



Choosing the Right Chart and Color

A practical field guide to matching charts to questions and
using color that carries meaning

A PRIMAFORJA FIELD GUIDE

PrimaForja · info@primaforja.com · www.primaforja.com

Most charts fail before a single number is plotted. They fail at the choice: the wrong form for the question, a color scheme that fights the data, an axis that misleads. The math can be flawless and the chart still lie — or, more commonly, just bore the reader into missing the point.

This guide is a practical, tool-agnostic reference for making that first choice well. It covers how to match a chart to the question you're answering, a catalog of the common chart types and what each is actually good for, how to use color so it carries meaning instead of noise, and the craft details that separate a chart people trust from one they squint at. None of it is specific to any one product or dataset — it's the working craft of visualization, in one place.

Keep it on your desk, share it with your team, and use it the next time you're staring at a blank panel wondering what to put there.

Part 1 — Start with the question

The single most useful habit in data visualization is this: **decide what question the chart answers before you decide what it looks like**. The form follows the question. If you can't state the question in one sentence, no chart type will save you.

Almost every analytical question falls into one of a handful of families. Name the family, and the shortlist of good charts writes itself.

The question you're asking	Question family	Reach for	Avoid
"Who is biggest / smallest?"	Comparison	Bar / column, dot plot	Pie, 3-D anything
"Which way is it trending?"	Change over time	Line, area, column (few periods)	Pie, radar
"What's it made of?"	Part-to-whole	Stacked bar, treemap, one donut	Many pies, exploded pie
"How spread out is it?"	Distribution	Histogram, box plot, violin	Bar of averages alone
"Do these move together?"	Relationship	Scatter, bubble	Dual-axis line
"Where is it happening?"	Geospatial	Choropleth, symbol map	3-D map, rainbow map
"How does it flow / convert?"	Flow & process	Sankey, funnel, waterfall	Stacked bar for stages
"What's the current number?"	Single value / KPI	Big number, bullet, gauge (sparingly)	Gauge farms
"What's driving the total?"	Ranked contribution	Pareto, waterfall	Unsorted bar

A few principles sit above the table:

- **One panel, one question.** If a chart is trying to answer two questions, split it into two charts. Overloaded panels are the most common cause of an unreadable dashboard.
 - **Comparison beats reading values.** People are far better at judging *longer vs. shorter* and *higher vs. lower* than at reading exact numbers off a chart. If the exact number matters, add a label or use a table.
 - **Position and length are the strongest encodings.** We read position along a common axis (scatter, dot plot) and length (bars) most accurately; then angle, area, and color; then shading and volume. Prefer the strong channels for the values that matter most.
-

Part 2 — The chart catalog

For each chart: what it is, what it's best for, and the mistakes that most often ruin it.

Bar and column charts

What it is. Rectangles whose length encodes a value. Columns are vertical; bars are horizontal. The workhorse of comparison.

Best for. Comparing a value across categories — sales by region, headcount by department, error rate by service. Horizontal bars are better when you have many categories or long labels (labels sit flat and readable).

Watch out for.

- **Bars must start at zero.** Length is the encoding; a truncated baseline exaggerates differences and is the most common way charts mislead.
- **Sort by value, not alphabetically,** unless the categories have a natural order (months, sizes). A sorted bar answers "who's biggest?" instantly.
- Cap the number of bars. Past ~15–20, switch to a horizontal layout, a Pareto, or a "top N + Other" grouping.

Grouped (clustered) bar charts

What it is. Bars for two or three sub-categories placed side by side within each main category.

Best for. Comparing a small number of series across categories — this year vs. last year by region; plan vs. actual by month.

Watch out for. More than three series per group turns into a picket fence. If you have many series, use small multiples (one mini-chart per series) instead.

Stacked bar and 100% stacked bar

What it is. Sub-values stacked into a single bar. Plain stacked shows totals *and* composition; 100% stacked shows composition only (every bar fills the height).

Best for. Part-to-whole across a few categories or time periods — revenue mix by channel across quarters.

