

CHAPTER ONE

On Beauty

Noah Iliinsky

THIS CHAPTER IS AN EXAMINATION OF WHAT WE MEAN BY BEAUTY in the context of visualization, why it's a worthy goal to pursue, and how to get there. We'll start with a discussion of the elements of beauty, look at some examples and counterexamples, and then focus on the critical steps to realize a beautiful visualization.*

What Is Beauty?

What do we mean when we say a visual is beautiful? Is it an aesthetic judgment, in the traditional sense of the word? It can be, but when we're discussing visuals in this context, beauty can be considered to have four key elements, of which aesthetic judgment is only one. For a visual to qualify as beautiful, it must be aesthetically pleasing, yes, but it must also be *novel*, *informative*, and *efficient*.

Novel

For a visual to truly be beautiful, it must go beyond merely being a conduit for information and offer some novelty: a fresh look at the data or a format that gives readers a spark of excitement and results in a new level of understanding. Well-understood formats (e.g., scatterplots) may be accessible and effective, but for the most part they no longer have the ability to surprise or delight us. Most often, designs that delight us do

* I use the words *visualization* and *visual* interchangeably in this chapter, to refer to all types of structured representation of information. This encompasses graphs, charts, diagrams, maps, storyboards, and less formally structured illustrations.

so not because they were designed to be novel, but because they were designed to be effective; their novelty is a byproduct of effectively revealing some new insight about the world.

Informative

The key to the success of any visual, beautiful or not, is providing access to information so that the user may gain knowledge. A visual that does not achieve this goal has failed. Because it is the most important factor in determining overall success, the ability to convey information must be the primary driver of the design of a visual.

There are dozens of contextual, perceptive, and cognitive considerations that come into play in making an effective visual. Though many of these are largely outside the scope of this chapter, we can focus on two particulars: the *intended message* and the *context of use*. Keen attention to these two factors, in addition to the data itself, will go far toward making a data visualization effective, successful, and beautiful; we will look at them more closely a little later.

Efficient

A beautiful visualization has a clear goal, a message, or a particular perspective on the information that it is designed to convey. Access to this information should be as straightforward as possible, without sacrificing any necessary, relevant complexity.

A visual must not include too much off-topic content or information. Putting more information on the page may (or may not) result in conveying more information to the reader. However, presenting more information *necessarily* means that it will take the reader longer to find any desired *subset* of that information. Irrelevant data is the same thing as noise. If it's not helping, it's probably getting in the way.

Aesthetic

The graphical construction—consisting of axes and layout, shape, colors, lines, and typography—is a necessary, but not solely sufficient, ingredient in achieving beauty. Appropriate usage of these elements is essential for guiding the reader, communicating meaning, revealing relationships, and highlighting conclusions, as well as for visual appeal.

The graphical aspects of design must primarily serve the goal of presenting information. Any facet of the graphical treatment that does not aid in the presentation of information is a potential obstacle: it may reduce the efficiency and inhibit the success of a visualization. As with the data presented, less is usually more in the graphics department. If it's not helping, it's probably getting in the way.

Often, novel visual treatments are presented as innovative solutions. However, when the goal of a unique design is simply to be different, and the novelty can't be specifically linked to the goal of making the data more accessible, the resulting visual is almost certain to be more difficult to use. In the worst cases, novel design is nothing more than the product of ego and the desire to create something visually impressive, regardless of the intended audience, use, or function. Such designs aren't useful to anyone.

Learning from the Classics

The vast majority of mundane information visualization is done in completely standard formats. Basic presentation styles, such as bar, line, scatter, and pie graphs, organizational and flow charts, and a few other formats are easy to generate with all sorts of software. These formats are ubiquitous and provide convenient and conventional starting points. Their theory and use are reasonably well understood by both visual creators and consumers. For these reasons, they are good, strong solutions to common visualization problems. However, their optimal use is limited to some very specific data types, and their standardization and familiarity means they will rarely achieve novelty.

Beautiful visualizations that go on to fame and fortune are a different breed. They don't necessarily originate with conventions that are known to their creators or consumers (though they may leverage some familiar visual elements or treatments), and they usually deviate from the expected formats. These images are not constrained by the limits of conventional visual protocols: they have the freedom to effectively adapt to unconventional data types, and plenty of room to surprise and delight.

