


PLAYBOOK

# Designing Dashboards with AI

The best thinking on dashboard design, turned into a method you can practise, with AI doing the heavy lifting and you keeping the judgement.

- D** Decide
- E** Establish metrics
- S** Structure the layout
- I** Illustrate
- G** Guard
- N** Nurture



# Start where you are.

This playbook distils the best-known work on dashboard and visualisation design into one method, then shows exactly how to work with an AI assistant at every step. Pick the path that fits today.

## Starting a new dashboard

Walk the DESIGN loop, page 5

Most dashboards go wrong before the first chart.

## Fixing one that isn't used

Score it with the critic, page 18

Then declutter, page 17.

## Building it with AI

Read the pact, page 19

Then borrow the conversations.

## Wanting the ideas behind it

Open the library, page 24

Every idea, with where it comes from.

- 
- |       |                                                                    |       |
|-------|--------------------------------------------------------------------|-------|
| 1     | <b>Five seconds</b>                                                | p. 3  |
|       | Why dashboards die, a test, and the same data redesigned           |       |
| <hr/> |                                                                    |       |
| 2     | <b>The DESIGN loop</b>                                             | p. 5  |
|       | Follow one dashboard from blank page to launch, a stage per page   |       |
| <hr/> |                                                                    |       |
| 3     | <b>The field guide</b>                                             | p. 12 |
|       | The loop at a glance, four kinds, the metric card, layout          |       |
| <hr/> |                                                                    |       |
| 4     | <b>The workshop</b>                                                | p. 15 |
|       | Spot the flaws, chart from question, declutter, colour, the critic |       |
| <hr/> |                                                                    |       |
| 5     | <b>Working with AI</b>                                             | p. 19 |
|       | Eight roles, six promises, and where AI lives in your tools        |       |
| <hr/> |                                                                    |       |
| 6     | <b>Conversation starters</b>                                       | p. 20 |
|       | Eighteen openings, each with a push-back                           |       |
| <hr/> |                                                                    |       |
| 7     | <b>DESIGN-ready?</b>                                               | p. 22 |
|       | The launch checklist and the playbook on a page                    |       |
| <hr/> |                                                                    |       |
| 8     | <b>The library of ideas</b>                                        | p. 24 |
|       | Thirty-one concepts and the books behind them                      |       |
| <hr/> |                                                                    |       |

## 1 WHY THIS MATTERS NOW

# AI makes dashboards fast. It also makes bad ones fast.

Most dashboards don't fail loudly. They launch, get opened twice, and quietly die. The causes are almost always the same three.

1

## Built for data, not decisions

Every metric goes on the page because it exists, not because anyone acts on it.

2

## Too much, too clever

Twenty charts, five colours, three fonts. The one number that matters gets lost.

3

## Numbers nobody trusts

No definitions, no source, no date. One wrong figure and people return to spreadsheets.

### FRAMEWORKS GIVE YOU

## The right dashboard

Decades of research on how people read charts, find answers and make decisions.

### AI GIVES YOU

## Built and checked, faster

Drafts, critiques, code and accessibility checks in minutes, once you tell it what good looks like.

## How this playbook works

You start with a test of your own eye, then follow one dashboard, a public helpline called Sahayak, through six stages of the DESIGN loop. The field guide keeps the reference pages together. The workshop lets you practise. The pact sets out what to hand to AI and what to keep. The launch checklist and the library close the book.


*Whatever your field, sales, health, schools, climate, public services, the loop is the same. Only the examples change.*

## 1 BEGIN WITH A TEST

# You have five seconds.


That is roughly how long a busy reader gives a dashboard before deciding whether it is worth their time. Cover the right-hand dashboard, look at the left one for five seconds, and answer: which region needs help most?

*before: everything, all at once*



**On the left,** the answer is in a text table, sixth row. Most readers have to hunt, and many give up. The design failed them, not the other way round.


*after: one finding, then the evidence*



**On the right,** the answer is the first thing you read. Same data, same truth. The difference is the method on the next pages.

*Run this test with a colleague before every launch. If they can't recall the main point after five seconds, the hierarchy needs work.*

Meet Sahayak, a public helpline that takes thousands of requests a month. Its managers asked for "a dashboard". Over the next six pages, watch it take shape one stage at a time. The same steps work for a sales team, a clinic network or a school system.



THE DESIGN LOOP · STAGE 1 OF 6

## Decide

Who decides what, how often?

Sahayak's managers asked for "everything about calls". That is a data request, not a purpose. The first job is to find the decision the dashboard serves, and the people who make it.

Write one sentence: this dashboard helps [who] decide [what], every [when]. For Sahayak it became: district managers decide where to send extra staff, every Monday. Then list the four or five questions they ask on that Monday.

Sahayak, so far

*just a sentence*

### Sahayak helpline

This dashboard helps district managers decide where to send extra staff, every Monday.

Are we on time?

Where is it worst?

Is the backlog growing?

Who is waiting longest?

**What changed:** A purpose sentence and the four questions managers ask. No charts yet.

### HOW THE CONVERSATION WITH AI WENT

YOU

Interview me, one question at a time, to find the purpose of a dashboard for our helpline managers.

AI, A TYPICAL REPLY

Who opens this dashboard first on a typical Monday, and what do they do differently after reading it?

YOU

Good. Now push me: what am I assuming about these managers that I haven't checked?

