



Visual Analysis Best Practices

Simple Techniques for Making Every
Data Visualization Useful and Beautiful

Table of Contents

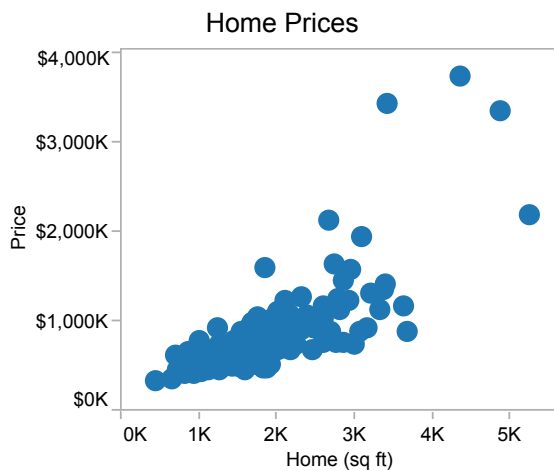
Start with Questions.....	5
Choose the Right Chart Type.....	7
Trends over Time.....	7
Comparison and Ranking.....	9
Correlation.....	10
Distribution.....	11
Part to Whole.....	13
Geographical Data.....	14
Create Effective Views.....	15
Emphasize the most important data.....	16
Orient your views for legibility.....	17
Organize your views.....	18
Avoid overloading your views.....	20
Limit the number of colors and shapes in a single view.....	21
Design Holistic Dashboards.....	22
General guidelines.....	23
Give the wheel to your users with interactivity.....	24
Highlighting.....	25
Filters.....	26
Hyperlinking and using the power of the web.....	30
Sizing: Making sure your visualization is visible.....	31
Scrollbars.....	31
Scrunched views.....	32
Fitting.....	32
Perfecting Your Work.....	33
Color TV looks better than black and white: Use color visualizations.....	34
Fonts: Can you read this? (Can you read this?).....	35
Tooltips tell the story.....	36
Make intelligent axes or you'll just have pretty rectangles.....	37
Labeling your visualization in all the right places.....	38
Evaluating Your Masterpiece.....	39

About this paper


You made a visualization! Congratulations: you are part of a small but growing group that's taking advantage of the power of visualization. However, bringing your visualizations from "good" to "great" takes time, patience and attention to detail. Luckily, we have compiled a short but important list of techniques to get you started. Happy visualizing!

Note: This document does not provide basic instructions for building visualizations; we assume you know how to do that. The purpose of this document is to provide tips for making your visualizations more effective. Below are examples of two visualizations: one that's good and another that's great. Why is one good and one great? Simply because the latter visualization helps you understand more about the data—and more quickly. There are many ways to do this. We'll touch on the different techniques throughout this paper.

Good visualization




Great visualization



Start with Questions


What are you trying to say? The single most important step you can take to make a great visualization is to know what you're trying to say. With the amount of data available today, anyone can get lost in a world of scatter plots and geocoding. It is vital that your visualization has a purpose and you are selective about what you include in your visualization to fulfill that purpose.

How do you know if your visualization has a purpose? Well, ask some questions to find out. Who is your audience? What questions do they have? What answers do you find for them? What other questions does it inspire? What conversations will result? The point is that your viewers should take something away from the time they spend with your visualization.

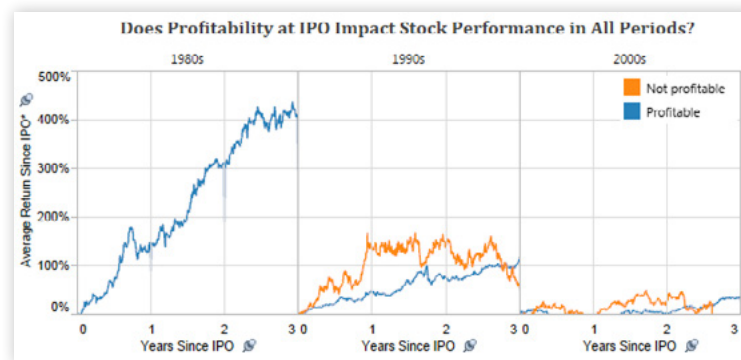


It is vital that your visualization has a purpose and you are selective about what you include in your visualization to fulfill that purpose.

Let's say you work for a stock broker who focuses on IPO investments, and you want to make a visualization to help him decide where to invest. You might ask a question like "Does profitability at IPO affect stock performance?" That might lead you to produce this view:




From this view, you would discover that profitability at IPO has a huge effect on its later performance. However, this dataset contains information on all software company IPOs over the past three decades. You might wonder if the trend you've discovered holds true throughout all periods. A second view could help you answer that question:



You can see from this view that the trend only applies to the 1990s. Furthermore, you can now make two more discoveries: 1) All companies were profitable at IPO in the 1980s, and 2) Profitability at IPO didn't have a huge impact on stock performance in the 2000s. Does this mean that modern investors are more risk-prone than their predecessors? Or is it that companies that were not profitable at IPO have equal likelihood of future success as those that are profitable? You can explore further to find out.

Choose the Right Chart Type

Once you have a purpose for your visualization in mind, it is time to think about what types of analysis will help you achieve that purpose. In this section, we will discuss some of the most commonly used analysis types and which chart types are most suitable for them.

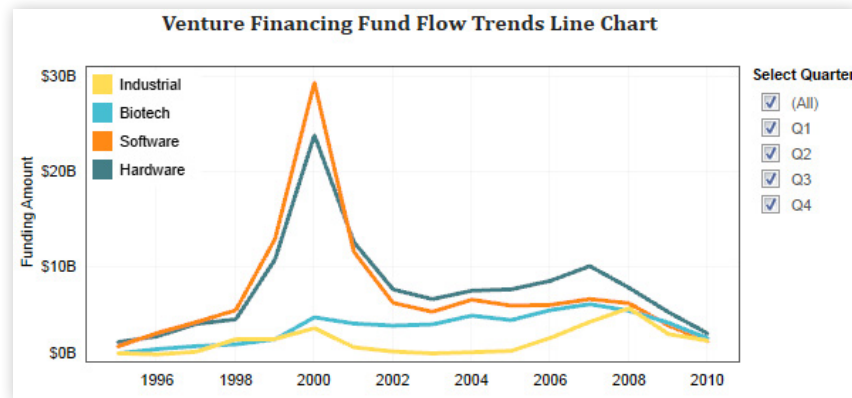


Some of the best visualizations for showing trends over time are line charts, area charts, and bar charts. In addition, you should try to put time on the X-axis and the measure on the Y-axis.



Trends over Time

One of the most frequently used methods for analyzing data is to track a trend over time. In the example below, we want to see trends over time, by sector, in the flow of venture-financing funds. In our experience, some of the best visualizations for showing trends over time are line charts, area charts, and bar charts. In addition, you should try to put time on the X-axis and the measure on the Y-axis; that will help your view cater to our cultural conventions on trending.

