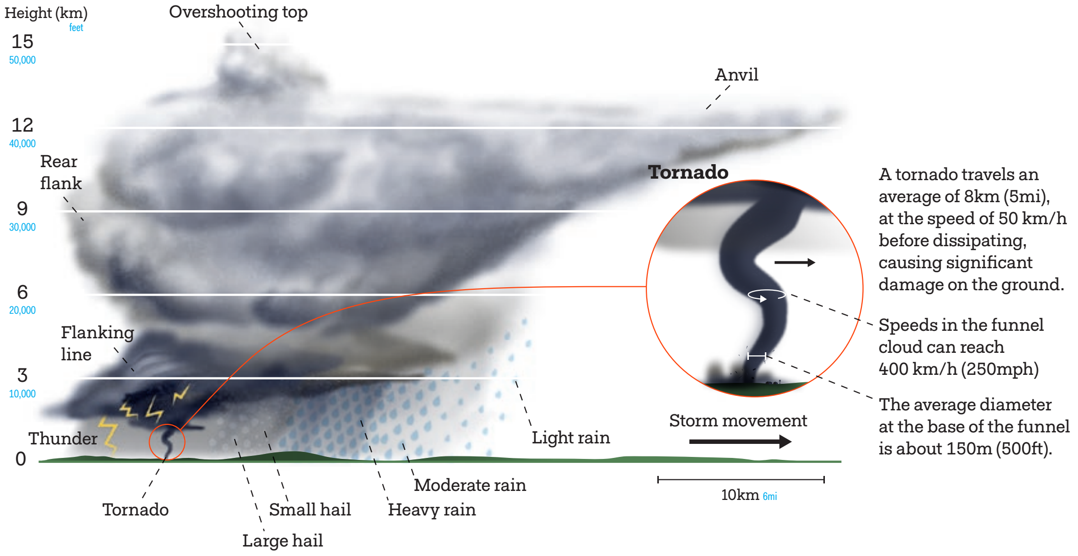

This explanatory graphic, adapted from a high school geography texbook, shows the structure of a supercell, a type of a storm cloud. The choice of data presented here derives directly from the high school curriculum, and the design has been chosen to best convey that information.

The primary purpose of **explanatory** graphics is to communicate information between people. They are used to declare, explain and affirm the facts. The creator of the graphic already knows the information at hand, and the primary design challenge is to find a way to convey that knowledge to the audience.

In contrast, the purpose of **exploratory** graphics is primarily to facilitate discovery and analysis of the information. They can be used in communication, but the primary aim of exploratory graphics is not to convey a message that has been determined in advance by the creator of the graphic. Their function is to act as a tool that enables the reader to find interesting features in the data,[10] and the creator of the graphic does not know in advance what the visualization will reveal.

---

[8] Kairenius 2012

[9] Berinato 2016, pp. 54–63

[10] See Cairo 2016, p. 31

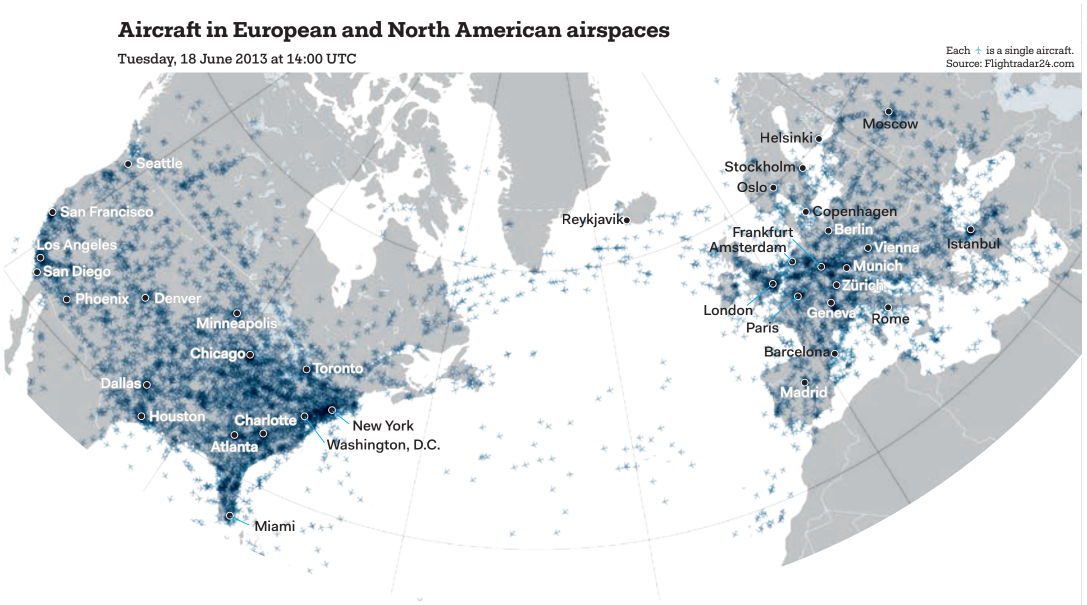

This exploratory graphic has also been adapted from a geography textbook. It shows the positions and bearings of the aircraft flying in the European and North American airspaces at a specific time in 2013. Besides showing the incredible number of simultaneously airborne planes, the visualization reveals interesting patterns. Not surprisingly, densely populated areas get more air traffic, but there are also exceptions to the rule. For example, Denver, a city in the relatively sparsely populated Mountain West part of the United States, stands out on the map thanks to its airport serving as a hub for several airlines.

Exploratory graphics are usually created with a computer and they are often, but not always, interactive. According to the typology outlined above, the information shown tends to be **measurable**. The graphic is based on a *scheme*—an agreed-upon set of rules for converting the information into visual form, as opposed to hand-crafting each element, which is common in explanatory graphics.

The difference between the categories is that an explanatory graphic *tells a story*, while an exploratory graphic is a tool for the reader to *find their own story* in the data.

There is a longstanding tradition of using the term **infographics** for explanatory graphics and the term **visualization** for exploratory graphics. This is for example how Alberto Cairo, a foremost authority on the topic, defines the terms in his book *The functional art*.[11] However, this distinction is less than universal, and the word infographics is also used to mean other things, such as to describe what in this book are termed *combination charts* (see p. 93).[12] Sometimes the term is also used to describe graphics showing **conceptual** as opposed to measurable information (according to our typology).

[11] Cairo 2012, pp. xv–xvi & 18. See also Cairo 2016, pp. 27–40.

[12] As does, indeed, Alberto Cairo in his more recent book *The truthful art* (2016).

In this book, we consider *visualization* to be a catch-all term for all graphics showing **measurable** information, whether explanatory or exploratory, and *infographics* to be a partially overlapping term that encompasses all **explanatory** graphics. The hypernym, or generic, overarching term for both is *information graphics*.

Whatever terminology is used, the division into explanatory and exploratory graphics is not black and white. They are better viewed as a continuum, rather than two clearly defined categories. Graphics often have both explanatory and exploratory features. These categories should therefore be considered as two approaches to presenting information, with both common and divergent features.