*An AI can interview you. It cannot interview your readers. Talk to two of them before you design anything.*

The ideas behind this: [Utilization-focused evaluation](#) · [Jobs to be done](#) · [The nested model](#)

### Your turn

for your own dashboard

- ☐ Write: this helps [who] decide [what] every [cadence]
- ☐ List three to five questions in the readers' words

- ☐ Pick the kind: strategic, operational, analytical or communicative
- ☐ Interview two real readers first



THE DESIGN LOOP · STAGE 2 OF 6

# Establish metrics


Which numbers answer those questions, exactly?

Each question needs a number, and each number needs a written definition. "On time" sounds obvious until two teams compute it differently. Sahayak wrote it down: requests closed within seven days, divided by all requests closed.

Then pair each result with an early warning. The on-time rate tells you how last month went; the size of the backlog tells you what next month will look like.

Sahayak, so far

four numbers, each defined



**What changed:** Four tiles, each with a written definition underneath.

## HOW THE CONVERSATION WITH AI WENT

YOU

For each of these four questions, propose one metric with a definition, formula, data source and a leading indicator. Flag anything our data can't support.

AI, A TYPICAL REPLY

On time: closed within 7 days ÷ all closed, from the case export. Leading indicator: open requests older than 2 days. I can't compute "who waits longest" by applicant group; the export has no group field.

YOU

Then we add that field next quarter. Write the spec cards and mark that metric as planned.

**Mind Goodhart's law. Once "on time" is a target, requests may get closed early. Track reopened requests too.**

The ideas behind this: [SMART indicators](#) · [Leading and lagging indicators](#) · [Goodhart's law and Campbell's law](#) · [Theory of change and results chains](#)

## Your turn

for your own dashboard

- ☐ Derive metrics from your theory of change, logframe or OKRs
- ☐ One card per metric: definition, formula, source, owner
- ☐ Pair each lagging outcome with a leading indicator
- ☐ Plan equity breakdowns before collecting data



THE DESIGN LOOP · STAGE 3 OF 6

# Structure the layout


What does the reader see first, second, third?

Readers scan from the top left, so the answer goes there, the evidence below it and the detail last: an inverted pyramid. Shneiderman said it as overview first, zoom and filter, then details on demand.

Sahayak's wireframe has one headline sentence, four tiles, a trend and a breakdown. Staff rosters and call durations moved to a second tab that nobody is forced to read.

Sahayak, so far

*a wireframe, nothing more*



**What changed:** Grey boxes only: headline, tiles, trend, breakdown. Staff rosters moved to another tab.

## HOW THE CONVERSATION WITH AI WENT

YOU

Sketch a one-screen wireframe for these four metrics, and tell me what you would remove.

AI, A TYPICAL REPLY

Headline sentence on top, four tiles, the trend on the left and regions on the right. Remove the staff-on-shift tile: none of your questions uses it.

YOU

Agreed. Move it to an operations tab rather than deleting it.

**AI layouts love to fill every gap. Ask what to take away, not what to add.**

The ideas behind this: [The inverted pyramid](#) · [The visual information-seeking mantra](#) · [Gestalt principles](#) · [The single-screen dashboard](#)

## Your turn


for your own dashboard

☐ Inverted pyramid: headline, KPIs, trends, breakdowns, detail

☐ Most important number top left

☐ Core view on one screen

☐ Group related items; same metric, same place



THE DESIGN LOOP · STAGE 4 OF 6

# Illustrate


Which chart shows each relationship most clearly?

Choose each chart by the relationship it shows. Nine months of change is a line; six regions compared is a sorted bar. Grey everything, then colour only what the reader must see.

Titles do the talking. "On time by region" names a topic; "Hills resolves only 51% on time" states a finding. Every Sahayak chart got a finding for a title.

Sahayak, so far

real charts, one colour



**What changed:** A line for change, a sorted bar for ranking, and red only on Hills.

## HOW THE CONVERSATION WITH AI WENT

YOU

Recommend a chart for each metric using the FT Visual Vocabulary, plus one chart to avoid.

AI, A TYPICAL REPLY

Trend: a line with the 80% target marked. Regions: horizontal bars, sorted, Hills highlighted. Avoid a pie for regions; six slices are hard to compare.

YOU

Now write a finding title for each, under twelve words.

*Position and length are read most accurately; angle and area much less. That is why bars beat pies.*

The ideas behind this: [The Visual Vocabulary](#) · [Ranking of visual encodings](#) · [Data-ink ratio](#) · [Pre-attentive attributes](#)

## Your turn

for your own dashboard

- ☐ Choose charts by relationship
- ☐ Grey everything, colour only what matters

- ☐ Bars start at zero; label directly
- ☐ Titles state the takeaway



THE DESIGN LOOP · STAGE 5 OF 6

# Guard


Can everyone read it, and trust it?

Guard two things: whether everyone can read it, and whether anyone can trust it. Text needs 4.5:1 contrast and chart marks 3:1; colour never carries meaning alone; every number shows its source and date.

Sahayak added a source line, definitions on hover, labels beside every coloured bar, and one honest sentence: drafted with AI, checked by the operations team. One sub-office with six cases was merged, because a group that small can identify people.

Sahayak, so far

safe to share



**What changed:** A source line, the update time and an honest note about AI help.

## HOW THE CONVERSATION WITH AI WENT

YOU

Audit this dashboard for contrast, colour-only meaning, missing sources, and any group small enough to identify a person.

AI, A TYPICAL REPLY

Bar colours pass 3:1. The red and green status dots fail for red-green colour blindness; add text labels. The Hills sub-office has 6 cases; suppress or merge it.

YOU

Now recompute the three headline numbers from the raw file and show your working.