Watch out for. Only the bottom segment sits on a common baseline, so all the others are hard to compare precisely. Put the series you most want to compare on the bottom. Keep segments to ~5 or fewer. If precise comparison of the middle segments matters, use grouped bars or small multiples instead.

Line charts

What it is. Points connected in order, almost always over time.

Best for. Trends and change over time — revenue by month, latency over the day, active users by week. Excellent for showing several series' trajectories at once (keep it to ~4–5 lines before it tangles).

Watch out for.

- The x-axis should be continuous and evenly spaced (time, or an ordered quantity). Don't use a line for unordered categories — that implies a progression that doesn't exist.
- A zero baseline is **optional** for lines: the shape of the trend is the point, and a zoomed y-range is legitimate — but label it clearly so no one misreads the magnitude.
- Too many lines = spaghetti. Highlight one or two, gray the rest, or split into small multiples.

Area and stacked area charts

What it is. A line with the region beneath it filled. Stacked area sums several series.

Best for. A single volume over time, or part-to-whole *that changes over time* (stacked area) — like traffic sources by month.

Watch out for. Stacked areas share the bottom-baseline problem: upper bands are hard to read. Use them for the gestalt ("the mix is shifting toward mobile"), not for reading precise values off the middle bands. Never overlap filled areas with transparency and expect anyone to decode it.

Combo / dual-axis charts

What it is. Two chart types (often bars + line) on one plot, sometimes with two different y-axes.

Best for. One genuinely useful case: a measure and its rate — revenue (bars) with margin % (line) — where they share an x-axis and the pairing is meaningful.

Watch out for. Dual axes let you imply any correlation you like by rescaling one axis; they're easy to abuse and hard to read. Use them sparingly, label both axes unmistakably, and when in doubt split into two aligned charts instead.

Pie and donut charts

What it is. A circle sliced by proportion. A donut is a pie with the center removed (handy for a total in the hole).

Best for. A single part-to-whole snapshot with **very few slices** (ideally 2–4), where you want to signal "this is a share of a whole" at a glance.

Watch out for. People judge angles and areas poorly. Beyond a few slices, a sorted bar is always easier to read. Never use multiple pies to compare across categories or time (that's a stacked bar or small multiples),

never explode the slices, never go 3-D. If your pie needs a legend and percentage labels to be readable, it wanted to be a bar chart.

Scatter plots

What it is. Points positioned by two quantitative variables (x and y).

Best for. Relationships and correlation — price vs. volume, defect rate vs. delivery time, spend vs. return. It's the best way to reveal clusters, outliers, and the shape of a relationship.

Watch out for. With thousands of points, overplotting hides density — use transparency, smaller marks, or a density/hex-bin plot. Add a trend line only if a linear relationship is actually plausible.

Bubble charts

What it is. A scatter plot where a third variable sets the size of each mark (and sometimes a fourth sets color).

Best for. Three-variable comparison — vendors plotted by quality vs. speed, sized by spend.

Watch out for. Size is a weak, easily-overestimated encoding; **scale bubbles by area, not diameter**, or big values look wildly bigger than they are. Keep the size range modest and add a size legend.

Histograms

What it is. Bars showing how many observations fall into each range (bin) of a single continuous variable. Looks like a bar chart but means something different: it's a distribution, not a comparison of categories.

Best for. Understanding the shape of one variable — order sizes, response times, ages. Reveals skew, gaps, and multiple peaks that an average hides.

Watch out for. Bin width changes the story; try a few. Don't confuse it with a bar chart — the x-axis is a continuous range, and the bars touch.

Box plots and violin plots

What it is. A compact summary of a distribution: median, quartiles, and outliers (box), optionally with the full density shape (violin).

Best for. Comparing distributions across categories — cycle time by team, salary by role — especially when you care about spread and outliers, not just the average.

Watch out for. Box plots are unfamiliar to many business audiences; label the parts or add a one-line legend. For small samples, show the actual points (a strip/jitter plot) instead of hiding them behind a box.