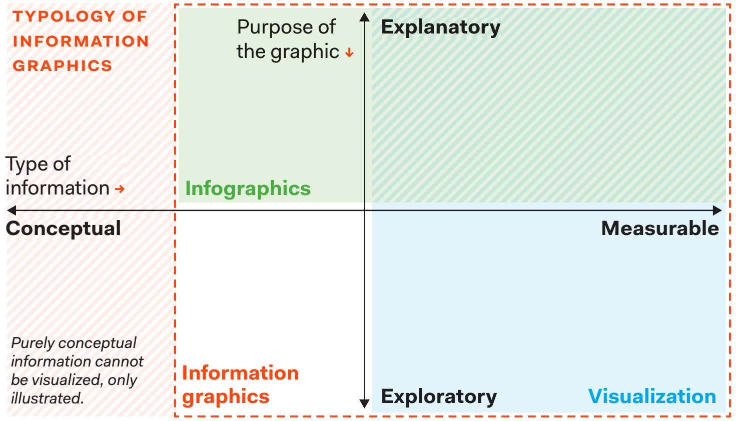

The visualization researcher Robert Kosara has suggested[13] the following definition for a visualization (exploratory, in particular):

- It is based on (non-visual) **data**,
- it produces an **image**, and
- the result is **readable and recognizable.**

Visualizations are therefore based on data that is abstract in nature, such as statistics, or that is not visible under normal circumstances, such as the structure of internal organs. An image that directly imitates visual perception, such as a landscape painting, is not a visualization.[14] The main result of the visualization process is always an image—though it can be supplemented with explanatory text for example. The image should be recognizable as an information graphic, and the information it contains should be presented as unambiguously as possible. This definition excludes, for example, data art (see next spread).

+++

[13] Kosara 2007b

[14] A representational image may, however, be *part of* an information graphic. It has been said (for example, Sullivan 1987, p. 41) that the simplest form of information graphic is a photograph with an arrow or a circle marked on it, indicating a feature of interest in the image. Some types of information graphics (see pp. 119–125), such as artist's renderings and identification pictures, again, come very close to a style that directly imitates visual perception, and therefore do not meet the definition presented by Kosara. They can be considered information graphics, but not visualizations.

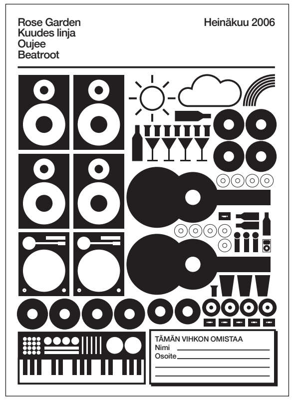

The difference between an infographic and an illustration is one of substance and function, not of style. The illustration on the cover of this brochure has a visual style similar to many infographics, but it does not convey any facts. What it aims to do instead is to convey a feeling or an impression.

Not all visual communication aims to convey information

Information design, then, is about presenting information in the clearest possible way. Nevertheless, not all communication is about conveying information. Messaging can aim to convey feelings, moods, values, or other less tangible things instead. If the aim is not to communicate concrete data, information graphics are the wrong tools for the job. According to Alberto Cairo, the purpose of an infographic or visualization "is not to make numbers 'interesting,' but to transform those numbers (or other phenomena) into visual shapes from which the human brain can extract meaning."[15]

This does not of course mean that information graphics could not *also* function as a visual element that adds color to a page, evokes a feeling, and attracts the reader. These tasks are however subordinate to the main purpose of information graphics as a tool for conveying of information. (See pp. 101–103.)

In the minds of many communicators, information graphics are regrettably often confused with illustrations. The concept is quite loosely defined. The word "illustration" may also refer to one of the genres of information design (see information illustration, pp. 119–135), but here the term refers to a type of image which Harold Evans, a pioneer in research on visual journalism, calls *flavor graphic*—differentiating it from *fact graphic*, that is information graphics and visualizations.[16] Defined in this way, illustrations seek to communicate not concrete facts, but things that are more difficult to identify, such as emotions and values, or simply to attract people to read the publication. In journalism—especially in periodicals—illustrations have been a key part of the visual style since the emergence of modern magazines, and they have a clearly defined role, which differs from that of fact graphics.

Data art in turn, uses data as its starting point, but deals with that data as material for artistic expression. Although it often applies presentation methods that resemble those used in visualization, data art does not seek to provide an unambiguous interpretation of data, but to create esthetic experiences and to test new methods of presentation. The new presentation methods that emerge within data art are often later adopted as part of the toolkit of visualizations.

The boundary between data art and visualization is not always clear-cut, but as a rule of thumb, visualization seeks to explain the structures in data, while data art takes the data and creates new kinds of visual structures from it. If then the presentation method is decided based on data, it is a visualization, but if data is picked to fit a chosen presentation method, it is data art.[17]

[15] Kuenn 2013

[16] Salo 2000, p. 155

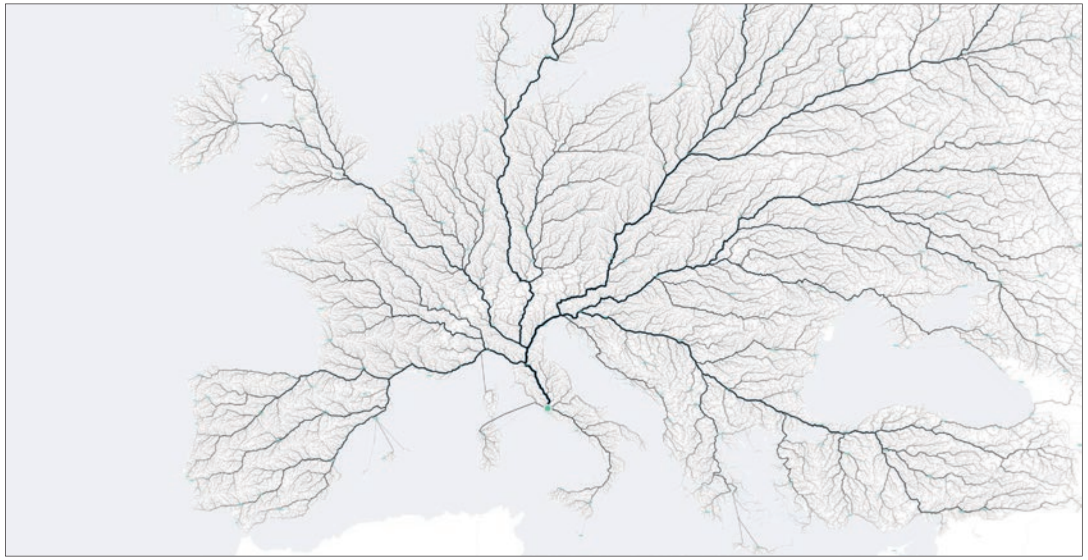

Roads to Rome (2013) is a data art project that explores the question of whether "all roads lead to Rome," as the saying goes. Benedikt Groß, Philip Schmitt, and Raphael Reimann created a computer program which used OpenStreetMap data to calculate the shortest route by road to Rome from each point in the network. The end result resembles a web of blood vessels branching out from the heart of the ancient Roman Empire all the way to its periphery. The group has since utilized the same technique in more practical applications, such as for analyzing urban structures. roadstorome.moovellab.com

When should data be presented in visual form and when not?