Most importantly, beautiful visualizations reflect the qualities of the data that they represent, explicitly revealing properties and relationships inherent and implicit in the source data. As these properties and relationships become available to the reader, they bring new knowledge, insight, and enjoyment. To illustrate, let's look at two very well-known beautiful visualizations and how they embrace the structure of their source data.

The Periodic Table of the Elements

The first example we'll consider is Mendeleev's periodic table of the elements, a masterful visualization that encodes at least four, and often nine or more, different types of data in a tidy table (see Figure 1-1). The elements have properties that recur periodically, and the elements are organized into rows and columns in the table to reflect the periodicity of these properties. That is the key point, so I'll say it again: the genius of the periodic table is that it is arranged to reveal the related, repeating physical properties of the elements. The structure of the table is directly dictated by the data that it represents. Consequently, the table allows quick access to an understanding of the properties of a given element at a glance. Beyond that, the table also allows very accurate predictions of undiscovered elements, based on the gaps it contains.

		Group																		
		I	II											III	IV	V	VI	VII	VIII	
Period	1	1 H																		2 He
	2	3 Li	4 Be											5 B	6 C	7 N	8 O	9 F	10 Ne	
	3	11 Na	12 Mg											13 Al	14 Si	15 P	16 S	17 Cl	18 Ar	
	4	19 K	20 Ca	21 Sc	22 Ti	23 V	24 Cr	25 Mn	26 Fe	27 Co	28 Ni	29 Cu	30 Zn	31 Ga	32 Ge	33 As	34 Se	35 Br	36 Kr	
	5	37 Rb	38 Sr	39 Y	40 Zr	41 Nb	42 Mo	43 Tc	44 Ru	45 Rh	46 Pd	47 Ag	48 Cd	49 In	50 Sn	51 Sb	52 Te	53 I	54 Xe	
	6	55 Cs	56 Ba	*	72 Hf	73 Ta	74 W	75 Re	76 Os	77 Ir	78 Pt	79 Au	80 Hg	81 Tl	82 Pb	83 Bi	84 Po	85 At	86 Rn	
	7	87 Fr	88 Ra	**	104 Rf	105 Db	106 Sg	107 Bh	108 Hs	109 Mt	110 Ds	111 Rg	112 Uub	113 Uut	114 Uuq	115 Uup	116 Uuh	117 Uus	118 Uuo	
	8	119 Uun																		
		* Lanthanides		57 La	58 Ce	59 Pr	60 Nd	61 Pm	62 Sm	63 Eu	64 Gd	65 Tb	66 Dy	67 Ho	68 Er	69 Tm	70 Yb	71 Lu		
		** Actinides		89 Ac	90 Th	91 Pa	92 U	93 Np	94 Pu	95 Am	96 Cm	97 Bk	98 Cf	99 Es	100 Fm	101 Md	102 No	103 Lr		
				Alkali metals		Alkaline earth metals		Lanthanides		Actinides		Transition metals								
				Poor metals		Metalloids		Nonmetals		Halogens		Noble gases								
State at standard temperature and pressure																				
Atomic number in red: gas																				
Atomic number in blue: liquid																				
Atomic number in black: solid																				
solid border: at least one isotope is older than the Earth (Primordial elements)																				
dashed border: at least one isotope naturally arise from decay of other chemical elements and no isotopes are older than the earth																				
dotted border: only artificially made isotopes (synthetic elements)																				
no border: undiscovered																				

Figure 1-1. A basic example of Mendeleev's periodic table of the elements

The periodic table of the elements is absolutely informative, arguably efficient, and was a completely new approach to a problem that previously hadn't had a successful visual solution. For all of these reasons, it may be considered one of the earlier beautiful visualizations of complex data.

It should be noted that the efficacy and success of the periodic table were achieved with the absolute minimum of graphical treatment; in fact, the earliest versions were text-only and could be generated on a typewriter. Strong graphic design treatment isn't a requirement for beauty.

The London Underground Map

The second classic beautiful visualization we'll consider is Harry Beck's map of the London Underground (aka the Tube map—see Figure 1-2). The Tube map was influenced by conventions and standards for visuals, but not by those of cartography. Beck's background was in drafting electrical circuits: he was used to drawing circuit layout lines at 45° and 90° angles, and he brought those conventions to the Tube map.