Heatmaps

What it is. A grid of cells colored by value — a table where color does the talking.

Best for. Patterns across two categorical dimensions — activity by hour × day of week, cohort retention, correlation matrices.

Watch out for. This is where color choice is everything (see Part 3). Use a single sequential ramp for magnitude; a diverging ramp only if there's a meaningful midpoint. Keep the grid legible — too many cells

and it's wallpaper.

Maps: choropleth and symbol maps

What it is. A **choropleth** shades geographic regions by value. A **symbol map** places sized/colored marks at locations.

Best for. Values that are genuinely geographic and that your audience thinks about spatially — sales by state, outage density by metro.

Watch out for. Choropleths are biased by area: big, empty regions dominate the eye. Normalize to a rate (per capita, per store) rather than raw counts, and consider a symbol map when the geography is about points, not areas. A map is worth it only when *location itself* is the insight — otherwise a sorted bar of regions is clearer.

Waterfall charts

What it is. A running total that rises and falls through a sequence of positive and negative steps.

Best for. Showing how you got from a starting value to an ending one — opening balance → inflows → outflows → closing; list price → discounts → net price → cost → margin.

Watch out for. Keep the step order meaningful (chronological or logical), and color increases and decreases with your reserved good/bad or up/down pair so the direction reads instantly.

Funnel charts

What it is. Stages of a process shown as progressively narrowing bars.

Best for. Conversion through an ordered pipeline — visitors → signups → trials → paid; leads → qualified → won.

Watch out for. Only valid when stages are genuinely sequential and monotonic (each stage is a subset of the last). Show the conversion rate *between* stages — that's the number people actually want.

Pareto charts

What it is. Sorted bars (descending) with a cumulative-percentage line overlaid — a purpose-built combo chart.

Best for. The 80/20 question: which few categories drive most of the total? Defect causes, revenue concentration, cost drivers.

Watch out for. It only tells a clean story when the distribution is actually concentrated. If the cumulative line rises evenly, say so rather than forcing a "vital few" narrative.

Gauges and bullet charts

What it is. A **gauge** shows one value against a range, dial-style. A **bullet chart** does the same job in a compact bar with target and qualitative bands — a far more efficient design.

Best for. A single measure against a target or threshold — SLA attainment, quota progress.

Watch out for. Gauges are seductive and space-hungry, and a wall of them says little. Prefer bullet charts, or a KPI tile with a target and a variance. Reserve gauges for the rare hero metric.

KPI tiles and sparklines

What it is. A large single number, usually with a comparison (vs. target, vs. prior period) and often a **sparkline** — a tiny, axis-free trend line — beside it.

Best for. The handful of numbers that matter most, at the top of a dashboard. The big number gives the "what"; the delta gives the "vs. what"; the sparkline gives the "how did we get here."

Watch out for. A number without a comparison is trivia. Always pair it with a reference (target, last period, benchmark) so the reader knows if it's good.

Treemaps

What it is. Nested rectangles sized (and often colored) by value — a space-filling part-to-whole.

Best for. Hierarchical composition with many items — product families and their SKUs, spend by department and team — where a pie would have too many slices.

Watch out for. Precise comparison of similarly-sized rectangles is hard. Treemaps are for the overview ("where does the mass sit?"), not for reading exact values. Label the big tiles; let the small ones be texture.

Sankey and flow diagrams

What it is. Bands whose width encodes a quantity flowing between stages or categories, splitting and merging along the way.

Best for. Flows and transfers — budget from sources to uses, users moving between states, energy or material flows.

Watch out for. Beautiful and easily overcrowded. Limit the nodes and flows, order them to minimize crossings, and reserve it for when the *flow* is the story.

Tables and heat tables

What it is. Rows and columns of values. A **heat table** adds background shading or in-cell bars to guide the eye.

Best for. When people need the **exact numbers**, want to look up a specific row, or the data is irreducibly multi-dimensional. Tables are underrated: they're often the honest answer.