When should data be visualized, and when is text enough? In short, if something can be expressed just as clearly or even more clearly in words, visualization is unnecessary—sometimes even counterproductive. In 1801, the inventor of statistical charts, William Playfair, wrote that a visual presentation "gives a simple, accurate, and permanent idea, by giving form and shape to a number of separate ideas, which are otherwise abstract and unconnected."[18] When this is indeed the case, the data usually should be visualized instead of using just text.

The success of a visualization is ultimately determined by whether its form and shape help the reader to perceive the data better, or whether it just adds an extra layer of code to be deciphered. **A figure is clear when the viewer understands what it shows and, through it, finds answers to questions or gains insights about features of the data.**

If these criteria are not met, the graphic is unsuccessful. The reason for this may be that the choice of presentation method

[17] It is not possible to discuss data art more extensively in this book, but for those interested in the topic, we recommend Casey Reas and Chandler McWilliams' book *Form+Code* (2010).

[18] Playfair 1801/2005 p. 30

is wrong or that it is poorly executed. Information graphics work best for showing spatial and geographical relationships, processes, chronologies and above all numerical data. The visualization of ideas, values, and abstract concepts is much more difficult—often even impossible. To paraphrase philosopher Ludwig Wittgenstein:[19] **Everything that can be shown can be shown clearly**. If a graphic seems hopelessly unclear, the reason can be that it attempts to convert into visual form something that cannot in fact be visualized.

The graphic below comes from a presentation discussing the United States Department of Defense Architecture Framework. It is an example of a visualization that is hopelessly confusing, because the topic is not one that can be fully converted into visual form. Blogger Paul Ford describes it as follows: "… this image could be used *anywhere* in *any* paper or presentation and *make perfect sense*. This is a graphic that defines a way of describing anything that has ever existed and everything that has ever happened, in any situation."[20] (It is an example of what cartographer Jacques Bertin calls a *pansemic* graphic: "In its attempt to signify 'everything' it no longer signifies anything precise.")[21]

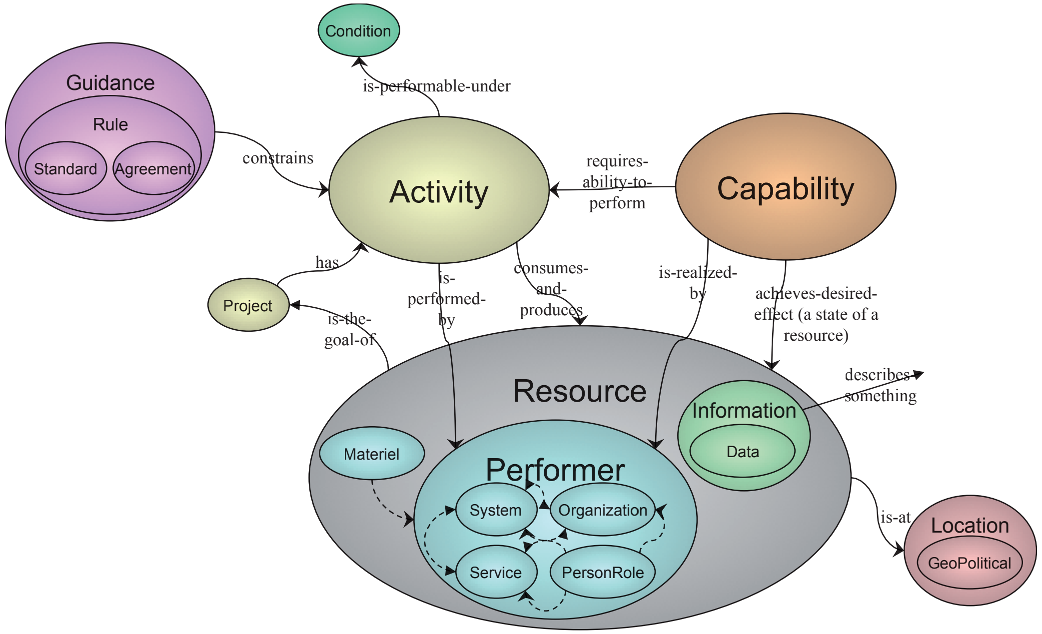

---

[19] Wittgenstein 1922, p. 45

[20] Ford 2014. Emphasis in original

[21] Bertin 1983/2011, p. 2

The grammar of information graphics

According to the computer scientist Leland Wilkinson, a language consisting only of words and no grammar expresses only as many ideas as there are words.[22] Grammar enables the formation of multi-word expressions and dramatically expands the scope of a language. (Almost) anything can be expressed by combining words, in accordance with the rules of grammar.

Seeing is *pre-lexical*—we usually identify what we see immediately without linguistic interpretation. For example, research[23] has shown that humans can reliably identify other individuals merely based on their *gait*, that is, the way they walk. Few of us however, are able describe someone's gait in any detail to another person. Visually identifying something or someone does not require the viewer to be able to describe what they see verbally.

In information design, our ability to intuitively understand what we see is crucial. The visualization of data helps us to identify features in the data that we may not necessarily be able to name. We do not need to know what words such as *correlation* or *outlier* mean to be able to visually identify the phenomena that they describe in a scatterplot for example.

The pre-lexical understanding of what is seen is nevertheless only the first level in interpreting perception. Understanding complex data, irrespective of the presentation method, always requires active interpretation, which is a much slower process than forming an initial overview based on perception.[24]

Symbols and glyphs

In his book *The design of everyday things*, Donald A. Norman, a pioneer of user-centered design, distinguishes between **additive** and **substitutive** dimensions in user interfaces.[25] Additive values such as length can be changed incrementally and are used to indicate varying *amounts* of something. Substitutive values such as color hue[26] cannot be changed in size but only substituted for another value and are used to indicate *categories*.

This basic logic is an example of what Norman calls *natural mapping*: "taking advantage of physical analogies and cultural standards" to visually encode abstract information. It is easy to understand that elements that may fluctuate in size are used to show measurable, quantitative values, while things that can only be substituted by one another correspond to differences in kind,

---

[22] Wilkinson 1999/2006, p. 1

[23] E.g. Stevenage, Nixon & Vince 1999

[24] According to Zimmerman (1989), perception is mostly processed unconsciously, and only a very small part of all sensory information is consciously processed: 40 bits per second for sight, 30 b/s for hearing, 5 b/s of for touch, and only 1 b/s for taste and smell.

[25] Norman 1988/2002, p. 23. Additive and substitutive scales were originally identified by psychologist Stanley Smith Stevens (see Norman 1991).

[26] Hue means the shade of a color, such as red or green. See p. 66.

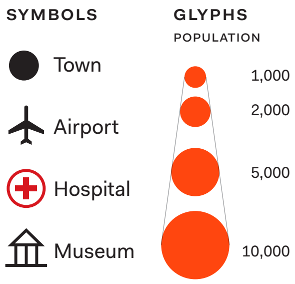

namely categories. We intuitively understand that a higher bar or a bigger circle indicates a greater number or amount when compared to a smaller element, and that a different-colored bar differs somehow from the other bars in a group. The interpretation of additive features is largely intuitive and based on the operational logic of lower-level visual perception (for more information, see *visual variables*, pp. 58–62).

