



Berner Fachhochschule
Haute école spécialisée bernoise
Bern University of Applied Sciences



Regional
Tourism
Alliance



^b
UNIVERSITÄT
BERN

Interaction Quality Dashboard

Dashboard User Manual

A step-by-step guide for destination management organisations

Version	1.0
Date	08.09.2026
Published by	RTA / FFT, University of Bern / BFH-AMI
Authors	Adrian Müller, Dominic Beutler
Platform	gastfreundlichkeit.ch
Support	General inquiries: adrian.mueller@unibe.ch Technical inquiries: mascha.kurpicz@bfh.ch

Inhaltsverzeichnis

1.	About This Manual	6
1.1	Purpose of this manual	6
1.2	Who this manual is for	6
1.3	How to use this manual	6
2.	Getting Started: Access and Login	7
2.1	Requirements	7
2.2	Logging in step by step	7
2.3	Changing your password	8
2.4	Logging out	8
2.5	Access rights: what you can see	8
3.	The Platform at a Glance	9
3.1	Navigation structure	9
3.2	The dashboard list	10
3.3	Tags	10
3.4	Anatomy of a dashboard	11
3.5	Exporting and sharing results	12
4.	Where the Data Comes From	14
4.1	Data source	14
4.2	From review to data point	14
4.3	Data cleaning and coverage	15
4.4	The metrics explained	15
4.5	Hospitableness and hardware: the key distinction	15
4.6	Topics	16
4.7	Aggregation levels	17
4.8	How current the data is	17
4.9	What the data can and cannot tell you	17
5.	Working with Filters	18
5.1	The filter bar	18
5.2	Filtering directly in a chart	18
5.3	Combining filters	19
5.4	Resetting filters	20

u^b

5.5	What to watch out for	21
6.	Dashboard Reference	22
6.1	Data Basis	22
6.1.1	Purpose	22
6.1.2	The visuals explained.....	23
6.1.3	What to watch out for	24
6.2	Points of Interest Map	24
6.2.1	Purpose	24
6.2.2	The visuals explained.....	24
6.2.3	What to watch out for	25
6.3	Metrics	25
6.3.1	Purpose	25
6.3.2	The visuals explained.....	27
6.3.3	What to watch out for	28
6.4	Hospitableness Metrics	29
6.4.1	Purpose	29
6.4.2	The visuals explained.....	30
6.4.3	What to watch out for	33
6.5	Hardware Metrics	34
6.5.1	Purpose	34
6.5.2	The visuals explained.....	35
6.5.3	What to watch out for	37
6.6	Overview	38
6.6.1	Purpose	38
6.6.2	The visuals explained.....	38
6.6.3	What to watch out for	39
6.7	POI Type Benchmarking	39
6.7.1	Purpose	39
6.7.2	The visuals explained.....	40
6.7.3	What to watch out for	41
6.8	Hospitableness Topic Benchmarking.....	41
6.8.1	Purpose	41
6.8.2	The visuals explained.....	42
6.8.3	What to watch out for	43

u^b

6.9	Hardware Topic Benchmarking.....	43
6.9.1	Purpose	43
6.9.2	The visuals explained.....	44
6.9.3	What to watch out for	45
7.	Typical Workflows	46
7.1	Preparing a conversation with a single service provider.....	46
7.1.1	How to get there	46
7.1.2	What you see	47
7.1.3	Reading the four values together	47
7.1.4	Before you use it	48
7.2	Preparing the annual meeting with the hoteliers association	48
7.2.1	How to get there	48
7.2.2	What you see	49
7.2.3	How to read it	50
7.2.4	Before you use it	50
7.3	Positioning the destination for a strategy or board meeting.....	50
7.3.1	How to get there	50
7.3.2	The three levels of comparison.....	51
7.3.3	What you see	51
7.3.4	How to read it	52
7.3.5	Before you use it	52
8.	Interpreting Results Responsibly	53
8.1	How many reviews stand behind a figure	53
8.2	Destinations have different data histories	53
8.3	Star ratings and text analysis measure different things	53
8.4	Not everything guests complain about is measured	53
8.5	Language effects.....	53
8.6	Differences are small, and that is expected	54
8.7	No causal conclusions	54
8.8	Communicating results to businesses	54
8.9	Data protection and confidential figures.....	54
9.	Glossary	55

10.	FAQ and Support	56
10.1	Frequently asked questions	56
10.2	Reporting a problem.....	57
10.3	Contact.....	57

u^b

1. About This Manual

1.1 Purpose of this manual

u^b

This manual explains how to work with the Interaction Quality Dashboard on gastfreundlichkeit.ch. It shows you how to log in, which dashboards exist, how to apply filters, and how to read what the charts show. It also explains, in plain language, where the underlying data comes from and what the figures can and cannot tell you.