**AI cannot see your screen's colours reliably. Use a contrast checker for the final call, and recompute numbers yourself.**


The ideas behind this: WCAG 2.2 contrast and colour · Colour vision deficiency and the Okabe-Ito palette · Small-cell suppression · Automation bias · Human in the loop

## Your turn

for your own dashboard

- ☐ Contrast 4.5:1 for text, 3:1 for marks
- ☐ Show source, freshness and definitions

- ☐ Never encode meaning in colour alone
- ☐ De-identify; suppress small groups



THE DESIGN LOOP · STAGE 6 OF 6

## Nurture


Is it used, and getting better?

A dashboard is finished when people use it, not when it ships. Sahayak watched five managers try three tasks each. Two couldn't find where the Hills figure came from, so definitions moved from a footnote onto the tiles.

Then they set a rhythm: track opens and filters, review every quarter, retire what nobody reads, and keep a short change log. The sticky notes on the dashboard are real feedback, turned into version 1.2.

Sahayak, so far

in use, getting better



**What changed:** Feedback from five readers, pinned on as sticky notes, became version 1.2.

### HOW THE CONVERSATION WITH AI WENT

YOU

Play five readers: a director, a district manager, an analyst, a funder and a new staff member. Each tries our four questions. Where does each get stuck?

AI, A TYPICAL REPLY

The funder doesn't know what first-call fix means. The new staff member can't tell whether 71% is good; there is no target line.

YOU

Turn those into three changes for version 1.2, smallest first.

**Synthetic readers find the obvious problems fast.**  
**Real readers find the ones that matter.**

The ideas behind this: [Testing with five users](#) · [The five-second test](#) · [Few's thirteen common mistakes](#)

### Your turn

for your own dashboard


- ☐ Test with five readers on three tasks
- ☐ Review quarterly; retire what nobody uses

- ☐ Track views, filters and questions
- ☐ Version it and log changes

## 2 THE RESULT

# Sahayak, version 1.2.

Every element on this dashboard can be traced to one stage of the loop.



**D Decide**  
The purpose sentence, written before anything else

**E Establish**  
Four numbers, each with a written definition

**S Structure**  
Headline, tiles, trend, breakdown: an inverted pyramid

**I Illustrate**  
A line for change, sorted bars for ranking, one accent colour

**G Guard**  
Source, freshness, labels, small groups merged, AI disclosed

**N Nurture**  
Feedback from five readers became version 1.2

*Staff rosters and call durations didn't disappear. They moved to an operations tab nobody is forced to read.*

## 3 THE FIELD GUIDE

# The loop at a glance.

Sahayak showed the loop in motion. This is the whole method on one page, to keep beside you while you build your own. Go round it more than once.

## D Decide

Who decides what, how often?

- Write: this helps [who] decide [what] every [cadence]
- Pick the kind: strategic, operational, analytical or communicative
- List three to five questions in the readers' words
- Interview two real readers first

*Stands on Patton's utilization-focused evaluation, Munzner's domain level, jobs to be done.*

## E Establish metrics

Which numbers answer those questions, exactly?

- Derive metrics from your theory of change, logframe or OKRs
- One card per metric: definition, formula, source, owner
- Pair each lagging outcome with a leading indicator
- Plan equity breakdowns before collecting data

*Stands on SMART indicators, results chains, Few's warning about deficient measures.*

## S Structure the layout

What does the reader see first, second, third?

- Inverted pyramid: headline, KPIs, trends, breakdowns, detail
- Core view on one screen
- Most important number top left
- Group related items; same metric, same place

*Stands on Shneiderman's mantra, Few's single screen, Gestalt, Z and F reading.*

## I Illustrate

Which chart shows each relationship most clearly?

- Choose charts by relationship
- Bars start at zero; label directly
- Grey everything, colour only what matters
- Titles state the takeaway

*Stands on The FT Visual Vocabulary, Tufte's data-ink, Knaflic, pre-attentive attributes.*

## G Guard

Can everyone read it, and trust it?

- Contrast 4.5:1 for text, 3:1 for marks
- Never encode meaning in colour alone
- Show source, freshness and definitions
- De-identify; suppress small groups

*Stands on WCAG 2.2, the Okabe and Ito palette, do-no-harm data ethics.*

## N Nurture

Is it used, and getting better?

- Test with five readers on three tasks
- Track views, filters and questions
- Review quarterly; retire what nobody uses
- Version it and log changes

*Stands on Usability testing, Few's thirteen mistakes, Munzner's validation.*

## 3 THE FIELD GUIDE

# Pick the kind. Define every number.

|        | Strategic                            | Operational                        | Analytical                             | Communicative                       |
|--------|--------------------------------------|------------------------------------|----------------------------------------|-------------------------------------|
| READER | Leadership, boards, funders          | Managers and field teams           | Analysts, MEL and research teams       | Public, partners, communities       |
| RHYTHM | Monthly or quarterly                 | Daily or weekly                    | When a question comes up               | Read once                           |
| SHOWS  | Four to six outcomes against targets | What needs action now, and where   | Comparisons, drivers, breakdowns       | A few findings, explained           |
| DESIGN | Sparse, big numbers, one takeaway    | Alerts, thresholds, sortable lists | Filters, drill-downs, detail on demand | Narrative, annotations, plain words |

*Mixing these on one page is the most common design mistake of all. If you need two, build two pages.*  
*Sahayak is operational: district managers act on it every Monday.*