Visual elements that include only substitutive dimensions are called **symbols**. According to Colin Ware, they represent an object through a simple "this is X" relation.[27] The expressive power of a visual code that only uses symbols is rather limited. In the vein of Wilkinson's definition presented above, it can be viewed as a *grammarless language*, in which all possible expressions and their meanings must be defined beforehand.

Visual elements that include at least one additive dimension are called **glyphs**.[28] An example of a glyph is a bar in a bar chart.

The meaning of a data visualization emerges above all from the **interrelationships** between symbols and glyphs, not from these parts themselves.

+++

The cartographer Jacques Bertin, who developed much of the theoretical foundation of information design in the 1960s, divides visual signs into two main categories: monosemic (having a single meaning) and polysemic (having multiple meanings). The meaning of monosemic signs such as mathematical symbols is given—known in advance—and their interpretation is therefore unambiguous.[29] The meaning of polysemic signs, such as the individual details that make up a drawing, follows from the other signs that are present: a circle may denote very different things from one picture to another. Polysemic signs are read "between the sign and its meaning," and thus their interpretation is always ambiguous and debatable to some extent.

According to Bertin, information graphics are read "among the given meanings" of monosemic signs.[30] The meaning of each element in isolation is known in advance thanks to legends, color keys, and so on, but most of the actual information in the graphic is embedded in the **relationships** between the signs—symbols and glyphs. This is called **induction:**[31] the reader is able to induce from a graphic much more than merely the individual pieces of information its designer has specifically included in it.

[27] Ware 2004/2013, p. 140

[28] Ward 2008. In visualization research, the term glyph is often used specifically to mean multivariate glyphs: visual objects that are used to show several data dimensions at once. Probably the best-known example of glyphs in this narrower sense are the so-called Chernoff faces (see, e.g., Kosara 2007a). Here we will use the term mostly for simple glyphs, which form most of the basic building blocks used in more complex visual presentations of data.

[29] Bertin 1983/2011, p. 2

[30] Bertin 1983/2011, pp. 2–3

[31] Robinson et al. 1995, p. 451

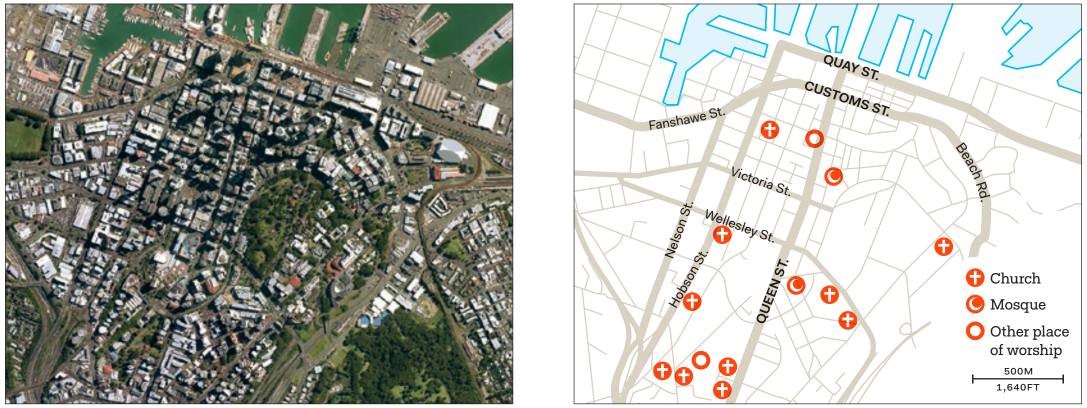

The aerial photo and map above both show the same area in Auckland, New Zealand. The photo includes much, much more information than the map. Most of this information is however useless in the majority of use cases. By leaving out most of the information in the photo, the map becomes more useful in its primary purpose of locating religious buildings in the area. For some other purposes however, such as differentiating between residential and industrial areas, the map is useless.

Simplify, compare and organize

Three factors above all define how easy or hard a visualization is to read and what insights can be gleaned from it: how the data is simplified, compared and organized.

Simplify

When irrelevant details are left out from a visualization, the remaining elements are easier to identify and their mutual comparison is easier. It is believed this is due to decreased need for visual decoding, which frees up more working memory capacity to process relevant information.

Consider maps: although a map also includes information that is missing from a photograph—place names, for example—an aerial photograph almost always includes vastly more information about the terrain than a map of the same area. Most of this information is useless to the map's user—the individual trees, cars, boats, roofing materials of buildings, weather, and the direction of light at the moment the photograph was taken, and so on. When such unnecessary information is removed, the information left in the map acquires greater visual prominence and reading the map becomes easier. It is often said that **the most important decision in making a map is not what to put in, but what to leave out.** The same is true of all information graphics.

It is a common mistake to cram so much information into a visualization that the eye is no longer able to pick out visual structures in the image. The information extractable from this crammed network diagram (see pp. 227–229)—a "hairball" in industry jargon—can be summarized in one sentence:

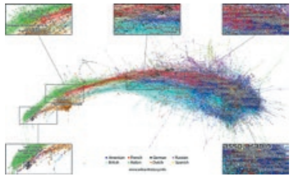

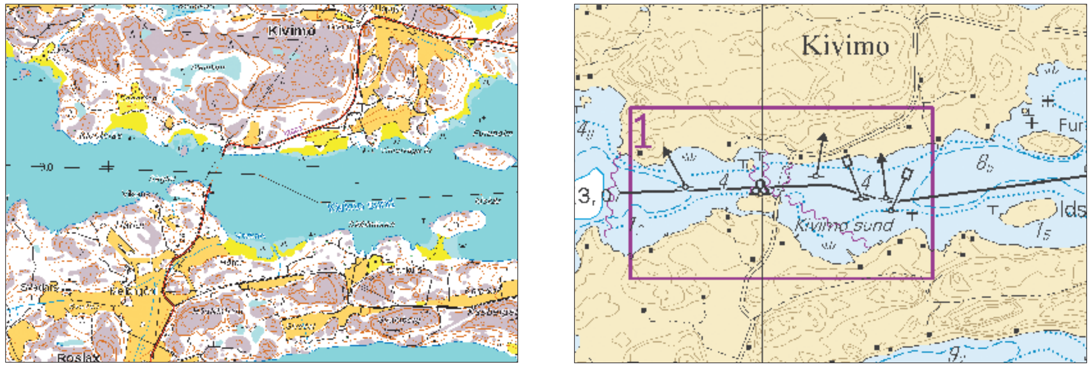

The makers of these two maps showing the same place in South-Western Finland, a "terrain map" on the left and a nautical chart on the right, have made very different choices in selecting what to leave out of their map.

"There are many connections." Something that can be expressed in one sentence does not deserve to be visualized.