Watch out for. A raw grid buries the story. Add light shading, in-cell bars, right-align numbers, use consistent decimals, and sort by what matters. A well-designed table beats a bad chart every time.

Small multiples (a technique, not a chart)

What it is. A grid of many small charts that share axes and scales, one per category.

Best for. Comparing a pattern across many groups — the same trend line for each of 12 regions. Once you learn to read one panel, you can read all of them, and outliers jump out.

Watch out for. The scales *must* be shared, or the comparison is a lie. Keep each panel minimal — one shared legend, sparse axes.

Part 3 — Color that means something

Color is the most abused channel in analytics. Used well, it encodes a variable or draws the eye to what matters. Used badly, it's decoration that adds noise and excludes readers. The governing rule: **encode, don't decorate. If color isn't carrying information, use less of it.**

Three palette types, three jobs

Matching the palette type to the data type is the whole game.

-  **Categorical (qualitative)** — for *unordered groups*: regions, products, channels. Use distinct hues of similar intensity, so no category looks more important than another. Cap it at **~6–8 colors**; beyond that, group into "Other" or switch encoding.
-  **Sequential** — for an *ordered quantity*: values on a map, intensity in a heatmap. A single hue ramped light → dark, where darker means more. Never a rainbow ramp — rainbows invent boundaries the data doesn't have and rank hues people don't rank.
-  **Diverging** — for a value with a *meaningful midpoint*: budget variance, forecast bias, change vs. last year, margin above/below target. Two hues meeting at a neutral center, so both direction and magnitude read at once.

The rules that save most dashboards

- **Reuse the same hue for the same thing, everywhere.** If Enterprise is teal on one panel, it's teal on every panel. Inconsistent category colors force the reader to re-learn the legend on each chart.
- **Reserve a semantic pair.** Keep one color for *good / on-target / revenue* and one for *bad / off-target / cost / risk* — and never use those two for anything else. If red is also "the third product category," red stops meaning "attention."
- **Gray is a color — use it.** Push everything that isn't the point to gray and let one accent color carry the message. The fastest way to make a chart clearer is usually to *remove* color, not add it.
- **Mind cultural meaning.** Red/green as bad/good is common in the West but not universal (and red is positive in some markets). Don't rely on a single color's connotation to carry meaning without a label.

Accessibility is not optional

- **Color-blindness is common.** Roughly **8% of men** (and ~0.5% of women) have some color-vision deficiency, most often red–green. A red-vs-green chart is invisible to a meaningful slice of your audience.
- **Never rely on color alone.** Reinforce it with position, shape, direct labels, or text. The test: *if the chart were printed in grayscale, would it still make sense?* If not, fix it.
- **Use color-blind-safe palettes.** Blue–orange diverges safely where red–green fails; established safe palettes (e.g. ColorBrewer, Okabe–Ito, Viridis for sequential) are free and well-tested. Reach for one rather than hand-picking.


- **Contrast matters.** Keep text and key marks well above the WCAG contrast minimums against their background. Pale gray on white may look elegant on your monitor and vanish on a projector.

Backgrounds, themes, and dark mode

- Pick a light or dark theme and keep it consistent; many analysts live in dark mode, so if you ship dark, make sure your palette still separates cleanly on a dark background (saturated colors can vibrate — desaturate slightly).
- Keep backgrounds neutral. Colored or textured backgrounds shift how every other color is perceived.
- Semi-transparency is for overlap (scatter density, overlapping areas), not for making a busy chart "softer."

A starter palette

You don't need a custom palette to do good work — but a small, deliberate one helps. A workable kit:

Role	Swatches
Categorical (up to 6)	
Sequential	
Diverging	
Semantic good / bad	
Neutral / de-emphasis	

Part 4 — The craft details

The chart type and palette get you most of the way. These details decide whether the result is trustworthy and clear.