First, let's look at the line chart below. We put year on the X-axis, funding amount on the Y-axis, and encoded sector type with color. What we can see from this view is that all sectors follow the same trend in funding over time. We can also see the trends of each individual sector and the differences between them. But what about the overall funding trend? Can you see exactly how much funding there is for all sectors in 2000 or any other given point in time?




The answer is no. Line charts do not have that capability. However, if answering these questions is important to you, you can explore area charts or bar charts (see below). Both of these chart types are great at amplifying total funding trends over time and how each individual sector contributes to the total over time. However, they do have a distinct difference: The area chart treats each sector as a single pattern while the bar chart focuses on each year as a single pattern.



Comparison and Ranking

Another method for analyzing data is comparison and ranking. We compare and rank countries, regions, business segments, salesmen and sports players based on one or a set of criteria. In many cases, this shows us where we are and how we are doing. A bar chart is great for comparison and ranking because it encodes quantitative values as length on the same baseline, making it extremely easy to compare values.

A bar chart is great for comparison and ranking because it encodes quantitative values as length on the same baseline, making it extremely easy to compare values.




Correlation

People are great at discovering relationships between measures—it's something we do all the time in data analysis. Does smoking cause cancer? Does the price of a product impact the amount that gets sold? Does outdoor temperature correlate with increased boating activity? Does government stimulus create more jobs?


Running a simple correlation analysis is a great place to start in identifying relationships between measures. Be mindful that correlation doesn't guarantee a relationship. Instead it only suggests a potential relationship. To confirm the relationship truly exists, a more sophisticated methodology is often required.

Here is an example of how to build a simple scatter plot to detect correlations between two factors. The data is from a Deli-food wholesale company. We put sales price on the Y-axis, sales quantity on the X-axis, and include monthly sales numbers on details. As we can see from the chart (specifically when we add a trend line), there is a clear negative correlation between sales price and quantity. When price is high, quantity is low—and vice versa. Does this mean that the company should lower prices to boost sales? Not necessarily. This is why we overlay the net profit onto the size of the circles; it looks like the company makes the greatest profit on both ends.

Running a simple correlation analysis is a great place to start in identifying relationships between measures. Be mindful that correlation doesn't guarantee a relationship.




But what about other chart types? Can we derive similar results from them? In this example, we combine two line charts with a bar chart. By putting trend lines for sales price and quantity side-by-side at the top, we guide the viewer's focus toward a comparison of these two trends. There is a clear negative correlation, right? Yes. Again, the net-profit bar chart at the bottom gives us additional information for decision-making—without interrupting the correlation analysis.



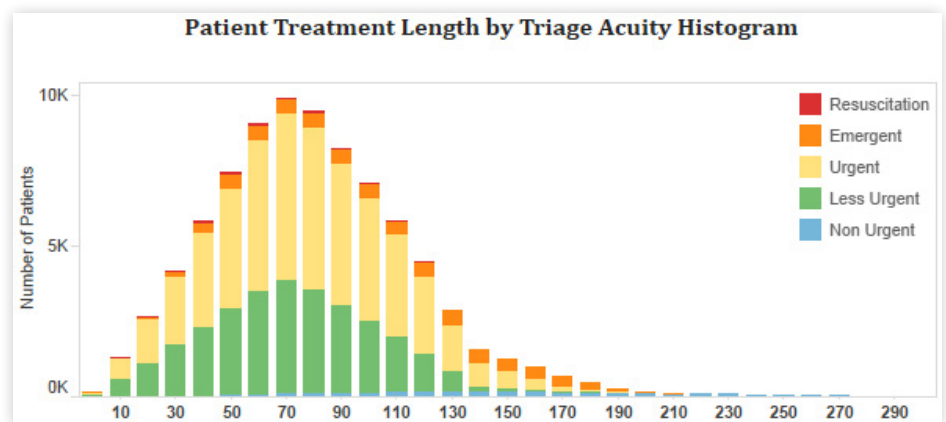
Distribution

Distribution analysis is extremely useful in data analysis because it shows how your quantitative values are distributed across their full quantitative range. For example, a hospital might want to look at their distribution of patient treatment duration. But what graph should they use? Well, there are a few options. In this example, we will discuss two of the most commonly used chart types for this purpose. One is a box plot, and the other is a histogram.

Box plots are excellent for displaying multiple distributions. They pack all your data points—in this case, minutes per patients—into a box and whisker display (see below). Now you can easily identify the low values, 25th-percentile values, the medians, the 75th-percentiles, and the maximum values across all categories—all at the same time. What really stands out from this box plot is that treatment length varies largely between patients in the Emergent and Non-Urgent categories because their boxes are much bigger. Why? Do some patients come in as emergent but then turn out not to be emergent? Or do they have significantly different problems and therefore receive significantly different treatments? Or maybe the doctors in the emergent rooms vary significantly in the way they treat the patients.




An alternative method for displaying distributions is to use a histogram. Instead of breaking up the data by Triage Acuity category and plotting the time that each patient spends in each category, in a histogram we break the data out by time segments (or bins) and count the number of patients in each segment. This graph also shows us that the peak (or most common) treatment length is 70 minutes. We can also color the bars to show the patient count varies by Triage Acuity category. Doing this allows us to see that there are patients in multiple categories in most of the time segments, and that “Urgent” and “Less Urgent” are the most common Triage Acuity categories.



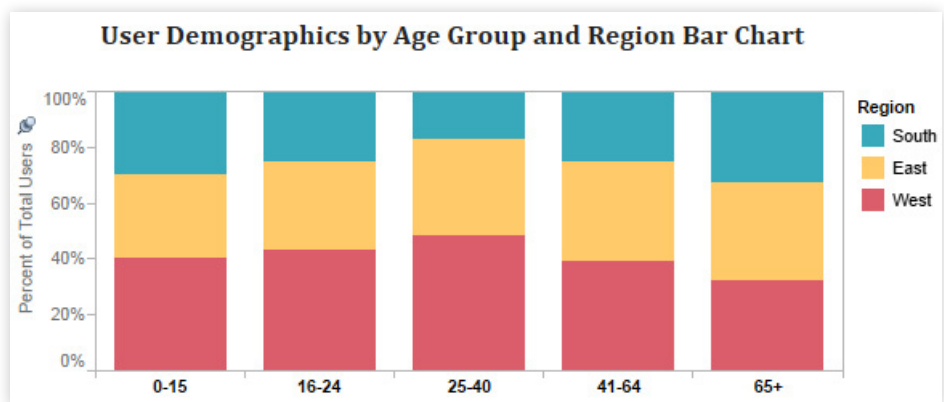
Part to Whole

There are occasions when you want to do a part-to-whole analysis. Although pie charts are commonly used in this type of situation, we suggest avoiding them for two reasons: 1) The human visual system is not very good at estimating area, and 2) You can only compare slices that are right next to each other. For example, in the chart below, can you tell which slice is the largest or how the Western region differs across age groups?

Although pie charts are commonly used in part-to-whole analysis, we suggest avoiding them.