The manual is not a methodological paper. It covers as much of the method as you need to use the dashboards correctly and explain the results to others. If you need the technical details behind the models, please refer to the separate methodology note.

1.2 Who this manual is for

This manual is written for people working in a destination management organisation who use the Interaction Quality Dashboard in their daily work: quality management, product development, market research, management. No prior knowledge of data analysis or of natural language processing is assumed.

Each participating destination has its own login. You will see the detailed views for your own destination, and you will see all participating destinations by name in the benchmarking views.

1.3 How to use this manual

- If you are opening the dashboards for the first time, read chapters 2 to 5 in order. They take you from the login to a working understanding of the filters.
- If you already know the platform, go straight to chapter 6, which documents each dashboard separately in the same structure.
- Chapter 7 shows three complete workflows across several dashboards.
- Read chapter 8 before you share results with partners or service providers. It explains the limits of the data and how to communicate findings fairly.

2. Getting Started: Access and Login

2.1 Requirements

- All common browsers are supported.
- There are no specific requirements regarding screen resolution.

Each destination receives one shared login. Accounts are created by BFH; there is no self-registration. To request access, or to have an account changed, contact the support address given in chapter 10.3.

2.2 Logging in step by step

1. Open gastfreundlichkeit.ch in your browser. Select the tab «Dashboard». Click the red button.




Figure 1: Landing page with the login entry point

2. Enter your username and password. Click the «Log in» button.




Figure 2: Login form

3. Click «Analytics», then «Dashboard». You will then see all the available dashboards.




Figure 3: Landing view after successful login

2.3 Changing your password

There is currently no self-service password reset. If you have lost your password or need it changed, send a request to the support address in chapter 10.3.

2.4 Logging out

The log out button can be found in the top-right corner by clicking the coloured circle. Since all DMOs share a single company login account, you should always log out after using the dashboard.

2.5 Access rights: what you can see

Your login provides detailed views of your own destination. In the benchmarking dashboards, you also see all other participating destinations by name, so comparison works in both directions: they can see you as well.

Administrator access is restricted to the project team. If a dashboard appears empty, this is normally a filter setting rather than a permission problem; see chapter 5.

3. The Platform at a Glance

The dashboards run on Kibana, a standard visualisation environment. If you have used Kibana elsewhere, the controls will be familiar.

3.1 Navigation structure

Everything you need sits on one screen: a navigation bar at the top, and the dashboard itself below.

There are three ways to move around.

The dashboard list. This is your starting point after logging in. It shows all dashboards available to you, with a short description for each. Select a dashboard to open it. To return to the list at any time, select **Dashboards** in the breadcrumb at the top left.

The breadcrumb. The bar at the top left always shows where you are, in the form **Dashboards** followed by the name of the dashboard you have opened. The name is also the quickest way to confirm which language version of a topic dashboard you are looking at.

Links inside a dashboard. The benchmarking dashboards carry their own navigation. A list on the left takes you directly to the other benchmarking views without going back to the dashboard list, and the entry you are currently on is highlighted. On the topic benchmarking dashboards, a second row of links switches between the four review languages. These language links open a different dashboard; they are not filters.

At the top right of every dashboard, you will find **Full screen**, which hides the browser interface for presentations, **Share**, and **Reset**, which returns the dashboard to its default state after you have changed filters.




Figure 4: The dashboard list – gives you an overview of all the available dashboards

3.2 The dashboard list

The dashboard list is not ordered by importance. The table below groups the dashboards as presented in this manual, from understanding your data basis through to benchmarking.