On the other hand, data can also be too narrow or simple to be visualized. Because visualizations are based on the visual comparison, a single number cannot be visualized.[32] If the data only includes a few numbers or other facts, in many cases a few sentences or a single table can convey the information as well as, or better than, a graphic. Edward Tufte, Professor Emeritus of statistics at Yale university, has suggested as a rule of thumb that "[t]ables usually outperform graphs in reporting small sets of 20 numbers or less."[33] The creation of a graphic requires more effort than a text or a table, so in cases like this, it often makes sense to allocate limited resources to something other than creating information graphics.

It is worth noting that simplification is something you do to the underlying data, not to the graphic itself. Simplified data does not necessarily require a minimalistic style of presentation. A decorative graphic can be clear and illustrative, as long as the reader is able to easily identify which of the elements in the image carry information and which are just decoration. The decorative elements should never fight for attention with the actual content. Often however, a minimalist style is a safer choice than a decorative one. It means there are fewer opportunities for mistakes that compromise the clarity of the visualization, as

32. To visualize a single number, a designer can choose a familiar object or comparison level as a reference point; for example, how many football fields would fit inside an area, or how much the share of a specific demographic in a city's population differs from the national average. In this way, it is no longer a question of visualizing one but two numbers, which is very much achievable.

33. Tufte 1983/2002, p. 56. The rule should not be taken too literally. There are many examples of datasets of merely two or three numbers that benefit from being visualized.

On the left is an example of a data set that is too simple to benefit from being visualized. The interesting thing here is the most common value (18 years), and the exceptions form it (Iceland and Denmark), not the percentual difference between numbers 18 and 20, or 18 and 16. The same information would be easier to read if presented as a list or a table.

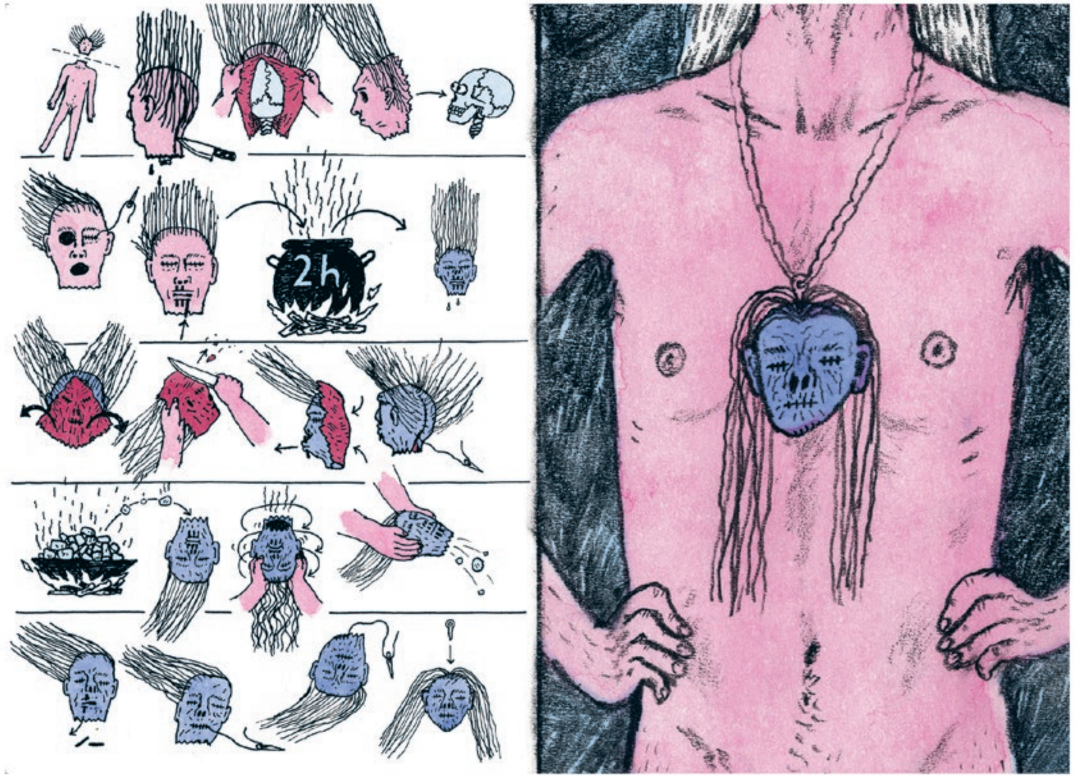

This infographic, created by illustrator Viktor Landström, explains the making of an Ecuadorean *tsantsa* (a shrunken head). The content is simplified, but the presentation style is anything but minimalistic.

it is unfortunately often the case that graphics are decorated at the expense of clarity. The importance of esthetic choices in the design of visualizations is discussed on pages 101–103.

Compare

Edward Tufte writes that "At the heart of quantitative reasoning is a single question: *Compared to what?*"[34] A visualization is at its heart, a tool for making comparisons. Thus the single most important question to answer when designing one is: **what comparisons should be enabled?**

Written and spoken language is linear in nature. Speech, text, and video proceed from beginning to end, at the pace and in the order determined by their creator. By contrast, data visualization hands control over to the viewer. They can explore the content at their own pace, only superficially or drilling down to details, and going through the elements in the any order they like. Someone reading a text or listening to a speech needs to rely on their working memory and mental calculation in order to be able to compare the numbers presented or the connections explained. Someone viewing an information graphic does not need to

34. Tufte 1990, p. 67

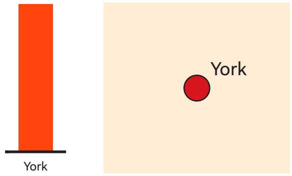

Meaningless bar chart and map.

engage in such brain-taxing activities, as a visual comparison of elements in the image takes place almost automatically.

On the other hand, text can use many different types of content structures, while an abstract visualization mainly just presents various **relationships between data points.**[35] For this reason, a single bar, map symbol or shape does not convey information, but only becomes meaningful by its relationship with other elements in the image—in other words, it is *polysemic*. As Vesa Kuusela, an expert on statistical charts, has said, "**A chart acquires its meaning from comparison.**"

Visual comparability requires that the visual encoding methods, such as colors, symbols and scales, are used consistently. (For more information, see pp. 83–86.)

A well-made visualization facilitates a fairly large number of different visual comparisons. Our visual system and cognitive capacity, however, set an upper limit to the complexity of presentations we can interpret. Because not all the methods available for visually encoding data are equally suitable (see pp. 58–62), the designer must also decide **which of the comparisons are the most important**. The best encoding methods, such as the positions of the glyphs, should be used for the most important comparisons; and the less effective ones, such as their surface area, to encode secondary or supplementary data.

+++

A visualization can show many types of relationships, which can be roughly divided into six main groups:

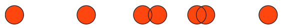

Numbers refers to either quantities (amounts) or sizes. Relationships between numbers are by far the most common type shown in information graphics. Most visualizations include at least some comparisons between numbers.

Rank (or ordering) is different from numbers only in that the hierarchical relationship of the data points—larger or smaller rank—is known, not the actual magnitudes or differences in values. (This has a variety of practical implications. See *ordinal scale*, p. 94.)

