




BUSINESS AND CUSTOMER ANALYTICS

Design Dashboards Around Real Decisions

Attach metrics to owners, thresholds, freshness and useful next actions



Sanjaay Gowda T S

FREE SAMPLE / INTRODUCTION + CHAPTER 1 / EDITION 1.0.0

sanjaaygowdata.ai

About this free sample

Design Dashboards Around Real Decisions

Sanjaay Gowda T S

First edition 1.0.0, 2026-10-13

Copyright 2026 Sanjaay Gowda T S. This is a labelled free sample, not the complete ebook. You may share this sample unchanged. It contains the cover, sample permissions, contents, introduction, Chapter 1 and a closing product link. The remaining chapters, worked solutions and companion files are not included. The full ebook text, diagrams and artwork have separate single-reader terms; this sample permission does not authorize sharing the full purchased PDF or EPUB.

The companion source code, tests and fictional data are separately MIT licensed. You may adapt or share that project with its included licence notice. MIT does not apply to the ebook text or book artwork. The books and companion projects have no technical activation requirement or buyer-specific watermark.

This first edition was prepared with AI assistance. The case studies are fictional teaching examples. Any measured outputs come from the supplied examples, not production customer data. The teaching scope, assumptions and limitations are explained in the chapters. Independent human technical, editorial and assistive-technology reader review is pending. Automated tests and accessibility checks do not establish certification or suitability for every use.

Run examples in a separate working folder using the supplied synthetic data. Review inputs and calculations before adapting the techniques to real operations. Model software, service access, subscriptions, ongoing updates and consulting are not included.

Author website: sanjaaygowdats.ai

Ebook terms: sanjaaygowdats.ai/ebooks/license

Inside this free sample

Start here.....	4
1. Start with the decision and owner	5
Continue with the full guide	7

This free sample contains the introduction and Chapter 1 only. The closing link opens the complete guide. The companion project and worked solutions are supplied with the complete ebook.

Start here

A useful dashboard helps a named person make a recurring decision. A collection of attractive charts can fail that test while showing perfectly correct arithmetic. The missing piece is often a decision contract: who reviews the evidence, when they review it, what action is available, and what uncertainty should stop that action. This guide starts with that contract and works toward the display.

The case study is Birch Support, a fictional team handling requests for a fictional equipment-rental service. Its conversations, staffing arrangements, thresholds and results are invented teaching data. No example represents a real customer's performance. Birch wants to notice a growing queue of late replies without treating every noisy percentage as an emergency. It also wants to show clearly when an apparently reassuring result is too old to support a decision.

You will define three metric cards and use a small Python rule evaluator to classify their evidence. The companion takes a local JSON file, calculates ratios and data age, and returns inspectable states. It does not draw a dashboard, connect to a business intelligence platform, send alerts or change a customer record. The design work remains portable to a spreadsheet, a reporting tool or a custom application.

Basic arithmetic and experience reading an operational report are sufficient for the design exercises. The optional companion uses Python's standard library and syntax compatible with Python 3.10 or later; use a supported Python release on your machine. Extract the project and run these commands from its folder:

```
python main.py --input example-input.json
python -m unittest discover -s tests -v
```

The example deliberately contains three different problems. Late replies breach a supplied rule. A completion percentage has too few observations for Birch's chosen guard. A reopened-case percentage is stale. These cases teach why a dashboard should distinguish a concerning measurement from insufficient evidence. The rules are policy choices for this lesson, not statistical anomaly detection or universal service targets.

Keep the original fixture while experimenting. Work through the nine chapters, record the decision each card supports, and compare your answers with the solutions. You should finish with a compact dashboard specification, a set of scenarios for reviewing it, and a clear explanation of what its green, review and unknown states actually mean.

1. Start with the decision and owner

Describe a meeting before drawing a screen

Birch's support lead reviews the unassigned queue at noon. The available decision is whether to investigate older conversations and temporarily redistribute work. That is more specific than asking for a customer-experience dashboard. It identifies a person, a time, an object of attention and a reversible action. It also makes clear that the dashboard is not deciding whether an individual employee has performed well.

Write a short decision statement: at noon, the support lead inspects evidence about late first replies and decides whether the queue needs investigation. Then ask what the lead would do if the measurement looked concerning. If the only answer is "watch the number," the design is incomplete. Watching may be appropriate, but give it a purpose and a next review time rather than presenting it as an automatic improvement.

An owner is accountable for the decision, not necessarily for every cause of the result. Birch's lead can redistribute conversations but cannot repair the data warehouse. A separate data owner must investigate a stale extract. Put both responsibilities in the specification. Otherwise an unknown state can bounce between teams while the visible card remains unexplained.

Connect evidence to an available action

The GOV.UK performance-metrics guidance starts from service purpose and connects measurements with goals and decisions [R1]. For this fictional case, turn that principle into a practical worksheet. List the decision, evidence needed, action available, person responsible and conditions that require stopping. The worksheet is your design, not a prescribed government reporting template.

Decision field	Birch example	Review question
Owner	Support lead	Can this person act on the queue?
Cadence	Noon operational review	Is the evidence current enough?
Evidence	Late first replies among eligible conversations	Are numerator and denominator aligned?
Action	Inspect oldest unassigned conversations	Is the action proportionate to uncertainty?
Stop condition	Stale or sparse evidence	Who repairs the evidence?




Figure 1. A decision starts before the chart. Original teaching diagram; all cases and values are fictional.

Limit the card's authority

The late-reply card can nominate a queue for review. It cannot establish why a delay occurred. A missing timestamp, a changed opening-hours policy, a sudden mix of complex requests and a real staffing shortage can produce similar visible results. Investigation should distinguish these possibilities before anyone acts on a causal story.

Record the cost of being wrong. A false review signal may use a few minutes of the lead's time. Automatically cancelling bookings or judging employees would have much greater consequences. The same measurement can be adequate for the first decision and inadequate for the others. Match the evidence to the actual action, rather than declaring a dashboard generally trustworthy.

Exercise: reject a vague requirement

Rewrite "show whether support is good" as two different decision statements. Give each a named role, a review cadence, an available action and a stop condition. Explain why neither statement can be implemented solely by adding more charts to the page.

Continue with the full guide

This free sample ends after the introduction and Chapter 1. The complete guide contains nine teaching chapters, worked exercise solutions, primary references and a separate Python companion project. The full PDF, reflowable EPUB and project files are supplied together.

Use the complete guide to work through the remaining exercises and inspect the supplied examples. The companion code and fictional data are separately MIT licensed; the complete ebook has the single-reader terms described on its product page.

The examples are teaching fixtures, not production customer results. Independent human technical, editorial and assistive-technology reader review remains pending. Software access, subscriptions, consulting and guaranteed results are not included.

[View the complete guide and purchase options.](#)

You may share this labelled free sample unchanged. Keep the full purchased PDF and EPUB private. Thank you for reading.

Primary references

1. [R1] GOV.UK: How to set performance metrics for your service

<https://www.gov.uk/service-manual/measuring-success/how-to-set-performance-metrics-for-your-service>

Using performance data to improve a service and selecting meaningful service metrics.