It can be difficult to make these comparisons with pie charts. But what about bar charts? Below you'll see the same data, but plotted on a percent-total bar chart. Can we answer our earlier questions with this chart? Sure! We now see that the 25-40 age group in the Western region is the largest slice. In addition, we can now see the regional differences across all age groups much more easily.

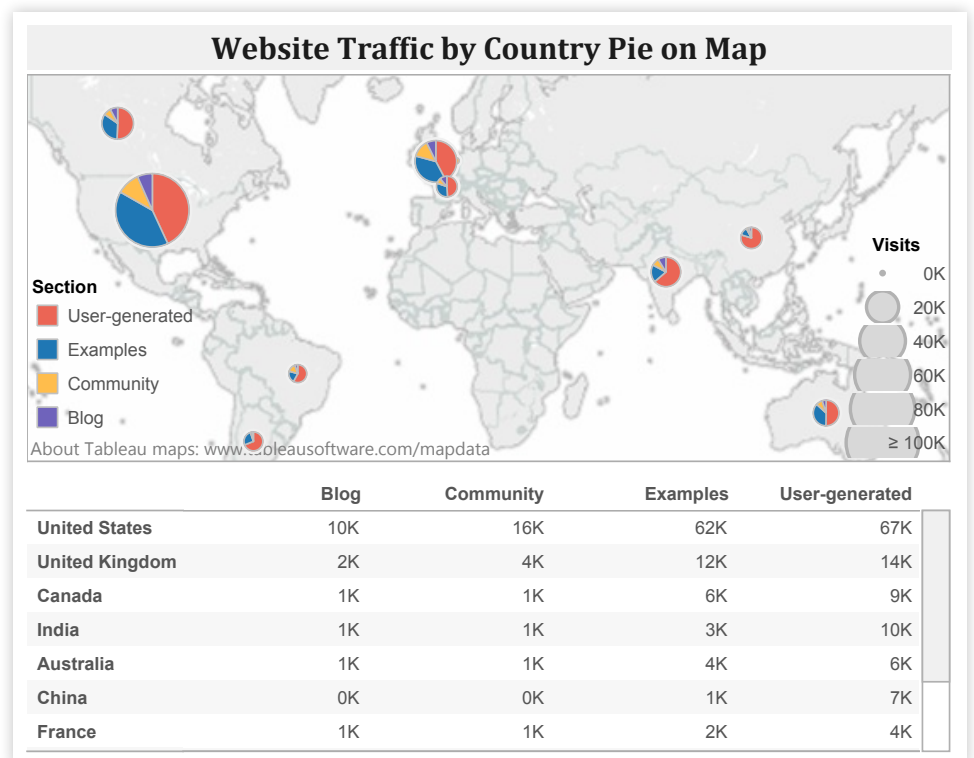


Geographical Data

When you want to show a location, use a map! Just remember that maps are often best when paired with another chart that details what the map displays—such as a bar chart sorted from greatest to least, a line charts showing the trends, or even just a cross-tab to show the actual data.


Although we do not recommend pie charts for part-to-whole relationship, we encourage the usage of pie charts on maps, such as in the website-traffic map below. By using pies on the map, we get a rough breakdown view of each country, which can be very useful when complemented by other chart types like the ones mentioned previously.

Maps are often best when paired with another chart that details what the map displays.



Create Effective Views


Even if you have chosen the best chart type(s) for your analysis, effective views do not always come naturally. Creating effective views requires effort, intuition, attention to detail—and trial and error. Below are some examples that demonstrate the effect of going the extra mile to make your views more effective and easy to understand.



Emphasize the most important data


Many chart types let you put multiple measures and dimensions in one view. In scatter plots, for example, you can put measures on the X- or Y-axis, as well as on the marks for color, size, or shape. Choosing where to put each measure depends on what kind of analysis you are doing and what you are trying to emphasize. However, a rule of thumb is to put the most important data on the X- or Y-axis and less important data on color, size, or shape.

Below, you will see a view we produced for home buyers. Its purpose is to help them understand the relationship between home price, home size, lot size, and the type of home they are interested in. What is the first relationship you see in this view?




Yes, the relationship between price and lot size is pretty clear. But is this the most important information for home buyers? Probably not; the relationship between price and home size probably takes precedence. For most homebuyers, finding a home with the right livable space takes priority over the lot size. This is why the chart below is more effective.

A rule of thumb is to put the most important data on the X- or Y-axis and less important data on color, size, or shape.



Orient your views for legibility


Sometimes, simple changes can go a very long way toward making your visualizations easy to interact with. For example, take a look at the view below:



Did you find it difficult to read? If so, that's probably because all of the labels are vertically oriented. This makes them difficult to read. If you find yourself with a view that has long labels that only fit vertically, try rotating the view. You can quickly swap the fields on the Rows and Columns shelves using the Swap toolbar button.


The same view is shown below—only this time with a horizontal orientation. This simple change makes the chart easier to read and make comparisons with.

If you find yourself with a view that has long labels that only fit vertically, try rotating the view.

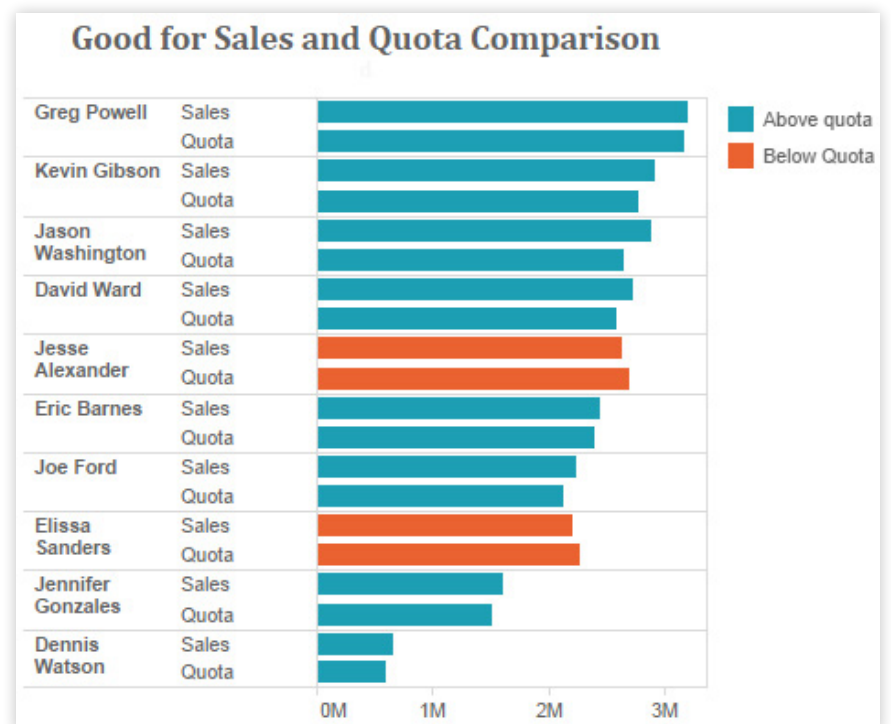


Organize your views

Let's say we want to evaluate a sales team by comparing their sales with their quotas. Our intuition tells us that we should put the two measures close together or side-by-side. This might lead to the chart shown below. But in this view, is it easy to see how well Greg Powell has performed? We know he is above quota... but by how much, exactly? With two side-by-side horizontal bars, it can be very difficult to make comparisons like this.