## METRIC CARD, SAHAYAK EXAMPLE

### On-time service rate

|            |                                                          |
|------------|----------------------------------------------------------|
| QUESTION   | Are we serving people within the promised time?          |
| DEFINITION | Share of requests closed within 7 days of receipt        |
| FORMULA    | $\text{closed\_within\_7d} \div \text{requests\_closed}$ |
| SOURCE     | Case system export, refreshed weekly                     |
| TYPE       | Lagging outcome; its leading indicator is backlog size   |
| TARGET     | 85% by December, set with the operations lead            |
| BREAKDOWNS | Region, request type, gender of applicant                |
| OWNER      | Operations lead; acts when below 75% two weeks running   |

- 1 Start from the results chain.**  
Metrics come from your theory of change, logframe or OKRs, not from what the database holds.
- 2 Pair leading with lagging.**  
Outcomes tell you if you won; leading indicators tell you in time to change course.
- 3 Every metric needs an action.**  
If nobody would act differently when it moves, it belongs in an appendix.
- 4 Plan breakdowns up front.**  
Decide equity cuts now so the data is collected to support them.
- 5 Context, always.**  
Show each number against a target, the last period or a comparison group.

*Hand your AI the theory of change and the data dictionary, and ask for six of these cards. Then ask which ones it could not compute from your data; that list is the real finding.*

## 3 THE FIELD GUIDE

# Answer first, details on demand.

Readers give a dashboard about five seconds to prove it is worth their time. Two classic patterns win those seconds.


## THE INVERTED PYRAMID



## SHNEIDERMAN'S MANTRA, 1996

- 1 Overview first**  
The whole picture on one screen
- 2 Zoom and filter**  
Readers narrow to what they care about
- 3 Details on demand**  
Tooltips, drill-downs and tables when asked

## A ONE-SCREEN WIREFRAME

**One screen**

The core view fits without scrolling (Few)

**Top left matters**

Readers scan in a Z or F; put the key number there

**Group by meaning**

Proximity and alignment show what belongs together


**Same metric, same place**

Consistency lets readers compare at a glance

## 4 THE WORKSHOP

# Spot the flaws.

This chart went to a board meeting. Before reading the answers, find the five classic mistakes yourself.



**1 The axis starts at 60%**

A six-point gap looks like a threefold difference. Bars start at zero.

**3 A legend far from the data**

Readers bounce between key and bars. Label each bar directly.

**5 A topic title, no source or date**

Say the finding, and show where the data comes from and when.

**2 Six colours for one measure**

Colour should mean something. Use one, and highlight only what matters.

**4 A pie with seven slices**

Angles are hard to compare past two or three slices. Use a sorted bar.

*Ask your AI to play critic on your own charts: "Find every way this chart could mislead a hurried reader."*

## 4 THE WORKSHOP

# Chart from question.

Finish the sentence "I want my reader to see...", and the chart follows. Built on the Financial Times Visual Vocabulary.




...how something changed over time

## A line chart

**Reach for:** Lines, columns or a slope chart from first to last. Mark the target.

**Steer clear of:** Pie charts and tables, which hide the shape of change.




...who is ahead, and who is behind

## A sorted bar chart

**Reach for:** Horizontal bars sorted from highest to lowest, or a dot plot.

**Steer clear of:** Bars in alphabetical order, which make the reader do the sorting.




...how big things are, side by side

## A column chart

**Reach for:** Columns or bars starting at zero, with values on top.

**Steer clear of:** Truncated axes, which exaggerate small gaps.




...what a total is made of

## A stacked bar

**Reach for:** A single 100% bar, a treemap or a waffle chart.

**Steer clear of:** Pies with more than three slices.




...how far each one is from a target

## A diverging bar chart

**Reach for:** Bars either side of zero or the target, with the gap labelled.

**Steer clear of:** Colour alone to show above and below.




...whether two things move together

## A scatter plot

**Reach for:** A scatter with a trend line; label the outliers.

**Steer clear of:** Reading cause into a pattern.




...how values are spread out

## A histogram

**Reach for:** A histogram, box plot or strip plot, with the median marked.

**Steer clear of:** A single average, which hides the spread.




...where something happens

## A tile grid map

**Reach for:** Equal-sized tiles per region, or a choropleth with a ranked bar beside it.

**Steer clear of:** A map when location is not the point.



...how people move between stages

## A funnel or Sankey

**Reach for:** A funnel with the share lost at each step, or a Sankey for many paths.

**Steer clear of:** Too many tiny paths to read.


## 4 THE WORKSHOP

# Declutter, then check the colour.

Seven moves inspired by Tufte's data-ink and Knaflic's "eliminate clutter, focus attention".  
Same data, very different read.


## before

Monthly requests by region, January to June



## after seven moves

South **doubled its requests** in six months, overtaking every other region



- 1 Remove border and background
- 2 Remove gridlines
- 3 Drop markers and labels
- 4 Direct labels, no legend
- 5 Push every line to grey
- 6 Colour the line that matters
- 7 Title states the finding

**Pre-attentive attributes.** The eye spots colour, size, position, length and enclosure before reading a word. Use one of them, once, on the thing that matters.

**Graphical integrity.** Bars start at zero, scales stay consistent across small multiples, and areas are never used for one-dimensional values.

About 1 in 12 men and 1 in 200 women have some form of colour vision deficiency. Here is a status palette as most people see it, and as someone with deuteranopia might.

### risky: red and green

Typical vision

Deuteranopia (simulated)

On track and off track collapse into similar tones.

### safer: Okabe and Ito

Typical vision

Deuteranopia (simulated)

Blue stays distinct; orange and vermilion drift closer, so every bar also carries a text label.