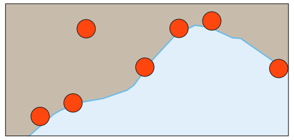

Location (or position) is the basic type of relationship shown on maps. Though it usually refers to geographic and astronomical position, location can also refer to locations such as those within the human body from which a biological sample was taken.

---

[35] This is not necessarily true of information graphics that visualize *conceptual* information. See pp. 219–225.

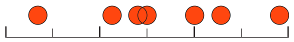

Time refers to relations between positions in time, such as dates. It is one of the basic relationships shown in time-series charts, such as line charts, and in timelines.

Category refers to some qualitative similarity between data points. For example, companies can be categorized by industry, or municipalities by region.

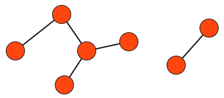

Connection refers to various links between the data points. The connections can be hierarchical (directed), as in a food web, or non-hierarchical (undirected), as in a network of friends. (For further information, see *network diagrams*, pp. 227–229.)

Qualitative data that does not include these relationships must, in order to be visualized, first be transformed to fit one of the above groups, for example by giving it a classification or score. An example of this is the *Martin–Quinn score*, developed by the political scientists Andrew D. Martin and Kevin M. Quinn. It is a scoring system for U.S. Supreme Court justices based on their voting record on the bench, giving each an ideological score ranging from –5 (extremely liberal) to 5 (extremely conservative). This conversion from a highly qualitative data set—the court's decisions—to numbers enables the data set to be analyzed using statistical methods and visualized.

The ideological spread of SCOTUS justices has slightly narrowed since the late 1970s

Martin-Quinn scores for Supreme Court justices

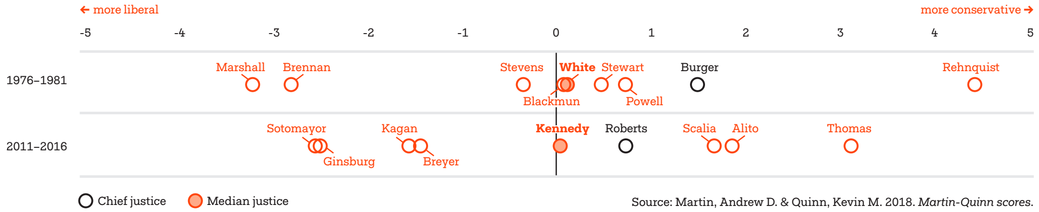

Source: Martin, Andrew D. & Quinn, Kevin M. 2018. *Martin-Quinn scores*. Center for Empirical Research in the Law, Washington University

Organize

Besides simplification and comparison, the third defining characteristic in how a visualization is read and understood is how the data is visually organized.

In his book *Information architects*,[36] Richard Saul Wurman, designer, author, and the founder of the TED conference, has proposed a set of organizing principles for information, which he calls the "Five Hat Racks," and provides the acronym **LATCH** as a mnemonic. Wurman himself has since disowned[37] this system, but it serves as a good starting point for an improved taxonomy.

---

[36] Wurman 1996

[37] Wurman & Grimwade 2016

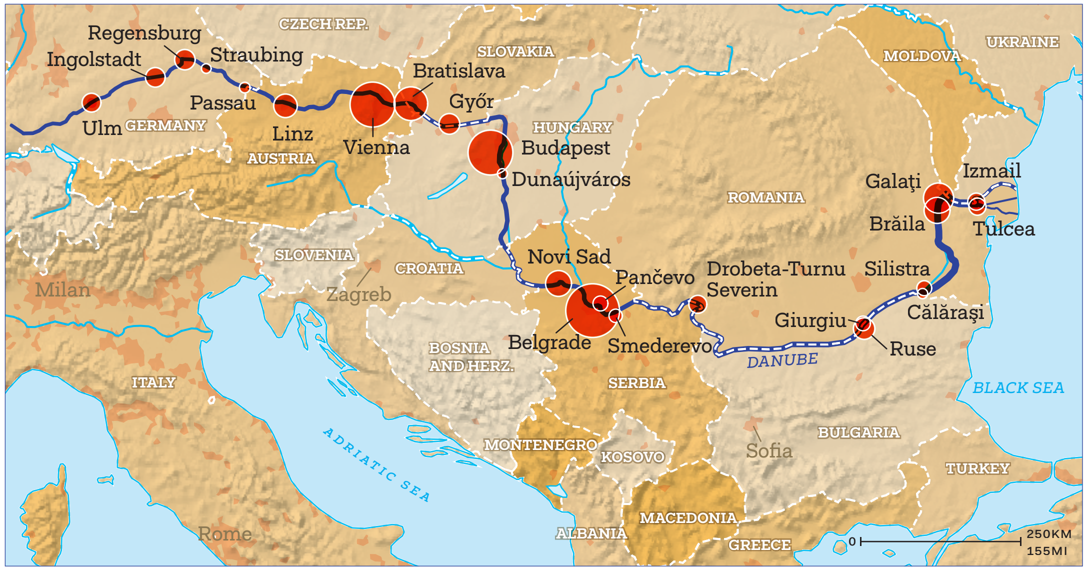

A map of cities located along the river Danube. The size of the bubble indicates the population of the city. On the next spread, the populations are visualized as a set of horizontal bar charts which have been ordered using each of the different organizing principles outlined in the text. Depending on the organizing principle used, each chart gives a different impression of the data and reveals different patterns in it.

Wurman's Five Hat Racks are:

- **Location**
- **Alphabet**
- **Time**
- **Category**
- **Hierarchy**

In an earlier book, *Information anxiety*, published in 1989, Wurman uses the term **continuum** instead of hierarchy. The term has apparently been changed mainly so that the initials of the organizing principles form the memorable acronym LATCH. Below, we use the original term continuum, as it better captures the nature of this organizing principle.

An attentive reader may notice that the organizing principles bear similarities to the comparable relationships in data described above. This is no coincidence. The relationships are closely related to the organizing principles, so that organizing the data in a particular way in a visualization strongly emphasizes one relationship in the data over others.

By making a few additions and clarifications to Wurman's Five Hat Racks, we can match each organizing principle with one of the relations, and create our own mnemonic—the **Seven C's**:

Continuum is the most typical way of organizing elements in a graphic. It is usually the most natural way of organizing data points in a visualization, unless there is a *specific* reason to do otherwise. Continuum has two subtypes:

ORGANIZATION PRINCIPLE

RELATIONSHIP	FIVE HAT RACKS	SEVEN C'S	BASED ON
Numbers	Continuum (Hierarchy)	Continuum by magnitude	Based on data
Rank		Continuum by rank	
Location	Location	Coordinates	
Time	Time	Chronology	
Categories	Category	Categories	
Connection		Connection	
(none)	Alphabet	Convention	Based on convention

Continuum by magnitude, in which the elements are ordered in a descending or ascending order based on the magnitude of their value, in other words, from the largest to the smallest value or vice versa. The particular advantage in organizing the data like this is that the values with the smallest differences are always located adjacent to each other, which enables very small differences to be detected.

Continuum by rank, in which the ordering is based on the structure of the scale itself. For example, if the data deals with educational attainment, it usually makes sense to order the data by the level of education from primary to secondary to tertiary level (continuum by rank), not the number of people at each level (continuum by magnitude).