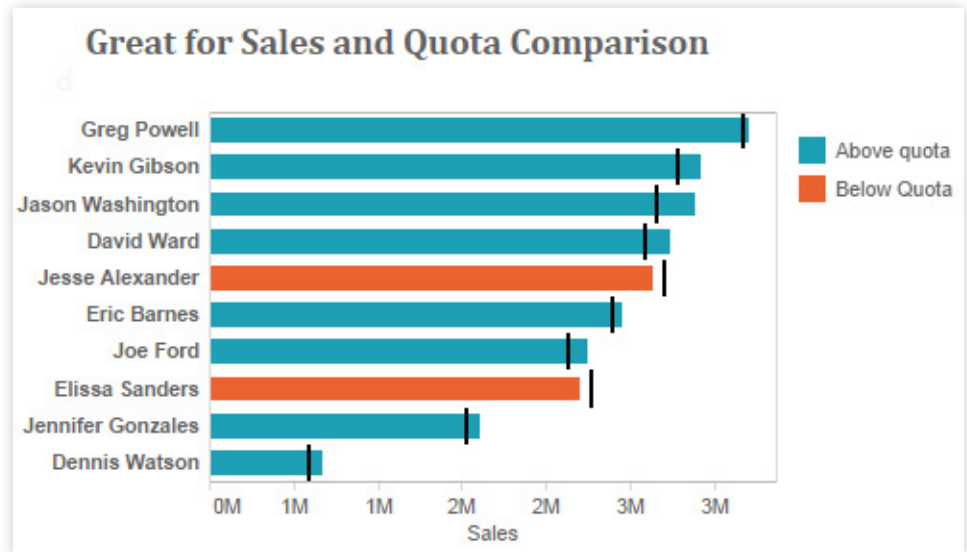


But what if we look at it like this? Instead of putting sales and quota data into columns, we put them into rows. In doing so, we create a shared baseline for the sales bar and the quota bar—which makes comparison far easier. Now we can see that Greg Powell is above quota, but only by a small amount.



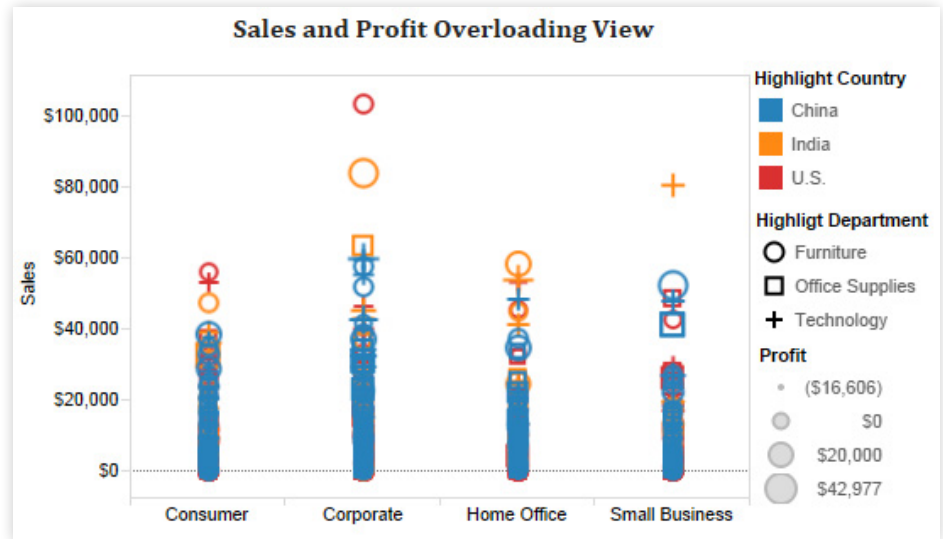
Bullet chart combines a bar chart with reference lines to create a great visual comparison between actual and target numbers.

However, there is even better visualization type for displaying this data: the bullet chart. This chart type combines a bar chart with reference lines to create a great visual comparison between actual and target numbers. In this instance, “actual” is sales (bars) and “target” is quota (vertical reference lines). Not only can we easily see how well each sales person is performing to his/her quota, but we can also cut down 50% of the bars by displaying data in reference lines.




Avoid overloading your views

Overloading a view is one of the most common mistakes people make in data visualization. Look at the view below. Can you tell how well India is doing in terms of sales and profits per customer and department? Well, probably not because there are simply too many measures and dimensions shown in this view.




Instead of stacking countries, departments and profit into one condensed view, break them down to small multiples. Now we can see and understand all of the relevant information in seconds. This is a great example of where smart usage of visualization in combination with traditional cross-tabs can have clear advantages.

Instead of stacking many measures and dimensions into one condensed view, break them down to small multiples.




Limit the number of colors and shapes in a single view

Effective use of color and shape can help us see and understand patterns more easily. However, too many colors and shapes in one view usually defeats that purpose. There are 24 colors visible in the view below. And with all the colors and lines clustered together, it is almost impossible to tell which country is on which line—let alone what pattern that country shows in terms of order quantity. In addition, some countries have the same or similar colors because there are simply not enough distinguishable colors available




Instead of putting Countries on color, let's put Departments on color and see what a difference it makes. Can you see the trends for each department more clearly? Absolutely! Limit the number of colors and shapes in one view to 7-10 so that you can distinguish them and see important patterns.

Limit the number of colors and shapes in one view to 7-10 so that you can distinguish them and see important patterns.