## 4.5:1

Text contrast  
WCAG 2 · 1.4.3 · AA

## 3:1

Chart marks and key UI  
WCAG 2 · 1.4.11 · AA

## ≠ colour

Never colour alone  
WCAG 2 · 1.4.1 · A

## Alt

A text summary per chart  
WCAG 2 · 1.1.1 · A

## 4 THE WORKSHOP

# The critic.

Stephen Few's thirteen common dashboard mistakes, as a self-check. Tick what your dashboard already avoids.

- |                                                                    |                                                                    |
|--------------------------------------------------------------------|--------------------------------------------------------------------|
| 1 <input type="checkbox"/> Fits on a single screen                 | 2 <input type="checkbox"/> Gives enough context for every number   |
| 3 <input type="checkbox"/> Avoids excessive detail or precision    | 4 <input type="checkbox"/> Uses measures that answer the question  |
| 5 <input type="checkbox"/> Uses the right chart for each job       | 6 <input type="checkbox"/> Avoids variety for its own sake         |
| 7 <input type="checkbox"/> Uses well-designed charts, not defaults | 8 <input type="checkbox"/> Encodes quantities accurately           |
| 9 <input type="checkbox"/> Arranges data in a meaningful order     | 10 <input type="checkbox"/> Highlights what matters most           |
| 11 <input type="checkbox"/> Has no useless decoration              | 12 <input type="checkbox"/> Uses colour sparingly and with meaning |
| 13 <input type="checkbox"/> Is visually clean and attractive       |                                                                    |

Score: \_\_\_\_ / 13

## then hand the gaps to your AI

You are a dashboard design critic who follows Stephen Few's principles. Review the attached dashboard. For each item I did not tick, explain what is wrong, why it matters to the reader, and the exact fix. Then rank the fixes by their effect on the reader's decision.

## A twenty-minute test with five readers

### Recruit five

Real readers, not the team that built it.  
Small, repeated rounds find most problems.

### Give three tasks

- Find this month's headline figure
- Which region is furthest behind?
- Where does this number come from?

### Watch, don't help

Note where they hesitate or misread. Ask afterwards what they expected to see.

5 WORKING WITH AI

# The AI pact.

AI can draft almost everything on a dashboard. It cannot be accountable for any of it. Eight jobs it does well, what you keep in each, and six promises that hold whatever the tool.

## Eight roles AI can play. One it can't.

**1**  
**Interviewer**  
Surfaces purpose, readers and decisions.  
You keep: confirming with real users

**2**  
**Metric architect**  
Turns questions into metric cards.  
You keep: signing off definitions

**3**  
**Layout sketcher**  
Drafts wireframes and reading order.  
You keep: choosing what to cut

**4**  
**Chart chooser**  
Matches relationships to charts.  
You keep: the final visual call

**5**  
**Builder**  
Writes DAX, SQL, formulas or chart code.  
You keep: recomputing the headlines

**6**  
**Critic**  
Reviews against Few, Tufte and Knaflc.  
You keep: weighing trade-offs

**7**  
**Accessibility checker**  
Flags contrast, colour-only meaning and jargon.  
You keep: testing with a contrast tool

**8**  
**Narrator**  
Writes takeaways, summaries and alt text.  
You keep: checking every claim

The one it can't play: the person who answers for the numbers.

## Six promises to keep

**1**  
**Every number traces to a row**  
Nothing you can't follow to source data and a formula.

**2**  
**Headlines get recomputed**  
A person checks the top three figures first.

**3**  
**No personal data goes in**  
Aggregate and de-identify, unless your organisation approved the tool.

**4**  
**Definitions come first**  
AI fills a gap with a plausible definition. Give it yours.

**5**  
**A person signs off**  
A named owner for numbers, wording and release.

**6**  
**AI's help is disclosed**  
One line: what AI drafted, what people checked.

## Where AI lives in your tools

| Tool                  | AI help for dashboards                                                                                       | Best for                        |
|-----------------------|--------------------------------------------------------------------------------------------------------------|---------------------------------|
| General assistants    | Claude, ChatGPT, Gemini and Copilot draft metric cards, layouts, chart code and critiques, and analyse files | Design thinking, critique, code |
| Power BI with Copilot | Generates report pages, suggests DAX measures, writes narrative summaries                                    | Microsoft 365 teams             |
| Tableau Pulse         | Personalised metric summaries, proactive insights, alerts by email or Slack                                  | Metric monitoring               |
| Looker with Gemini    | Conversational analytics, a visualisation assistant, calculated fields                                       | Google Cloud teams              |
| ThoughtSpot Spotter   | Conversational analysis and AI-assisted dashboard creation                                                   | Search-style self-service       |
| Metabase Metabot      | SQL generation, chart summaries, query debugging                                                             | Open-source, lean teams         |
| Spreadsheets          | Copilot in Excel and Gemini in Sheets suggest formulas, pivots and charts                                    | Small teams, trackers           |

Features change fast; check your tool's current release notes. The DESIGN loop works with any of them.

## 6 BORROW THESE

# Conversation starters.

Eighteen openings for any AI assistant, three for each stage. Replace the highlighted words, then push back on the first answer; the second is usually better.

## D Decide

### Find the purpose

Interview me, one question at a time, about who uses this dashboard, what they decide and how often. Then write a one-sentence purpose and five key questions.

*Then push back: "What am I assuming about these readers that I haven't checked?"*

### Choose the type

Given this purpose: [purpose], should this be a strategic, operational, analytical or public-facing dashboard? Explain the trade-offs.