That freed the map of any attachment to accurate representation of geography and led to an abstracted visual style that more simply reflected the realities of subway travel: once you're in the system, what matters most is your logical relationship to the rest of the subway system. Other maps that accurately show the geography can help you figure out what to do on the surface, but while you're underground the only surface features that are accessible are the subway stations.



Figure 1-2. The London Underground ("Tube") map; 2007 London Tube Map © TfL from the London Transport Museum collection (used with permission)

The London Underground map highlighted the most relevant information and stripped away much of the irrelevant information, making the pertinent data more easily accessible. It was executed with a distinctive and unique graphical style that has become iconic. It is widely recognized as a masterpiece and is undoubtedly a beautiful visualization.

Other Subway Maps and Periodic Tables Are Weak Imitations

Due to the success of the periodic table and the London Underground map, their formats are often mimicked for representations of other data. There are periodic tables of just about everything you can imagine: foods, drinks, animals, hobbies, and, sadly, visualization methods.* These all miss the point. Similarly, Underground-style

* See http://www.visual-literacy.org/periodic_table/periodic_table.html.

maps have been used to represent movies of different genres,^{*} relationships among technology companies,[†] corporate acquisition timelines,[‡] and the subway systems of cities other than London.

Of these examples, the only reasonable alternate use of the latter format is to represent subways in other cities (many of these—Tokyo, Moscow, etc.—are quite well done). All the other uses of these formats fail to understand what makes them special: their authentic relationships to and representations of the source data. Putting nonperiodic data into a periodic table makes as much sense as sorting your socks by atomic number; there's no rational reason for it because the structure you're referencing doesn't exist. Casting alternate data into these classic formats may be an interesting creative exercise, but doing so misses the point and value of the original formats.

How Do We Achieve Beauty?

Given the abundance of less-than-beautiful visualizations, it's clear that the path to beauty is not obvious. However, I believe there are ways to get to beauty that are dependable, if not entirely deterministic.

Step Outside Default Formats

The first requirement of a beautiful visualization is that it is novel, fresh, or unique. It is difficult (though not impossible) to achieve the necessary novelty using default formats. In most situations, well-defined formats have well-defined, rational conventions of use: line graphs for continuous data, bar graphs for discrete data, pie graphs for when you are more interested in a pretty picture than conveying knowledge.

Standard formats and conventions do have their benefits: they are easy to create, familiar to most readers, and usually don't need to be explained. Most of the time, these conventions should be respected and leveraged. However, the necessary spark of novelty is difficult to achieve when using utilitarian formats in typical ways; defaults are useful, but they are also limiting. Defaults should be set aside for a better, more powerful solution only with informed intent, rather than merely to provide variety for the sake of variety.

Default presentations can also have hidden pitfalls when used in ways that don't suit the situation. One example that I encountered was on a manufacturer's website, where its retailers were listed alphabetically in one column, with their cities and states in a second column. This system surely made perfect sense to whoever designed it, but the design didn't take into account how that list would be used. Had I already known the names of the retailers in my area, an alphabetical list of them would have been useful.

^{*} See <http://blog.vodkaster.com/2009/06/25/the-top-250-best-movies-of-all-time-map/>.

[†] See <http://informationarchitects.jp/wtm4/>.

[‡] See <http://www.meettheboss.com/google-acquisitions-and-investments.html>.

Unfortunately, I knew my location but not the retailer names. In this case, a list sorted by the most easily accessible information (location) would have made more sense than a default alphabetic sort on the retailer name.

Make It Informative

As I mentioned earlier, a visualization must be informative and useful to be successful. There are two main areas to consider to ensure that what is created is useful: the intended message and the context of use. Considering and integrating insight from these areas is usually an iterative process, involving going back and forth between them as the design evolves. Conventions should also be taken into consideration, to support the accessibility of the design (careful use of certain conventions allows users to assume some things about the data—such as the use of the colors red and blue in visuals about American politics).

Intended message

The first area to consider is what knowledge you're trying to convey, what question you're trying to answer, or what story you're trying to tell. This phase is all about planning the function of the visual in the abstract; it's too early to begin thinking about specific formats or implementation details. This is a critical step, and it is worth a significant time investment.