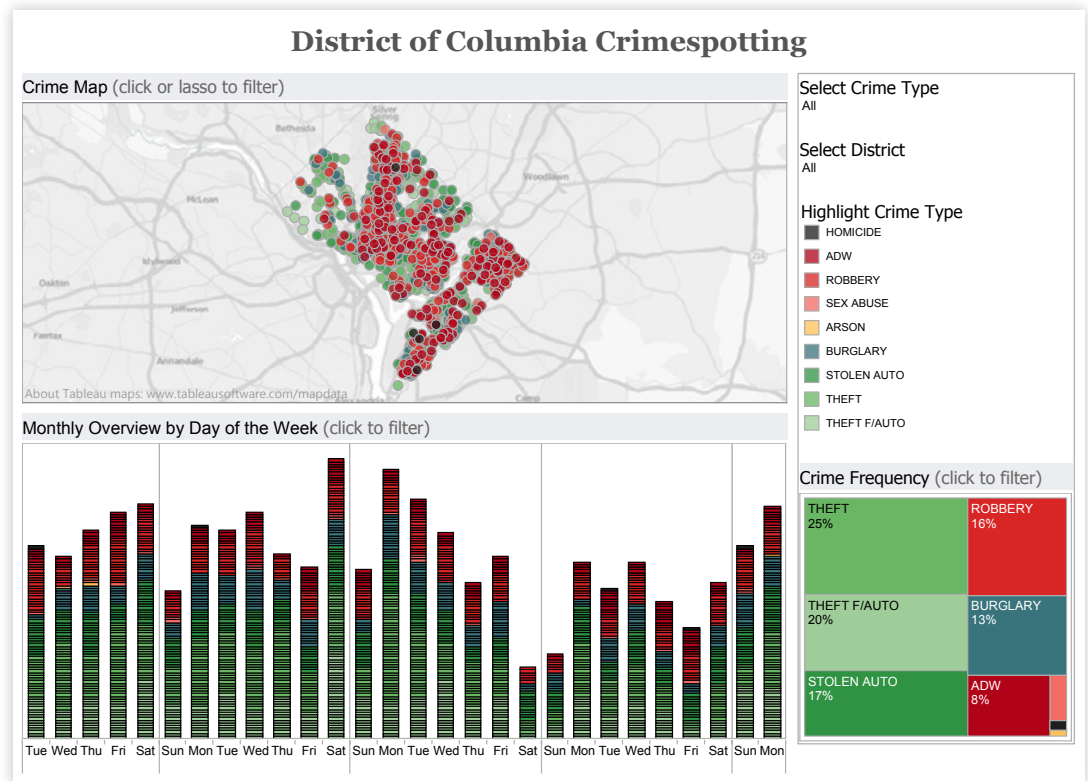
Design Holistic Dashboards

A Tableau dashboard is a collection of several related visualizations shown on a single page—and usually tied together through interactivity. Dashboards increase the analytical power of your visualization by showing multiple perspectives in the same location. They can also be used to combine multiple types of data in a single location.



General guidelines

When designing a dashboard it's important to structure it in a way that is accessible to your audience. For example, the dashboard below is a great example of interactive views that tell a single story together. This dashboard is accessible because it guides you sequentially through each important piece of the story: the crime locations, the day of the week, and the crime frequency. In addition, it's hard to miss the interactive panel on the top right and the interactive instructions that are subtly embedded in the titles.



Use interactive views only when it is necessary: you need to guide a story, encourage user exploration, or there is too much detail to show all at once.

The following guidelines will help you design great dashboards:

- Place the most important view at the top of your dashboard, or in the upper left corner. When looking at a dashboard, your eye is usually drawn to that corner first.
- If your visualization has chained interactivity (the first view filters the next view which filters the last view, etc.), structure them from top to bottom and left to right. That way, the final view to be filtered will be on the bottom, or bottom right.
- Unless there is an absolute need to add more, limit the number of views in your dashboard to three or four. If you add too many views, the big picture can get lost in the details. Remember, you can always use multiple dashboards to tell one story!
- Avoid using multiple color schemes in a dashboard—unless there are natural and independent color schemes in your data.
- If you have multiple filters, try to group them together with a layout container. A light border around them gives a subtle visual cue that they have shared features. The right, top, and left sides of the dashboard are all great places to put your filters.
- If a legend applies to all of your views, place them together with all of your filters. If a legend applies to one or a few more views, place it as close to those views as possible.

Give the wheel to your users with interactivity

Interactivity can make the difference between a horribly confusing visualization and an all-star analysis. However, unnecessary interactivity can turn an all-star analysis into a useless one. This section will teach you how to use interactivity well, but remember that interactive views should only be used when it is necessary: you need to guide a story, encourage user exploration, or there is too much detail to show all at once.

Whatever interactivity you build into your visualization, make sure your viewers know that they can interact with it—and understand where to look for the changes that their interaction will bring about. You can also include subtle instructions for your viewers. In the views below, sub-headings suggest interactivity by using verbs such as “Select,” “Highlight” and “Click.”

Highlights let you quickly show relationships between values in a specific area or category—even across multiple views.

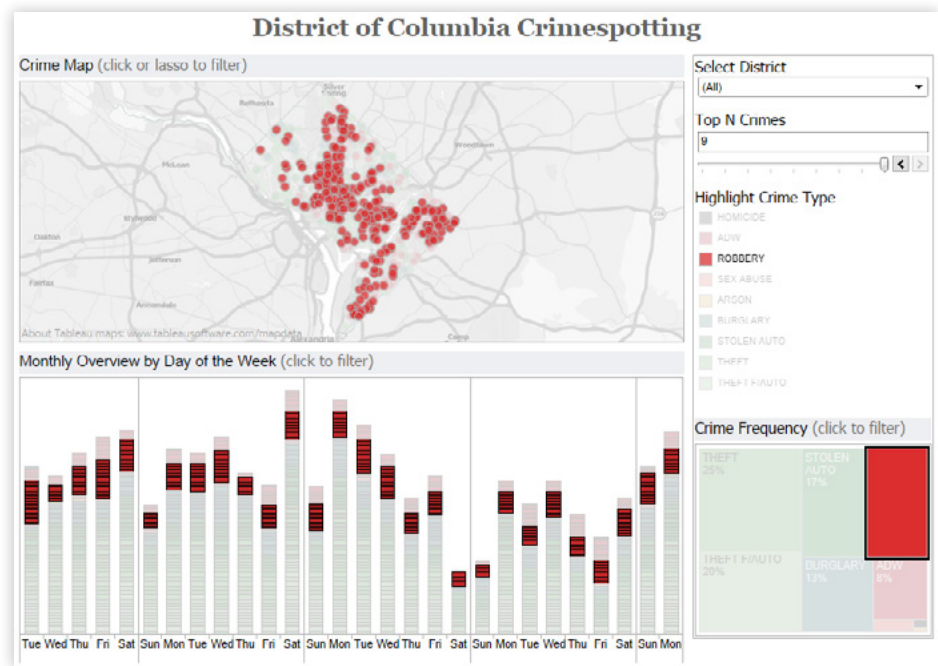
Highlighting

Highlights let you quickly show relationships between values in a specific area or category—even across multiple views. One of the best things about highlighting is that it preserves the context of the rest of the points (unlike filtering which we will discuss later).

To add highlighting functionality to your visualization, either click the highlight icon in the legend or add a highlight action through the Dashboard Actions menu. The difference between these two methods is that the latter gives you more control by allowing you to select specific source sheets, target sheets, and data fields.

While designing highlights on a dashboard, ask yourself these questions: What are your users interested in? Will a highlight function help them see patterns in your data more easily? Are there relationships in your data that you'd like to emphasize? Which views or data fields do you want to build your highlights upon? Would it help you make a particular point by having something highlighted initially when you publish?

In this example, we make the highlight function available across all views by crime type, meaning that when you click on a particular crime type in any view, the data related to that crime type gets highlighted in all the other views. Now we can more easily see where robbery occurs in the city—as well as the number of occurrences each day and its percentage of total crime—all at the same time.




Filters are great ways to enable multi-level data exploration and user-driven data analysis.