*Then push back: "What would I lose if I chose the second-best option?"*

### Map the readers

List the readers of a dashboard on [topic], what each one decides, and which questions they share.

*Then push back: "Which reader would be hurt most by a wrong number?"*

## E Establish

### Write the specs

Turn these questions into metric specs with definition, formula, source, leading or lagging, target and owner: [questions]. Our columns are: [columns].

*Then push back: "Which definition would two teams most likely compute differently?"*

### From theory of change

From this theory of change: [paste], propose output, outcome and leading indicators, and flag gaps in our data.

*Then push back: "Which of these could be gamed once it becomes a target?"*

### Plan equity cuts

Which breakdowns of [metric] would show who is being left behind? Suppress any group under [10] people.

*Then push back: "Who should we ask before publishing these breakdowns?"*

## S Structure

### Sketch the screen

Draft a one-screen wireframe: a headline sentence, up to six tiles, two trends and one breakdown. Explain the reading order.

*Then push back: "What would you remove if I allowed only four elements?"*

### Cut it down

Here is my plan: [list]. Remove everything that doesn't answer the five key questions, and say why each cut is safe.

*Then push back: "Where should the removed items live instead?"*

### Split by audience

Split this into pages for [audience A] and [audience B], keeping shared metrics in the same place on both.

*Then push back: "What should both pages always show identically?"*

## I Illustrate

### Pick the charts

For each metric, recommend a chart using the FT Visual Vocabulary, one alternative and one to avoid: [metrics].

*Then push back: "Which of these could mislead a hurried reader?"*

### Write finding titles

Rewrite these chart titles as one-sentence findings under twelve words: [titles and key numbers].

*Then push back: "Is any title claiming more than the data shows?"*

### Declutter

Review this chart using Tufte and Knaflitz: list what to remove, what to grey out and what to highlight.

*Then push back: "Show me the before and after as a checklist."*

## G Guard

### Accessibility audit

Audit this dashboard for WCAG 1.4.3 and 1.4.11 contrast, colour-only meaning, missing text alternatives and jargon. List each issue with a fix.

*Then push back: "Which issue affects the most readers?"*

### Check the numbers

Recompute these headline figures from the attached data and show your working: [figures]. Flag any mismatch.

*Then push back: "What assumptions did you make in the calculation?"*

### Privacy scan

Scan this dataset and dashboard for anything that could identify a person: names, IDs, small groups, rare combinations.

*Then push back: "What threshold would you suggest for group size, and why?"*

## N Nurture

### Synthetic readers

Play five readers: [roles]. Each tries our key questions. Report where each one gets stuck.

*Then push back: "Rank the problems by how often a real reader would hit them."*

### Usability script

Write a twenty-minute usability test for this dashboard: three tasks, what to watch for, and five follow-up questions.

*Then push back: "What should I not say during the test?"*

### Quarterly review

Here is usage data: [opens, filters, comments]. Recommend what to keep, change and retire.

*Then push back: "Draft a short change log for readers."*

## 7 BEFORE YOU LAUNCH

# DESIGN-ready?

Twelve checks across the loop. Print this page and tick it off for the dashboard you are building.

|                                                                             |                                                                             |
|-----------------------------------------------------------------------------|-----------------------------------------------------------------------------|
| <input type="checkbox"/> Purpose confirmed with a real reader <b>D</b>      | <input type="checkbox"/> Dashboard kind chosen <b>D</b>                     |
| <input type="checkbox"/> Every metric has a card and an owner <b>E</b>      | <input type="checkbox"/> Every number has context <b>E</b>                  |
| <input type="checkbox"/> Core view fits on one screen <b>S</b>              | <input type="checkbox"/> Takeaway sits top left <b>S</b>                    |
| <input type="checkbox"/> Charts chosen by relationship <b>I</b>             | <input type="checkbox"/> One accent colour, the rest grey <b>I</b>          |
| <input type="checkbox"/> Contrast and colour checks passed <b>G</b>         | <input type="checkbox"/> Headlines recomputed by a person <b>G</b>          |
| <input type="checkbox"/> No personal data; small groups suppressed <b>G</b> | <input type="checkbox"/> Tested with five readers; review date set <b>N</b> |

**Score:** \_\_\_\_ / 12

## After launch, keep it alive.

### Track use

Views, filters used, comments and questions

### Review quarterly

Keep, change or retire every element

### Version it

Log what changed and why

### Close the loop

Tell readers what you changed from their feedback

*A dashboard is finished when people use it, not when it ships.*

## 7 KEEP THIS ONE

# The playbook on a page.

Pin it above your desk, or share it with the team before your next dashboard.

## D • Decide

Write who decides what, how often. Interview two real readers.

**Remember:** Most dashboards fail here, before the first chart.

## E • Establish metrics

A card per metric: definition, formula, source, owner, target.

**Remember:** Pair every outcome with a leading indicator.

## S • Structure the layout

Headline, tiles, trend, breakdown, detail. One screen.

**Remember:** Most important number top left.

## I • Illustrate

Chart by relationship. Grey first, then one accent colour.

**Remember:** Titles state the finding.

## G • Guard

Contrast 4.5:1 and 3:1, never colour alone, source and date shown.

**Remember:** Recompute headlines; suppress small groups.

## N • Nurture

Five readers, three tasks. Review quarterly; retire what nobody uses.