Dashboard	Group	Use it to	Chapter
Data Basis	Data basis	Check how many reviews your figures rest on	6.1
Points of Interest Map	Data basis	See which POIs are recorded for your destination	6.2
Metrics	Destination detail	Read Google rating and overall sentiment	6.3
Hospitableness Metrics	Destination detail	Read hospitableness scores and topics	6.4
Hardware Metrics	Destination detail	Read hardware scores and topics	6.5
Overview	Benchmarking	Compare your destination with the others	6.6
POI Type Benchmarking	Benchmarking	Compare POI categories across destinations	6.7
Hospitableness Topic Benchmarking	Topic benchmarking	Compare Hospitableness Topics by review language	6.8
Hardware Topic Benchmarking	Topic benchmarking	Compare hardware topics by review language	6.9

3.3 Tags

The tags shown in the dashboard list describe the type of analysis each dashboard contains: `data_basis` for review counts and distributions, `google_rating` for star ratings, `sentiment` for sentiment across the whole review, and `hospitality_sentiment` and `hospitality_topic_sentiments` for analyses restricted to hospitableness statements. They are an internal ordering system; you do not need them to find your way around.

3.4 Anatomy of a dashboard




Figure 5: Anatomy of a dashboard

Every dashboard is built from the same components. Once you recognise them, you can find your way around a dashboard you have never opened before.

Query bar (1). The top row holds the search field, which accepts KQL, the query language used by Kibana: enter a condition such as a field name followed by a value, and every panel on the dashboard is restricted to the matching reviews. The two symbols to the left let you build the same restrictions without typing. The funnel symbol creates your own filter by selecting a field, an operator and a value; the plus sign adds a further one. Filters you add appear as small labels below the bar, where you can see, edit or remove each one.

Time range selector (2). At the top right of the query bar, you set the period the reviews were written in. Choose one of the predefined ranges, or enter your own start and end date. The setting applies to the whole dashboard, and every value you see is calculated only from reviews in that period. **Refresh**, next to it, reloads the dashboard with the current settings.

u^b

Controls (3). The second row holds ready-made controls for the restrictions you will need most often. **Select language(s)** limits the analysis to reviews written in one or more of the four analysed languages. **Select POI category(ies)** limits it to highlights or to the wider set of tourism-relevant POIs. **Select POI type(s)** limits it to one or more of the seven POI groups; note the reminder in the label that most POIs carry several types, so a POI can still appear when you filter for a different one. Each control shows **Any** when no restriction is active. The **Time slider** on the right moves a window through the recorded period, with a play symbol to animate it and arrows to step forward and back. Take care with it: the slider replaces the time range you set at the top right, so the two never apply at the same time. Not every dashboard has this second row. Where it is missing, use the query bar instead.

Dashboard title (4). Below the controls, the title names your destination. It is fixed for your login, so every detailed view you open shows your own destination.

Panels (5). Each chart, gauge or table sits in its own panel with a title. Select it to open **Options**, which offers **Settings**, **Inspect** to see the figures behind the chart, **Maximize** to enlarge the panel, and, under **More**, the CSV export described in section 3.5. Where a panel has a legend, the legend is interactive: select an entry to focus on it, or hide entries to reduce a crowded chart to the destinations or categories you actually want to compare.

Charts as filters. You do not have to use the controls at all. Selecting a bar, a segment or a legend entry directly in a chart adds the corresponding filter to the whole dashboard, which is usually the fastest way to drill into a category.

Tooltips. Move the mouse over a bar, a segment or a point on a line, and the exact value appears together with the category it belongs to. Use this rather than reading values off the axis, particularly on the benchmarking dashboards, where the differences between destinations are small.

Navigation links (6). The benchmarking dashboards carry a list on the left that takes you directly to the other benchmarking views. On the topic benchmarking dashboards, a second row of links switches between the four review languages. These links open a different dashboard rather than filtering the current one, so any filter and any time range you have set is not carried over. If you are comparing two dashboards, set the time range again after switching.

3.5 Exporting and sharing results

You can take results out of the platform in two ways: the figures behind a chart as a CSV file, and the chart itself as a screenshot.

Exporting the data behind a panel

1. Open the dashboard you need and set the filters and time range you want the export to reflect. The file contains exactly what is currently displayed.
2. Move the mouse over the panel you want to export, such as a chart, a gauge or a metric.
3. Select the options menu, the three dots at the top right of the panel.
4. Select **More**.