Filters

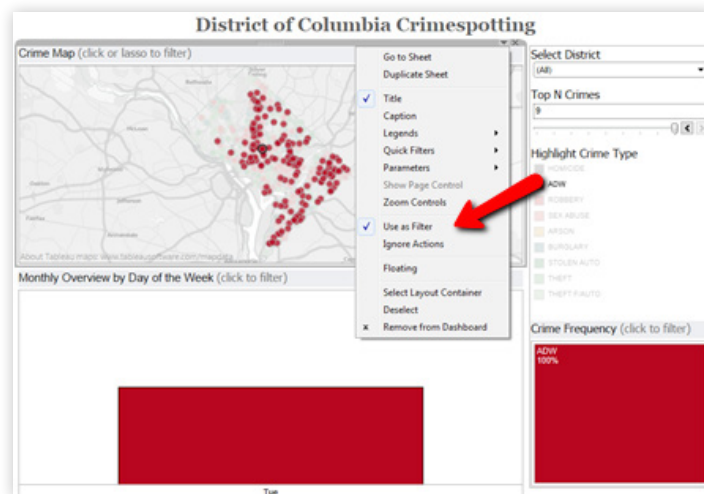
Filters let you slice data from different angles or drill down to a more detailed level. They are great ways to enable multi-level data exploration and user-driven data analysis. Tableau provides many options for building powerful filters right into your dashboards; however, yet there is always the potential to confuse your audience if you do not use filters properly. The following steps will help you create effective filters:

1. **Think about what you want your filters to do.** Before you start adding filters to your dashboard, it is often useful to ask yourself these questions: What are your views? How much flexibility do you want your users to have? Which filters will bring maximum value to your views? Are these filters already part of your views? What happens when you apply these filters? Are these filters necessary for your viewers to obtain certain information? How do your filters work together with your highlight actions? Do you want your filters to apply to one view, a few selected views, or to all of your views?
2. **Determine which types of filters to add.** Tableau provides four ways to add filters to your dashboards. Determining which types of filters to add usually depends on the answers to your previous questions. Below is a summary of the different filters and how they work:


- **Quick filters:** Applying a quick filter is the easiest way to add filters to your views. Simply right click a field or drag a field onto the filters shelf in your worksheet and then bring them up on the dashboard. In this example, notice that the filter is a sub-area of the default dashboard view –the filter is District within DC and the default geographical area is the entire DC. In other words, this filter adds great value to the dashboard by enabling users to drill down views from the entire District of Columbia into a subarea of the district. Also note that the default scope of this filter is to one view only. To apply it to additional views, you will need to manually adjust the filter.



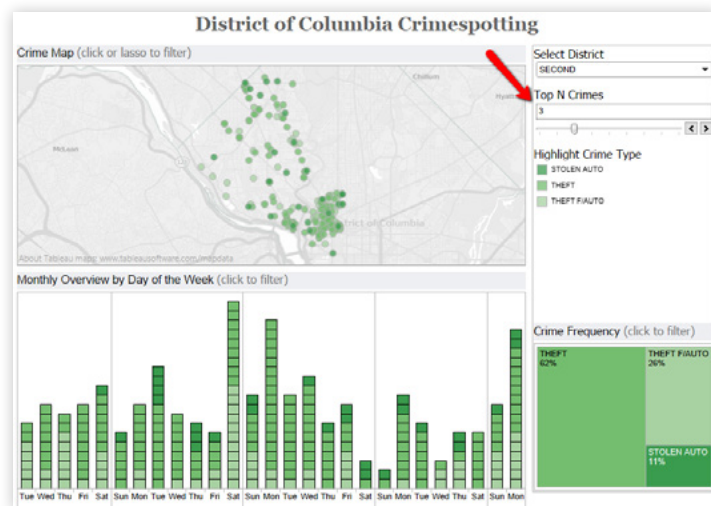
- **Use a view as a filter:** This is another easy way to add filters to your dashboard. Unlike with the quick filter, though, the default for this filter is to apply to all views and fields on the dashboard. Below is an example of what you might see if you select the “Use as Filter” option on the map view; in this case, the other views filter down by crime and day of the week when you select a data point on the map. By seeing this result, we realize that the value of seeing each crime on each day of the week is less informative than seeing each crime on all of the day of week. This is where the next filter type comes in handy.



- Filter action:** This option gives you great control and flexibility over your filters; you can choose your source views, target views, which fields to filter by, how your users activate the filters, and what you want to happen after you clear the filters. You can either create a filter action from the dashboard action menu or just edit the action that is automatically generated with the “Use as Filter” option. Below is the new and more informative view that you would see after changing your existing map filter into a filter by crime type.



- Filter with a parameter:** Compared to the previous three filter types, this option is the most powerful. Filtering with a parameter gives you the ability to make more flexible and interesting filters which can't be achieved by other type of filters. It even lets you filter across different data sources. Below is a simple example of a filter with a parameter; in this view, users can choose a top number of crimes to filter the view by.



In addition to these steps for creating effective filters, the following tips will help you optimize the efficacy of your filters:

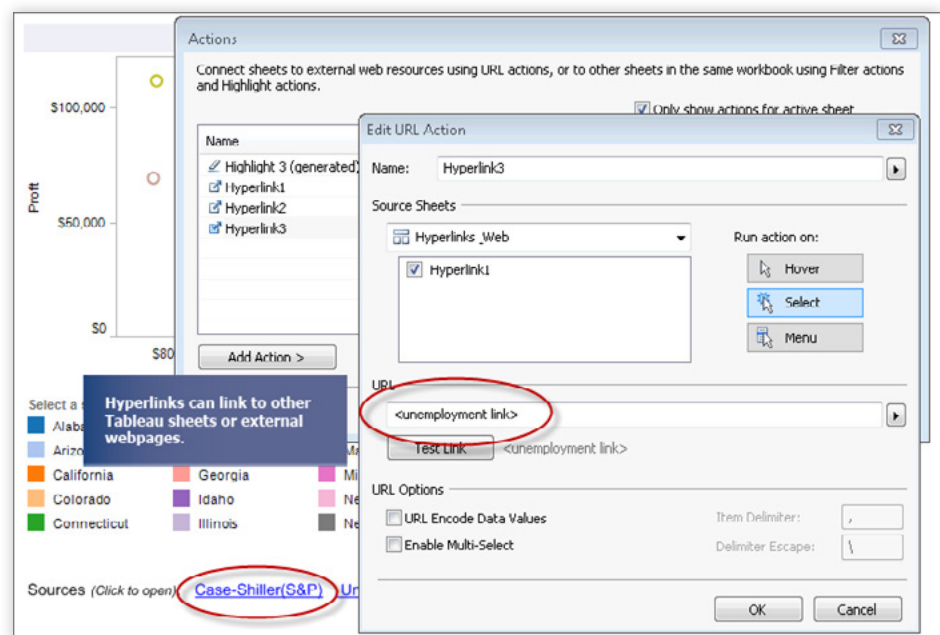
- Try to apply filters to all the views in one dashboard unless there is a strong reason to have separate filters for each view--in which cases filters should be placed as close as possible to the views it filters. In situations where you have multiple dashboards, a conscious decision has to be made as to how you would like your filters to apply to your views (whether each filter applies to all dashboards or each dashboard has its own separate set of filters).
- Arrange your filters in meaningful orders, such as by date, country, state, city, or business segment. Make sure you turn on the Show Less Value button if your filters have a cascading effect. For example, you have State as the first filter and then City as the second filter, you should turn on Show Less Value button on the City filter—that way, you will only see cities in the state your users select.
- Make sure the values in the quick filter are ordered in a way that makes sense for your data. For example, instead of listing classes alphabetically, you might order them by popularity. You can specify the order of a quick filter by setting the default sort order for that field.
- Add dynamic titles that represent the current filter selections. That way your users always know what views are being filtered as well as what the selections are.
- Remember that a field does not have to be in use in order for you to filter by it. In other words, you can have a bar chart showing the GDP of twenty countries, and then add a filter on population to show only those countries with a population over 100 million. These “slicing” filters are very powerful.
- Hide filters from your audience if they don’t add value to your view. Not all filters have to be shown to your audience. For example, a common way to clean up extraneous data is to exclude null values. That’s a filter that you probably don’t want to show on the dashboard.
- Consider including a “apply” button to your filter if the list of choices in your filter is long and each click results in automatically update which causes a lag for the users.
- Decide if you want users to be able to select “All” in your filter—or just be able to go to one. Tableau allows for both, so figure out which makes sense and set your filter settings accordingly.