**Remember:** Version it and tell readers what changed.

### Five-second test

- Show it to a colleague for five seconds
- Ask for the main point
- If they miss it, fix the hierarchy

### Hand to AI

- Interviews, specs, wireframes
- Chart code, critiques, audits
- Drafts of titles and alt text

### Keep for yourself

- Talking to real readers
- Signing off every definition
- Recomputing the headline numbers

## Six promises for AI-built dashboards

### Every number traces to a row

Nothing you can't follow to source data and a formula.

### Headlines get recomputed

A person checks the top three figures first.

### No personal data goes in

Aggregate and de-identify, unless your organisation approved the tool.

### Definitions come first

AI fills a gap with a plausible definition. Give it yours.

### A person signs off

A named owner for numbers, wording and release.

### AI's help is disclosed

One line: what AI drafted, what people checked.

## 8 GO DEEPER

# The library of ideas.

The concepts, theories and practices this playbook stands on: where each comes from, and how to use it.

## ● Pre-attentive attributes

Perception · Anne Treisman, 1980; Christopher Healey

Colour, size, position, orientation and enclosure "pop out" from their surroundings without effort. Treisman's feature integration theory explains why one red bar among grey ones is found at once, while a red bar among many colours is not.

**On Monday:** Use one pre-attentive cue, once, on the single thing the reader must see.

## ● Gestalt principles

Perception · Wertheimer, Koffka and Köhler, 1920s

Proximity, similarity, enclosure, connection and continuity make the brain group things. Items placed close together are read as related, even without a box around them.

**On Monday:** Group related tiles with white space and alignment rather than borders and backgrounds.

## ● Ranking of visual encodings

Perception · William Cleveland and Robert McGill, 1984

Their experiments found people judge position on a common scale most accurately, then length, then angle and slope, then area, with volume and colour intensity least accurate.

**On Monday:** Prefer bars and dots to pies, bubbles and colour-shaded maps when precise comparison matters.

## ● Cognitive load

Perception · John Sweller, 1988

Cognitive load theory separates the effort the task needs from extra effort caused by poor presentation. Legends, unclear labels and decoration add effort without adding meaning.

**On Monday:** Remove anything that makes the reader work without telling them something.

## ● Data-ink ratio

Design · Edward Tufte, 1983

In *The Visual Display of Quantitative Information*, Tufte argues for erasing non-data ink and redundant data ink, within reason. Gridlines, borders and backgrounds are usually the first to go.

**On Monday:** Run the seven declutter moves in the workshop on your own charts.

## ● Chartjunk

Design · Edward Tufte, 1983

Heavy grids, 3D effects, shadows, gradients and clip art. Later research shows some embellishment can help memory, but it rarely helps accurate reading.

**On Monday:** If an element would not be missed when deleted, delete it.

## ● Graphical integrity and the lie factor

Design · Edward Tufte, 1983

The lie factor compares the size of an effect shown in the graphic with its size in the data. Truncated bar axes and area used for one-dimensional values inflate differences.

**On Monday:** Start bars at zero, keep scales consistent and label any break clearly.

## ● Small multiples

Design · Edward Tufte

A grid of identical charts on one shared scale lets the eye compare shape and level across regions or groups at a glance, without a crowded single chart.

**On Monday:** Use small multiples instead of one chart with eight coloured lines.

## ● The single-screen dashboard

Design · Stephen Few, 2006

In *Information Dashboard Design*, Few defines a dashboard as a display of the information needed to meet one or more objectives, arranged on a single screen so it can be monitored at a glance.

**On Monday:** If the core view needs scrolling, move detail to a second page.

## ● Few's thirteen common mistakes

Design · Stephen Few, 2006

They range from exceeding one screen and giving too little context to choosing the wrong measures, misusing colour and adding useless decoration.

**On Monday:** Score yourself with the critic in the workshop.

## ● The bullet graph

Design · Stephen Few, about 2005

One bar shows the actual value, a tick marks the target, and shaded bands show poor, fair and good ranges, all in a small space that lines up neatly with others.

**On Monday:** Use bullet graphs for KPIs with targets instead of dials.

## ● The visual information-seeking mantra

Design · Ben Shneiderman, 1996

From his paper *The Eyes Have It*, the mantra describes how people explore information: see the whole, narrow down, then inspect one item.

**On Monday:** Design the first screen as the overview; put filters and drill-downs behind it.

## ● The inverted pyramid

Design · Journalism practice

News writers lead with the conclusion, then the supporting facts, then background. On a dashboard: headline, key numbers, trends, breakdowns, detail.

**On Monday:** Write the headline sentence before you place any chart.

## ● The Visual Vocabulary

Design · Financial Times

The FT sorts charts into nine relationships: deviation, correlation, ranking, distribution, change over time, magnitude, part-to-whole, spatial and flow.

**On Monday:** Name the relationship first; the chart choice follows.

## ● The nested model

Design · Tamara Munzner, 2009

Domain problem, data and task abstraction, visual encoding and interaction, and algorithm. A beautiful chart built on the wrong problem still fails, so validate each level.

**On Monday:** Check you have the right problem and questions before polishing visuals.

## ● Storytelling with data

Design · Cole Nussbaumer Knaflic, 2015

Understand the context, choose an appropriate visual, eliminate clutter, focus attention where you want it, think like a designer, and tell a story.

**On Monday:** Write every chart title as the finding, not the topic.

## ● Utilization-focused evaluation

Decisions and use · Michael Quinn Patton, first edition 1978

Patton's approach starts by identifying the primary intended users and their intended uses, and designs everything around them. It applies directly to dashboards.