Once the message or goal has been determined, the next consideration is how the visualization is going to be used. The readers and their needs, jargon, and biases must all be considered. It's enormously helpful in this phase to be specific about the tasks the users need to achieve or the knowledge they need to take away from the visualization. The readers' specific knowledge needs may not be well understood initially, but this is still a critical factor to bear in mind during the design process.

If you cannot, eventually, express your goal concisely in terms of your readers and their needs, you don't have a target to aim for and have no way to gauge your success. Examples of goal statements might be "Our goal is to provide a view of the London subway system that allows riders to easily determine routes between stations," or "My goal is to display the elements in such a way that their physical properties are apparent and predictions about their behaviors can be made."

Once you have a clear understanding of your message and the needs and goals of your audience, you can begin to consider your data. Understanding the goals of the visualization will allow you to effectively select which facets of the data to include and which are not useful or, worse, are distracting.

Context of use. It's also important to recognize the distinction between visuals designed to reveal what the designer already knows, and visuals intended to aid research into the previously unknown (though the designer may suspect the outcome in advance). The former are tools for presentation; the latter are tools for examination. Both may take standard or unconventional formats, and both benefit from the same process and treatments. However, it is important to be clear about which type of visual is being designed, as that distinction affects all subsequent design choices.

Visualizations designed to reveal what is already known are ubiquitous, appearing wherever one party has information to convey to another using more than just text. Most graphs and charts that we encounter are meant to communicate a particular insight, message, or knowledge that is evident in the underlying data: how a team is performing, how a budget is divided, how a company is organized, how a given input affects a result, how different products compare to each other, and so on. The data might reveal other knowledge or insights as well, but if they aren't important for the purpose at hand, the design need not focus on revealing these other messages or trends. The process of designing these visualizations is therefore aided by having a well-defined goal.

Visualizations designed to facilitate discovery are commonly found in more specific, research-oriented contexts in science, business, and other areas. In these cases, the goal is typically to validate a hypothesis, answer a specific question, or identify any trends, behaviors, or relationships of note. Designing these visualizations can be more challenging if it's unclear what insights the data may reveal. In contexts where the shape of the answer is unknown, designing several different visualizations may be useful.

The periodic table is an interesting hybrid of these purposes, in that it was used to visualize both known and unknown information. The structure of the table was defined by the properties of the elements known at the time, so in that way it was a reference to existing knowledge, as it is used today. However, this structure resulted in gaps in the table, which were then used to predict the existence and behavior of undiscovered elements. In this latter mode, the table was a tool of research and discovery.

Make It Efficient

After ensuring that a visualization will be informative, the next step is to ensure that it will be efficient. The most important consideration when designing for efficiency is that every bit of visual content will make it take longer to find any particular element in the visualization. The less data and visual noise there is on the page, the easier it will be for readers to find what they're looking for. If your clearly stated goal can't justify the existence of some of your content, try to live without it.

Visually emphasize what matters

When you've identified the critically necessary content, consider whether some portion of it—a particular relationship or data point—is especially relevant or useful. Such content can be visually emphasized in a number of ways. It can be made bigger, bolder, brighter, or more detailed, or called out with circles, arrows, or labels. Alternately, the less-relevant content can be de-emphasized with less intense colors, lighter line weight, or lack of detail. The zones in the Tube map, for example, are visually deemphasized: they exist, but clearly aren't as relevant as the Tube lines and stations.

Note that this strategy of emphasizing relevance typically applies to presentation data, not research data: by changing the emphasis, the designer is intentionally changing the message. However, highlighting different facets or subsets of unknown data is a valid way to discover relationships that might otherwise be lost in the overall noise.

Use axes to convey meaning and give free information

One excellent method for reducing visual noise and the quantity of text while retaining sufficient information is to define axes, and then use them to guide the placement of the other components of the visualization. The beauty of defining an axis is that every node in a visualization can then assume the value implied by the axis, with no extra labeling required. For example, the periodic table is made up of clearly defined rows (periods) and columns (groups). A lot of information can be learned about an element by looking at what period and group it occupies. As a result, that information doesn't have to be explicitly presented in the element's table cell. Axes can also be used to locate a portion or member of the dataset, such as looking for an element in a particular period, southern states, or a Tube station that is known to be in the north-west part of London.