5. Select **Download CSV**.
6. The file downloads to your computer and opens in Excel, Google Sheets or any other spreadsheet application.

u^b

The file contains the aggregated values the panel displays, not the underlying reviews. It is the right basis for your own charts, for a report or for combining these figures with other data you hold. If you plan to reuse an export later, note the filter settings and the time range alongside it; the file itself does not record them.




Figure 6: The CSV export function

Exporting a chart

For the chart as an image, take a screenshot. **Full screen** at the top right of the dashboard hides the browser interface and gives a cleaner picture.

Crediting the source

When you use a chart or figures from the Interaction Quality Dashboard outside the platform, credit the source as follows:

Graphic from Interaction Quality Dashboard, University of Bern / BFH (gastfreundlichkeit.ch)

Section 8.9 explains what information you are permitted to share, particularly with regard to values that identify other destinations.

4. Where the Data Comes From

This chapter gives you enough background to read the figures correctly and to explain them to others. It is deliberately short on technical details; the separate methodology note covers the models themselves.

u^b

4.1 Data source

The Interaction Quality Dashboard analyses publicly available Google reviews. No guests are surveyed and no personal data is collected for the platform. Everything you see is derived from text that guests published themselves.

Reviews are collected for the points of interest, or POIs, that belong to your destination. Which POIs those are was defined together with experts from your own organisation, in a process the project calls regionalisation. Two groups are recorded: highlights, the POIs that shape the perception of the destination, and the wider set of tourism-relevant POIs. Businesses without tourism relevance are not recorded at all.

4.2 From review to data point




Figure 7: From review to score, processing overview

The first four steps were done once, when the method was developed. The last four run continuously whenever new reviews arrive.

A single review usually contains several statements at once. A guest may praise the staff, complain about the price, and mention the view, all in three lines. To measure hospitableness rather than overall satisfaction, the analysis therefore works at sentence level rather than at review level.

Each sentence is assigned to one of two content categories, or to neither. Statements about people and interaction count as hospitableness. Statements about the physical and material offering count as

hardware. Statements that refer to neither, such as price or process, are not counted towards either metric. Within each category, similar statements are grouped into topics, and the sentiment of each statement is calculated. The results are then aggregated to POI, POI category, and destination level.

4.3 Data cleaning and coverage

u^b

- Duplicates, empty reviews and reviews containing only emojis are removed.
- Four review languages are analysed: German, English, French and Italian. Reviews in other languages appear in the review counts but are not part of the metrics.
- From the moment a destination was regionalised, all new reviews have been collected continuously.

This last point matters for comparison. Destinations joined the project at different times, so the length of the recorded history differs from one destination to the next. Section 8.2 explains how to deal with this.

4.4 The metrics explained

All metrics are shown on a scale from 0 to 5, where 0 is the poorest and 5 the best possible value. This matches the familiar Google star scale, making the metrics easy to compare.

Metric	What it measures	Calculated from
Google rating	The star rating guests awarded	All reviews, including those consisting of stars only
Sentiment	How positive or negative a review reads overall	The full text of the review
Hospitableness	How positively guests describe interaction and human contact	Only those sentences classified as hospitableness
Hardware	How positively guests describe the physical and material offering	Only those sentences classified as hardware
Topic	How positively guests describe one specific theme within hospitableness or hardware	The sentences belonging to that topic

Note: all four metrics are calculated only from reviews of recorded POIs written in one of the four analysed languages. Within that set, the Google rating is the only metric that also includes reviews without any text, because it needs no text to exist. The Data Basis dashboard, by contrast, counts reviews in all languages, so the review totals you see there are larger than the basis of any metric. Section 8.3 returns to this.

4.5 Hospitableness and hardware: the key distinction

This distinction is what makes the Interaction Quality Dashboard different from a conventional review monitoring tool. A star rating tells you how well a business is judged. It does not tell you whether guests are talking about the people or the building. Once the two are separated, you can see whether an issue calls for staff development or investment.