Coordinates (location) is the organizing principle used in maps. It is less common in other types of information graphics, but can be used as an organizing principle in many types of statistical charts, for example.

Chronology is the organizing principle used in line charts and other time series. In a graphic created for a Western reader, time should usually be shown as moving from left to right.[38] In the domain of statistical graphics, the main exception to this rule are the graphics used immediately next to tables in financial statements, such as annual reports. In such tables, time can also be presented as moving from top to bottom or from right to left, and this same order can also be used for related graphics—excluding line charts, in which the direction of time should always move from left to right. In some visualizations other than statistical graphics, such as timelines, time can also be presented as moving from top to bottom.

[38] The direction of time in graphics is the same as the writing direction of the Latin alphabet. In the case of, for example, Arabic, Hebrew or other right-to-left languages, time can also be shown as moving from right to left in figures. In these languages however, practice varies more than in Western languages; Israeli newspapers feature time series that are ordered both from right to left and from left to right.

CONTINUUM BY MAGNITUDE

By population

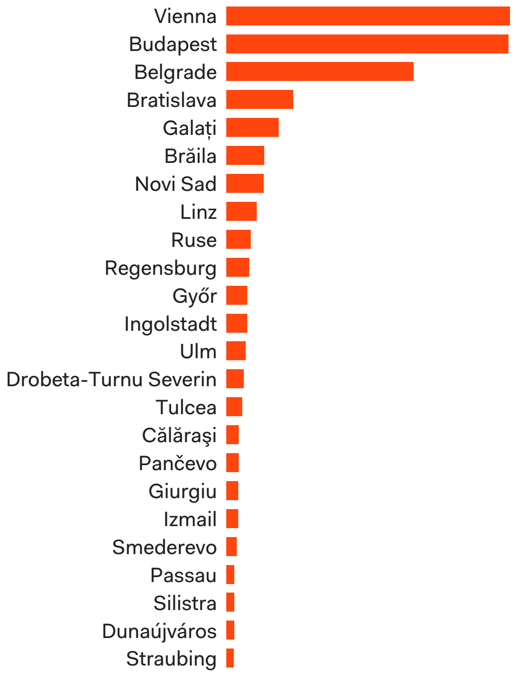

COORDINATES

By location along the Danube downstream from Ulm

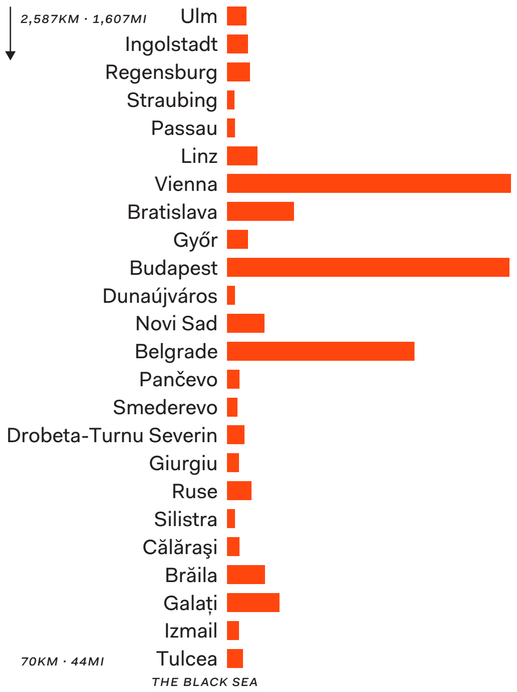

CHRONOLOGY

By city founding year, earliest to latest

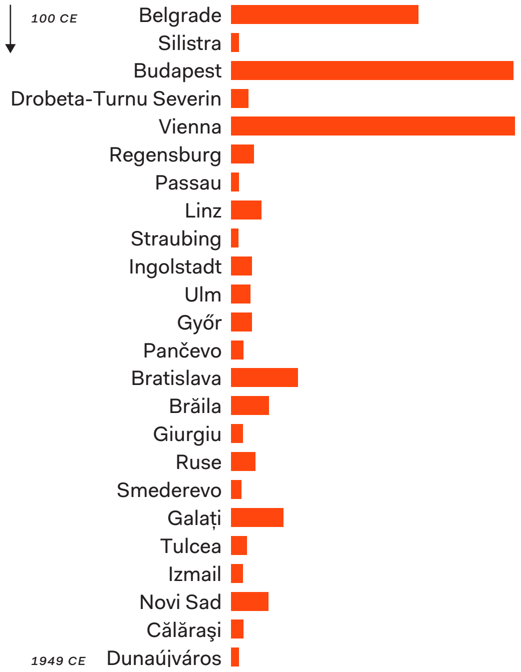

CATEGORIES

Cities in European Union member countries

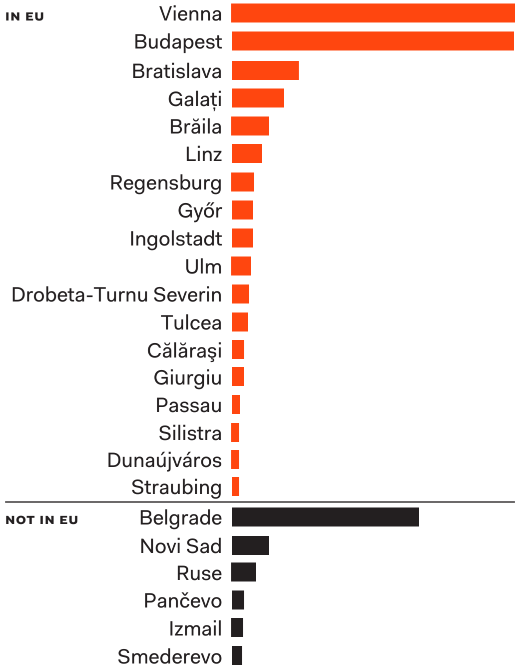

CONNECTION

By railroad connections

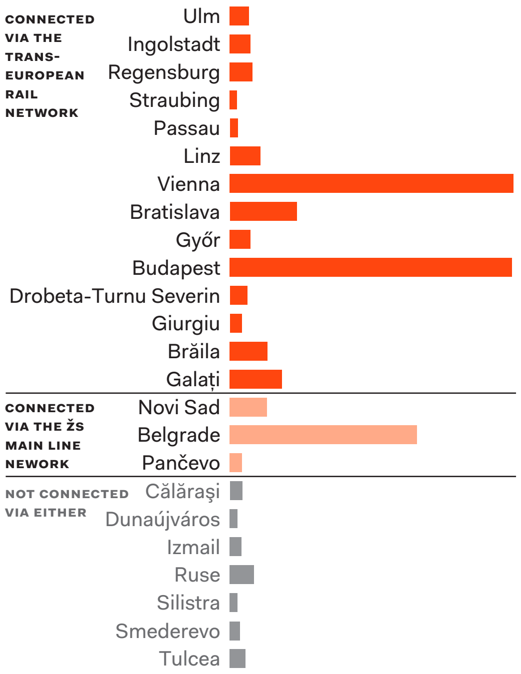

CONVENTION

Alphabetical order by name

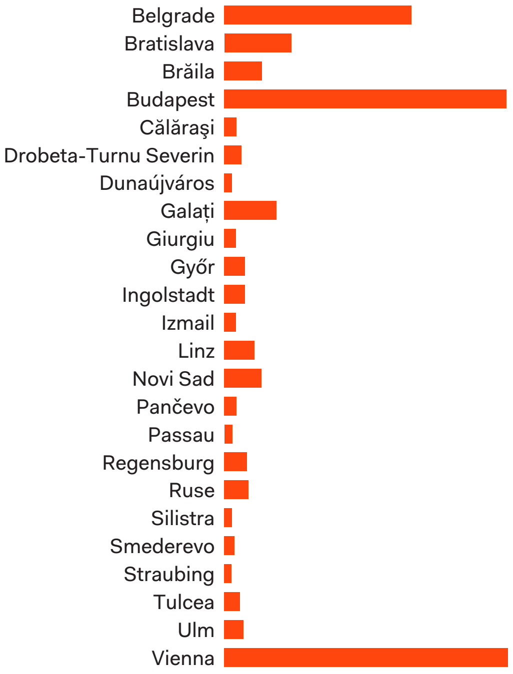

Category as an organizing principle means grouping data points based on some kind of similarity in content. Usually, the grouping is based on values on a *nominal scale* (see p. 94): people can be divided into employed, unemployed, and out of the workforce for example. Quantitative values can, however, also be classified into categories. For example, companies can be classified into small, medium-sized, and large, based on the number of employees or annual turnover.

The visual encoding of categorical data can generally involve the use of visual variables (see pp. 58–62) that are ill-suited for showing other classifications, such as color hue or shapes. For this reason, categorical ordering can often easily be combined with other organizing principles. Even when position is used for encoding categories, the items *within* the categories can be organized by position based on some of the other principles, such as continuum, as in the example on the left.

Connection is an organizing principle mostly used in network diagrams (see pp. 227–229). It is based on the links between the data points, which are often called *nodes* in this context. Often, nodes that are connected are also shown close to each other (see p. 228).

CHRONOLOGY

By city founding year, earliest to latest

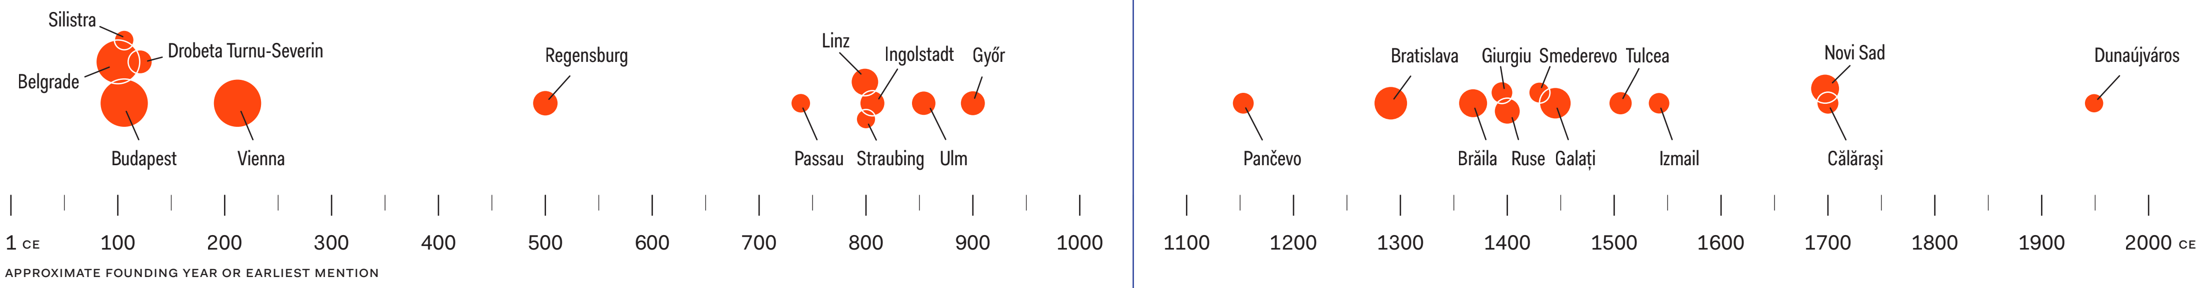

Another example of using chronology as an organizing principle. Here the data points are positioned along an *interval scale*, as opposed to the *ordinal scale* used for positioning the bars in the bar charts on the previous spread. (For further information about scales of measurement, see pp. 94–96.)