Well-defined axes can be effective for qualitative as well as quantitative data. In qualitative contexts, axes can define (unranked or unordered) areas or groupings. As with quantitative axes, they can provide information and support the search for relevant values.

Slice along relevant divisions

One last way to reduce visual clutter and make information more accessible is to divide larger datasets into multiple similar or related visualizations. This works well if the information available can be used independently and gains little (or infrequent) value from being shown in conjunction with the other data in the set. The risk here is that there may be relevant, unsuspected correlations among seemingly unrelated datasets that will only become evident when all the data is displayed together.

Use conventions thoughtfully

After the influences of the intended message, context of use, and data have been taken into consideration for your unique situation, it's worth looking into applying standard representations and conventions. Intentional and appropriate use of conventions will speed learning and facilitate retention on the part of your readers. In situations where a convention does exist, and doesn't conflict with one of the aforementioned considerations, applying it can be extremely powerful and useful. The examples we've examined have used default, conventional representations for element symbols, subway line colors, and compass directions. Most of these seem too obvious to mention or notice, and that's the point. They are easily understood and convey accurate information that is integrated extremely rapidly, while requiring almost no cognitive effort from the user and almost no creative effort from the designer. Ideally, this is how defaults and conventions should work.

Leverage the Aesthetics

Once the requirements for being informative and efficient have been met, the aesthetic aspects of the visual design can finally be considered. Aesthetic elements can be purely decorative, or they can be another opportunity to increase the utility of the visualization. In some cases visual treatments can redundantly encode information, so a given value or classification may be represented by both placement and color, by both label and size, or by other such attribute pairings. Redundant encodings help the reader differentiate, perceive, and learn more quickly and easily than single encodings.

There are other ways in which aesthetic choices can aid understanding: familiar color palettes, icons, layouts, and overall styles can reference related documents or the intended context of use. A familiar look and feel can make it easier or more comfortable for readers to accept the information being presented. (Care should be taken to avoid using familiar formats for their own sake, though, and falling into the same traps as the designers of the unfortunate periodic tables and Tube-style maps.)

At times, designers may want to make choices that could interfere with the usability of some or all of the visualization. This might be to emphasize one particular message at the cost of others, to make an artistic statement, to make the visualization fit into a limited space, or simply to make the visualization more pleasing or interesting to look at. These are all legitimate choices, as long as they are done with intention and understanding of their impact on the overall utility.

Putting It Into Practice

Let's look at one more example of a successful, data-driven visualization that puts these principles to work: a map of the 2008 presidential election results from the *New York Times*.^{*} Figure 1-3 is a standard map of the United States, with each state color-coded to represent which candidate won that state (red states were won by the Republican candidate, blue states by the Democratic candidate). This seems like a perfectly reasonable visualization making use of a default framework: a geographic map of the country. However, this is actually a situation in which an accurate depiction of the geography is irrelevant at best and terribly misleading at worst.

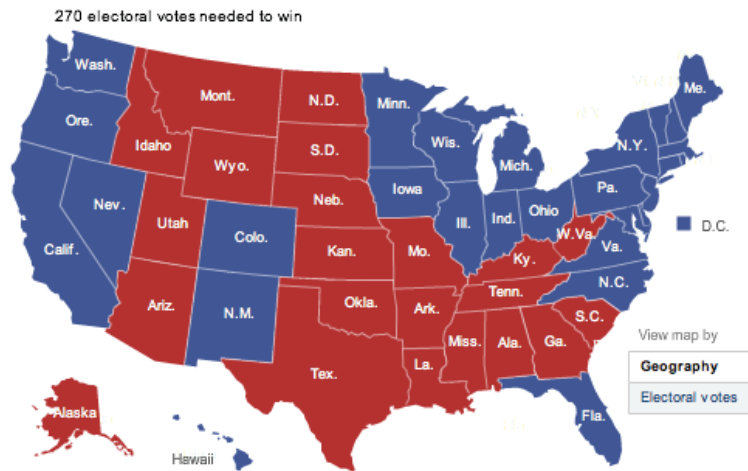


Figure 1-3. A geographically accurate electoral vote results map of the United States