**On Monday:** Name the primary readers and the decisions they make before choosing any metric.

## ● Jobs to be done

Decisions and use · Clayton Christensen and colleagues

Rather than asking what features users want, ask what progress they are trying to make. A manager hires a dashboard to decide where to send staff, not to look at charts.

**On Monday:** Ask readers: what were you trying to decide the last time you opened it?

## ● Theory of change and results chains

Measurement · Development and evaluation practice

Inputs lead to activities, outputs, outcomes and impact, with assumptions at each step. Logframes and OKRs play a similar role in other sectors.

**On Monday:** Pull dashboard metrics from your results chain, not from whatever the database holds.

## ● SMART indicators

Measurement · George Doran, 1981

Specific, measurable, assignable or achievable, realistic and time-related. Many variants exist; the point is precision.

**On Monday:** Run each metric spec card through the SMART test.

## ● Leading and lagging indicators

Measurement · Management practice

Lagging indicators, such as on-time resolution, confirm outcomes. Leading indicators, such as backlog size, move first and give time to act.

**On Monday:** Pair every outcome metric with one leading indicator.

## ● Goodhart's law and Campbell's law

Measurement · Charles Goodhart, 1975; Donald Campbell, 1970s

Often paraphrased, after Marilyn Strathern, as "when a measure becomes a target, it ceases to be a good measure". Campbell warned that the more a quantitative indicator drives decisions, the more it invites corruption.

**On Monday:** Track a counter-metric, such as reopened requests next to closure rate.

## ● WCAG 2.2 contrast and colour

Access and ethics · W3C, 2023

Text needs a contrast ratio of at least 4.5:1 (1.4.3), chart marks and key interface parts 3:1 (1.4.11), and colour must never be the only way meaning is shown (1.4.1).

**On Monday:** Test your palette with a contrast checker before launch.

## ● Colour vision deficiency and the Okabe-Ito palette

Access and ethics · Masataka Okabe and Kei Ito, 2008

Around 1 in 12 men and 1 in 200 women have some colour vision deficiency, most often red-green. Okabe and Ito proposed eight colours chosen to remain distinct.

**On Monday:** Avoid red-green pairs for status; add text labels whatever the palette.

## ● Small-cell suppression

Access and ethics · Statistical disclosure control

Statistics offices suppress or merge counts below a threshold, often 5 or 10, because a tiny group combined with other facts can reveal who someone is.

**On Monday:** Set a minimum group size for every breakdown on the dashboard.

## ● Data feminism

Access and ethics · Catherine D'Ignazio and Lauren Klein, 2020

The book applies feminist thinking to data: examine power, challenge binaries, make labour visible and consider context. It is a useful lens for which breakdowns a dashboard shows.

**On Monday:** Decide equity breakdowns with the people the data is about.

## ● Human in the loop

AI · AI governance practice

Rather than full automation, a person checks key steps: definitions, numbers and release. Responsibility stays with a named human.

**On Monday:** Name an owner for every AI-assisted dashboard.

## ● Automation bias

AI · Raja Parasuraman and Victor Riley, 1997

Research on humans and automation shows people accept automated suggestions even when they are wrong, especially when busy. Fluent AI answers make this more likely.

**On Monday:** Recompute AI-produced headline numbers yourself, every time.

## ● Errors in AI-generated charts

AI · Benchmarks such as VisEval, 2024

Benchmarks testing language models on chart generation find failures ranging from invalid code to wrong aggregations and misleading encodings.

**On Monday:** Treat AI charts as drafts until the numbers are checked against the source.

## ● Testing with five users

Testing · Jakob Nielsen and Thomas Landauer, 1993

Their model suggested about five users uncover most problems in a design, so several small rounds beat one big study.

**On Monday:** Test with five readers, fix, and test again.

## ● The five-second test

Testing · User research method

A quick way to check whether the main message lands at first glance. If readers can't recall the key point after five seconds, the hierarchy needs work.

**On Monday:** Run it with a colleague before every launch, as you did at the top of this page.



written by

# Shashidhar S A

Founder, HubForge Global · MERL & AI Innovator

- En-Intrapreneur
- Driving evaluation ecosystems through People, Process & Product
- Co-Chair, IOCE Comms & Storytelling
- Co-Lead, EvalYouth Global Champion Engagement

LinkedIn: [linkedin.com/in/shashidharsa](https://www.linkedin.com/in/shashidharsa)

## Also from HubForge

**Claude Dashboard Playbook:** building dashboards hands-on in Claude, from the first prompt to live data, filters, comments and sharing.

## On the shelf

The books, papers and standards behind the ideas, if you want the originals.

### Information Dashboard Design

Stephen Few

### The Visual Display of Quantitative Information

Edward Tufte

### Storytelling with Data

Cole Nussbaumer Knaflic

### The Big Book of Dashboards

Wexler, Shaffer and Cotgreave

### A Nested Model for Visualization Design

Tamara Munzner, 2009

### The Eyes Have It

Ben Shneiderman, 1996

### Visual Vocabulary

Financial Times

### Utilization-Focused Evaluation

Michael Quinn Patton

### WCAG 2.2

W3C

### Colour-blind safe palette

Masataka Okabe and Kei Ito

### VisEval benchmark

Evaluating AI-generated charts

### Claude Dashboard Playbook

HubForge



## Designing Dashboards with AI Playbook

Made by HubForge Global Impact Network Foundation, which builds evidence systems and ecosystems for the MERL sector. Share it freely with credit.

write to [hubforgeglobal@gmail.com](mailto:hubforgeglobal@gmail.com)

visit [www.hubforgeglobal.com](http://www.hubforgeglobal.com)