Convention refers to organizing principles that are not based on characteristics of the data, like the other six principles, but on an agreed-upon convention, such alphabetical order or, say, numbering of military units (1st Infantry Division, 2nd Infantry Division etc.). Organizing principles based on how the data has been collected, such as following the order in which questions were asked on a questionnaire, also fall in this category.

Convention-based orderings are better than a completely arbitrary order, but in almost all cases, the other six organizing principles are better options. The main advantage of alphabetical order is that the viewer can quickly look up an individual data point. This is useful in a table, but the purpose of a visualization is to reveal patterns and other larger-scale structures (and the exceptions to those patterns) in the data. A convention-based (as opposed to data-based) ordering does not reveal such patterns.

The various organizing principles are not mutually exclusive, and in many cases multiple principles are applied in parallel in the same presentation. Because different organizing principles reveal different patterns in the data, it is often a good idea to let the user select the order of items from a number of options, when technically feasible.

The golden rule of information design

In this book, we present a variety of guidelines and rules for the design of visualizations. Some of these are based on research, while others draw on the observations and experience of the authors of this book and of others working in the field. We believe that by observing them, designers can create clear, illustrative and interesting visualizations. Even as our understanding of human perception grows with new research, an information designer's work is still more art than science. The rules and guidelines we propose do not cover all the potential issues that a designer has to address, and they can, in some cases, conflict with each other. The best result is not achieved by mechanically following the rules. The designer must always use independent judgment when creating visualizations.

There is, however, one rule that an information designer should always follow: **choose the clearest presentation method available.** Disregard any other rules and guidelines we present in this book, when they are in conflict with this golden rule of information design. These rules should not, of course, be ignored without a **well-founded reason**, but such reasons may occasionally come up in a designer's work—as a reader carefully studying the graphics in this book may indeed notice.

A visualization should only make true statements about the real world, in the clearest possible way. The data should be correct and from a trusted source. The presentation method should be selected so as not to mislead the reader and take attention away from the important features of the data. It should show large differences as large, and small differences as small. It should draw the reader's attention first to the most significant features, and leave insignificant features in the background. A visualization should include as much information as possible, but not too much. Its visual style should be appropriate for the context, and it should be carefully executed down to the smallest detail. At its best, a visualization gives an overview of the topic quickly, but also rewards the reader who spends time exploring it in depth.[39]

In order for all these goals to be achievable, the various rules and guidelines set forth in this book should be applied on a case-by-case basis and should sometimes even be broken.

[39] Edward Tufte calls this "micro/macro readings." Tufte 1990, pp. 37–51