- Remember that slider filters are great for date and numerical values—while list filters are better for categorical data.
- Always test your filters after applying them to your dashboards. Try out as many different combinations (including weird, nonsensical selections) just to make sure nothing bizarre happens with filter interactions. You can easily get into a situation where the combination of filters returns no result – that can be disorienting for users and we should try to prevent it.
- Don't forget to check the initial state of your views before publishing. Be deliberate about the first impression you create for your viewers. Sometimes it is also a good technique to publish with selections so users discover interactivity sooner. If a single point is already clicked, they might be more likely to interact.

Hyperlinking and using the power of the web

URL actions allow you to embed a hyperlink that points from your dashboards to a Web page, file, or other web-based resource outside of Tableau. You can use URL actions to link to information hosted outside of your data source. To make the link relevant to your data, try using values of your data as parameters in the URL. For example, if you have a list of Twitter users that are encoded in your data as the field <username>, you can create a URL action that points to www.twitter.com/<username>. By triggering that action, a viewer will be able to see the selected user's profile. It's important to note that the link can open a new webpage, or load directly in the dashboard as a web object.

You can use URL actions to link to information hosted outside of your data source.



No matter where you are publishing, be sure to construct your visualization at the size that you will eventually publish at—and use the Range sizing feature to avoid scrollbars or scrunched views.

Sizing: Making sure your visualization is visible

Tableau dashboards are set to a fixed default size that is intended to work well on a typical computer desktop. However, when you publish (to the web, in a blog, for a presentation, etc.) you may find yourself more limited. No matter where you are publishing, be sure to construct your visualization at the size that you will eventually publish at—and use the Range sizing feature to avoid scrollbars or scrunched views.


Scrollbars


When you republish a visualization at a different size than the original, you often end up with scroll bars like the ones shown below. If you are changing the size of a visualization that you have already embedded in a website, make sure to re-embed it with the new code. While scroll bars can be acceptable in some list views, in general you should avoid scroll bars in the major views of your visualization (as shown below).



77	\$961	\$10,731	\$51
61	\$733	\$6,922	\$87
40	\$3,621	\$11,050	\$50
37	\$585	\$1,985	\$14
25	\$1,359	\$5,859	\$28
16	(\$68)	\$488	\$7
105	\$863	\$6,919	\$93
54	(\$4)	\$721	\$16
228	\$1,561	\$6,504	\$53
208	\$3,419	\$21,836	\$180
66	(\$135)	\$989	\$22


Scrunched views

Be careful while resizing your views; you don't want them to become “scrunched.” You should always leave enough space for full headers and labels—and to show all of the data in an understandable way. Once you've ensured enough space in your visualization, you should clear any manual sizing you've added unless it is absolutely necessary. You can clear manual sizing using the Clear button on the toolbar (). Quite often, clearing manual sizing will also prevent unnecessary scroll bars in your views.



Fitting

You can use the Fit options on the toolbar to specify how each view fits within the window. Select from the following options:




- **Normal:** Automatically resizes based on the data in the view and the size of the window.
- **Fit Width:** Scales the view horizontally to fill the width of the window. Vertical scroll bars are allowed.
- **Fit Height:** Scales the view vertically to fill the height of the window. Horizontal scroll bars are allowed.
- **Entire View:** Scales both the width and height to fill the window. Scroll bars will not be shown; all data will be scaled to fit within the window.

Select the fit that is most appropriate for your data and how it will be filtered. For example, you might select “Fit Entire View” for a simple cross-tab showing a fixed set of data. That way, it will always fill the entire space allotted to it. On the other hand, you might choose to stay with “Normal” if you're working with a view that is sometimes filtered to a few values. That will prevent the marks from being stretched to fill too large an area.

Perfecting Your Work

Although it might seem like a small detail, the formatting of your views can change everything about them. This is the most nit-picky part of making a visualization; often the only way to succeed is through trial and error. But if you put your mind to it and follow these steps, you will end up with a great-looking visualization.




Color TV looks better than black and white: Use color visualizations

Color can make the difference between a boring visualization and an inspiring one.



The following tips will help you create effect visualizations with color:

- Try to use no more than two color palettes. Make sure to use non-overlapping scales like the ones shown below.



- Consider changing the default grey title background (shown below) to white, or to a light color that does not clash with your other schemes.



- Consider how your use of color will be interpreted. Select semantically meaningful colors if they apply to the context of your data. If there are no relevant colors in the context of your data, use the Tableau color sets; they have been carefully selected to match and not clash. Once you have selected your color palettes(s), consider whether any of the colors in it have alternate meanings that do not align with your message. When using colors that have an inherent meaning, make sure you've assigned them to relevant values in your database. For example, in many cultures green is associated with positivity while red has a negative connotation.
- If the meaning behind your color choice isn't obvious, or your visualization does not obviously label the color, make sure to include a legend.
- When using a diverging color palette, the midpoint and end points should be meaningful. Zero is often a meaningful midpoint.
- Avoid adding color encoding to more than 12 distinct values.

Fonts: *Can you read this?* (Can you read this?)

Although there are dozens of fonts available in Tableau, there are only a few that you should use to optimize readability online. The following have been selected by our resident “visualization wizards” for their readability and visual appeal:

- Trebuchet MS or Verdana (especially for tables and numbers)
- Arial
- Georgia
- Tahoma
- Times New Roman
- Lucida sans

In addition, **Calibri** and **Cambria** are suitable for tooltips (see below), but are not recommended for use in any other part of a visualization.

In addition, it’s important to consider the color of your fonts. As a general rule, axes and labels should be dark grey (this keeps them from distracting viewers’ attention away from the visualization). Try to keep it to 2-3 (Font) colors per page.

If you use different fonts and styles throughout your visualization, make sure to check back over it to make sure the formatting is consistent. For instance, all of the filters should have the same style, as well as all of the titles—but the filters and titles can have different styles themselves.