New Jersey (that peanut-shaped state east of Pennsylvania and south of New York that's too small for a label) has an area of a little more than 8,700 square miles. The total combined areas of the states of Idaho, Montana, Wyoming, North Dakota, and South Dakota is a bit more than 476,000 square miles, about 55 times the area of New Jersey, as shown in Figure 1-4. If we were interested in accurate geography and the shape, size, and position of the states, this would be a fine map indeed. However, in the context of a presidential election, what we care about is relative influence based on the electoral vote counts of each state. In fact, the combined total of those five states is just 16 electoral votes, only one more than New Jersey's 15 votes. The geographically accurate map is actually a very inaccurate map of electoral influence.

^{*} Source: <http://elections.nytimes.com/2008/president/whos-ahead/key-states/map.html>.




Figure 1-4. *Relative size of five states versus New Jersey*

The surface area of a state has nothing to do with its electoral influence; in this context, an entirely different sort of visualization is needed to accurately represent the relevant data and meet the goal of the visualization. To this end, the *Times* also created an alternate view of the map (Figure 1-5), in which each state is made up of a number of squares equivalent to its electoral vote value. This electorally proportionate view has lost all geographic accuracy regarding state size, and almost all of it regarding shape. The relative positions of the states are largely retained, though, allowing readers to find particular states in which they may have interest and to examine regional trends. The benefit of sacrificing geography here is that this visualization is perfectly accurate when it comes to showing the electoral votes won by each party and each state's relative influence. For example, when we look at this new map, a comparison of the size of the five states previously mentioned versus New Jersey now accurately depicts their 16 to 15 electoral vote tallies, as shown in Figure 1-6.

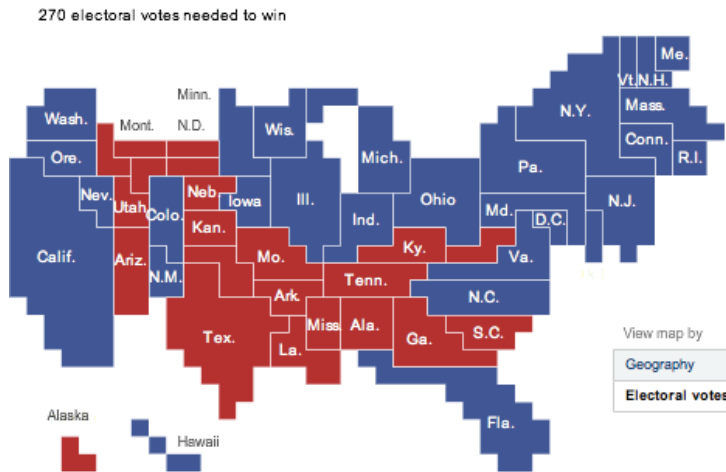


Figure 1-5. *A proportionally weighted electoral vote results map of the United States*




Figure 1-6. *Relative electoral vote influence of five states versus New Jersey*

You may have noticed that another trade-off was made here: because readers can't see the outlines of each individual square, they can't easily count the 15 or 16 squares in each of the areas we're comparing. Also, because a decision was made to retain the shape of each state to the extent possible, the aggregated red and blue blocks in Figure 1-6 are shaped very differently from each other, making it difficult to compare their relative areas at a glance. So, this is a great example of the necessary balancing act between making use of conventions (in this case, the shape of the states) and presenting data efficiently and without decoration.

The success of this visualization is due to the fact that the designers were willing to move away from a standard, default map and instead create a visual representation based primarily on the relevant source data. The result is a highly specialized image that is much more accurate and useful for its intended purpose, even if it's not very well suited for typical map tasks such as navigation. (In that way, it is similar to the Tube map, which is optimized for a very particular style of information finding, at the expense of general-purpose geographical accuracy.)

Conclusion

While this has been a brief treatment of some of the strategies and considerations that go into designing a successful visualization, it is a solid foundation. The keys to achieving beauty are focusing on keeping the visualization useful, relevant, and efficient, and using defaults and aesthetic treatments with intention. Following these suggestions will help ensure that your final product is novel, informative, and beautiful.