- **Zero baselines for bars, always.** Length is the encoding. Truncating a bar axis exaggerates differences and is the classic way to mislead. Lines and dot plots may use a zoomed range — the *shape* is the message there — but label it.
- **Sort for meaning.** Descending by value for comparisons; natural order for time or ordered categories. Alphabetical is almost never what the reader needs.
- **Maximize data-ink; cut chartjunk.** Remove heavy gridlines, boxes, redundant legends, drop shadows, and anything 3-D. Every drop of ink should carry information. When in doubt, take it out.
- **Direct-label instead of forcing a legend hunt.** Put the series name at the end of its line, the value on top of the bar. A legend is a lookup table; a direct label is the answer.
- **Annotate the "so what."** The best charts tell you what to notice: a callout on the spike, a reference line for target, a short title that states the finding ("Margin recovered in Q3") rather than the topic ("Margin by Quarter").
- **Titles should state the takeaway.** A chart title is prime real estate. Use it for the conclusion, and let the subtitle or axis carry the mechanics.

- **Format numbers like a human.** Round to the precision that matters (you rarely need six significant figures), use thousands separators, show units once, and keep decimals consistent within a column.
- **Mind the aspect ratio.** A trend's perceived steepness depends on the chart's shape. Very wide or very tall distorts the slope; aim for a ratio where the typical line segment sits near 45°.
- **Consistency across the dashboard, variety across question types.** When two panels answer the *same kind* of question, make them identical so they're instantly comparable. When they answer *different* questions, deliberately vary the form — a trend line next to a ranked bar next to a KPI tile — so each panel announces its kind before the reader parses a label.
- **Design for the medium.** A chart for a phone, a boardroom projector, and a printed report are three different charts. Check legibility at the size and contrast it will actually be seen.

Part 5 — Quick reference

Chart chooser cheat-sheet

I want to show...	First choice	Also consider	Skip
Ranking across categories	Sorted bar	Dot plot	Pie
A trend over time	Line	Column (few periods), area	Pie
Composition at one moment	Stacked bar / treemap	One donut (≤ 4 slices)	Many pies
Composition changing over time	Stacked area / stacked bar	100% stacked bar	Pie series
The shape of one variable	Histogram	Density plot	Bar of the average
Compare distributions	Box plot	Violin, strip plot	Bars of averages
Relationship between two measures	Scatter	Bubble (add a 3rd)	Dual-axis line
Something geographic	Choropleth (rates)	Symbol map	3-D / rainbow map
A process converting	Funnel	Sankey	Stacked bar
Building to a total	Waterfall	—	Stacked bar
The vital few	Pareto	Sorted bar	Unsorted bar
One number vs. target	KPI tile + delta	Bullet chart	Gauge farm
Exact values / lookups	Table (with light shading)	Heat table	A chart

Color checklist

- ☐ Palette type matches data type (categorical / sequential / diverging)
- ☐ $\leq 6-8$ categorical colors; same hue = same category everywhere
- ☐ A good/bad pair is reserved and used for nothing else
- ☐ No rainbow ramp for an ordered quantity
- ☐ Readable for color-blind viewers and in grayscale
- ☐ Color encodes something — nothing is colored just to decorate

Pre-publish QA checklist

- ☐ Each panel answers one clearly stated question
- ☐ Bars start at zero; any zoomed axis is labeled
- ☐ Sorted for meaning, not alphabetically
- ☐ Titles state the takeaway; key marks are directly labeled
- ☐ Numbers rounded and formatted consistently, units shown
- ☐ Chartjunk removed; no 3-D
- ☐ Legible at the size and on the device it will be viewed
- ☐ The chart still makes sense in grayscale

The whole thing, in one sentence

Pick the chart that answers your question, spend a small and consistent color system on meaning rather than decoration, start your bars at zero, and give the ink to the data. Do that, and most of the ways a chart can fail are already behind you.

This field guide was produced by PrimaForja. We build production-scale synthetic datasets and the dashboards that prove them out — every chart designed, and every metric defined, to the standard described here. Questions, or want to talk data visualization and synthetic data? Reach us at info@primaforja.com or visit www.primaforja.com.