Lastly, never make a change in adjacent text that modifies more than one attribute of a font (such as size, boldness, color, or serif quality).

Good Change

Bad Change



Country: United States
Number of Records: 1365
Value: 365

United States
 Number of Records: 1365
 Value: 365

United States
 Planes: **1365**
 Average Price: **\$365M**

Tooltips tell the story

Tooltips—the text boxes that pop up when you hover over an object—can make the difference between a user loving your visualization and not understanding it. Tableau automatically includes everything that might be possibly relevant – that means it repeats a lot of values you already have in your visualization. To modify it, go to “Edit” and then select “Tooltip.”

The following step-by-step example provides tips for improving a basic tooltip:

1. Begin with the basic tooltip. Use the front that anti-aliases properly (or does not get pixelated) online. Calibri and Cambria typically work well, but the default, Arial, is not bad either.
2. Next, identify the most important part of the tool tip and make that your title. In this example, the United States is clearly the subject of the tooltip, so it has been bolded and displayed in 16-point font. You can also add data to the title when applicable. For instance, if we were looking at states as well as countries, we could format the title to read “United States – Wyoming.” Think of it as you would a sports card (e.g., “Kobe Bryant – Los Angeles Lakers”).
3. Next, change the measure names to make them specific and understandable. “Number of records” does not mean much, but “Number of Planes” is very descriptive. “Value” is also abroad term that can raise more questions than answers. “Average Price could be better.”
4. Last but not least, make sure that units are included for all numbers in your tooltips.

Tooltips are a great place to add annotations or notes to your visualization (e.g., “This data point for 2009 only”). It is best to format notes slightly differently than the rest of your tooltip, such as with a lighter shade and slightly smaller font.

Make intelligent axes or you'll just have pretty rectangles

Although Tableau provides a great starting point for your visualization, axes are so important to the analytical experience that they require some deliberate attention. If your viewers do not have a good frame of reference, they will be flying blind (see the view below). The following components can help you create intelligent axes:

- **Fixed axes:** By default, the axis range automatically adjusts based on the data in your view. If that view is going to be filtered and changed (such as with quick filters or filter actions), your audience might not notice the resulting change in the axis range, and could therefore be misled. Changing axes also makes visual comparison very difficult. You can set the axis to a specific, fixed range to avoid any potential confusion.
- **Axis gridlines:** If it isn't practical to create fixed axes (such as if you're dealing with a large range of data), add gridlines to the view. Doing so can signal to the reader that the axes are changing. Reference lines can also be very effective at achieving the same result. However, they should be kept relatively subtle—you don't want to distract users from key elements of your view.
- **Axis labels:** Make sure that axis labels are appropriate and include units when necessary.



- **Axis tick labels:** Make sure that the values on your tick marks are also formatted properly (e.g., currency should have a symbol and the right number of decimal places).


Labeling your visualization in all the right places

Mark labels (the labels on your data points) can help you tell your story quickly and succinctly. It is often much easier to read a mark label than to mouse over a data point for its tooltip. Select **Format > Mark Labels** to turn on labels.

You can select from the following mark label options:


- **Labels on selection:** This labels the selected marks in the view. If several marks that are close together, you may want to select a different option to avoid clutter. This also applies to labels on highlight.
- **Labels on min/max:** This labels the outliers by marking the minimum and maximum values in the view.
- **Labels on highlight:** This labels the highlighted marks in the view.
- **Labels on line ends:** This labels the ends of the lines in the view. You can specify which ends to label (just the beginning, just the end, or both ends).

It is often much easier to read a mark label than to mouse over a data point for its tooltip.



Evaluating Your Masterpiece

Once you've designed each important component of your visualization, it's time to take a step back and evaluate your masterpiece. Are all the elements of your design working together well?



Use these questions like a check list to make sure:

What questions are you trying to answer? Pages 4-5

- Does this visualization answer all of your questions?
- Is the purpose of the visualization clearly explained in its title or surrounding text?
- Can you understand the visualization in 30 seconds or less, without additional information?
- Does your visualization include a title? Is that title simple, informative, and eye-catching?
- Does your visualization include subtitles to guide your viewers?

Do you have the right chart type for your analysis? Pages 5-12

- What types of analysis are you performing?
- Have you selected the most suitable chart type(s) for your types of analysis?
- Have you considered alternative chart types that could work better than the ones you have chosen?

Are your views effective? Pages 12-19

- Are your most important data shown on the X- and Y-axes and your less important data encoded in color or shape attributes?
- Are your views oriented intuitively—do they cater to the way your viewers read and perceive data?
- Have you limited the number of measures or dimensions in a single view so that your users can see your data?
- Have you limited your usage of colors and shapes so that your users can distinguish them and see patterns?

Your effort and best judgment are critical to producing great visualizations.

Is your dashboard holistic? Pages 19-29

- Do all your views fit together to tell a single story?
- Do all your views flow well from one to the next? Are they in a good order?
- Do your most important views appear in the top or top-left corner?
- Are secondary elements in your dashboard placed well so they support the views without interrupting them?
- Are your filters in the right locations?
- Do your filters work correctly? Do views become blank or downright confusing if you apply a filter?
- Do your filters apply to the right scope?
- Are your filter titles informative? Can viewers easily understand how to interactive with your filters?
- Are your legends close to the views they apply to?
- Is your legend highlight button set to “on” or “off” according to your preference?
- Do you have filter, highlight or URL actions? If so, do they work?
- Are your legends and filters grouped and placed intuitively?
- Do you have scrollbars in your views? If so, are they acceptable ones?
- Are your views scrunched?
- Do your views fit consistently well when you apply filters?


Did you perfect your work? Pages 19-29

- Do all the colors on your dashboard go together without clashing?
- Do you have less than 7-10 colors on your dashboards?
- Do you use fonts consistently in all of your views and no more than three different fonts on one dashboard?
- Are you labels clear and concise? Are they placed optimally to help guide your viewers? Make sure subtitles are formatted to be subordinate to the main title.
- Are you tooltips informative? Do they have the right format so that they’re easy for viewers to use?

In conclusion, your effort and best judgment are critical to producing great visualizations. By following these visual best practices, not only can you produce highly useful and effective visualizations, but you can also make your work visually appealing to your users. Congratulations! You are now on the way to become a true data visualization master.


About Tableau

Tableau Software helps people see and understand data. Tableau helps anyone quickly analyze, visualize and share information. More than 17,000 customer accounts get rapid results with Tableau in the office and on-the-go. And tens of thousands of people use Tableau Public to share data in their blogs and websites. See how Tableau can help you by downloading the free trial at www.tableausoftware.com/trial.



Additional Resources

[Download Free Trial](#)




Related Whitepapers

[Which chart or graph is right for you?](#)

[Selectinga Visual Analytics Application](#)

[See All Whitepapers](#)



Explore Other Resources

- [Product Demo](#)
- [Training & Tutorials](#)
- [Community & Support](#)
- [Customer Stories](#)
- [Solutions](#)