Hospitableness covers the human, interaction-based side of the guest experience: how staff, hosts and other guests behave, and how that behaviour is perceived. It is not limited to reception staff; guides, instructors and waiters are all part of it. Hardware covers the tangible side of the offering: rooms, buildings, facilities, cleanliness, location, surroundings and views, and the quality of food and drink.

u^b

The examples below are real sentences from guest reviews.

Example statement	Counts as
The staff was very attentive and friendly .	Hospitableness
The guide was extremely knowledgeable and polite .	Hospitableness
Waiters very rude and impatient .	Hospitableness
The room was clean and comfortable .	Hardware
The location is great and easily accessible .	Hardware
The food was pretty disappointing .	Hardware
Beautiful hotel with super helpful and friendly staff .	Both
Great fondue and great staff .	Both
The return process was also very easy.	Neither
For this price I would have expected more but ok	Neither
Excellent!!	Neither

A sentence is only counted when its reference is unambiguous. Statements about price, waiting times, booking procedures and organisation, and general praise or criticism without a subject are counted towards neither metric. A genuine complaint about a slow process will therefore not lower either score. Keep this in mind when a metric looks better than your own impression of a business.

The word service is the most common borderline case, because guests use it for both the people and the process:

Example statement	Counts as	Why
Very friendly service	Hospitableness	Describes how people behaved
Self-service.	Hardware	Refers to a facility, not to people
Good service	Neither	No indication of what is meant
Poor customer service	Neither	A judgement without a subject

4.6 Topics

Within hospitableness and within hardware, statements that describe the same theme are grouped into topics. We did not define the topics in advance. They were derived from what guests actually wrote, and the project team then reviewed the resulting groups manually.

That manual step matters for how you read the dashboards. The model often produces several groups that describe the same theme, for example, one containing positive statements about staff and another containing negative ones, or two groups whose difference is too subtle to act on. Groups of

that kind were merged, and each remaining topic was given a name that points to a recognisable management dimension rather than to a technical cluster number. The result is a smaller set of topics that are easier to interpret and that can be compared consistently across destinations.

Merging positive and negative groups is deliberate. Because each topic now contains both, the sentiment shown for a topic tells you how guests judge that theme, rather than simply reflecting whether the topic name itself sounds positive.

A separate model runs for each review language. The topics therefore differ between languages, in number as well as in content: what German-speaking guests emphasise is not identical to what English-speaking guests emphasise. Each topic benchmarking dashboard lists the topics it covers at the top of the page. Compare topics within a language, not across languages.

4.7 Aggregation levels

Results are built up in steps: from the individual sentence to the review, from the review to the POI, from the POI to one of seven POI categories, and from there to the destination as a whole.

The dashboards show aggregated values. Individual businesses are not ranked or listed; single POIs appear only on the map when you hover over a marker.

4.8 How current the data is

Collection runs automatically. Reviews are retrieved per POI, and how often a POI is queried depends on how quickly it accumulates new reviews: busy POIs are checked every eight hours, quiet ones once a week. The system reassigns a POI between the two intervals as its review activity changes, so a business that becomes popular is picked up more frequently on its own.

The dashboards do not currently display a data-as-of date. To see how recent your data is, open the Data Basis dashboard and look at how many reviews arrived in the most recent weeks.

4.9 What the data can and cannot tell you

- It can show how guests who write reviews describe your destination, and how that compares with other destinations and with earlier periods.
- It can show whether guests are talking about people or about the offering, and which themes dominate.
- It cannot show what all your guests think. People who write reviews are not necessarily a representative sample of visitors.
- It cannot explain why a value moved. It shows what was written, not the reason it was written.
- It cannot judge an individual business. The figures are aggregates and are meant to open a conversation, not to close one.

5. Working with Filters

5.1 The filter bar



Figure 8: The filter bar

Everything that controls what you see sits in the two rows above the dashboard title. The numbers below refer to Figure 8.

1. **Filter symbol.** Builds your own filter without typing: choose a field, an operator and a value. The filter is then shown as a label below the bar, where you can edit, disable or delete it.
2. **Plus sign.** Adds a further filter alongside the first.
3. **Query field.** Accepts KQL, the query language used by Kibana. Type a condition directly if you know the field name. Anything you enter here applies to every panel on the dashboard at once.
4. **Time range.** The period the reviews were written in. Select one of the predefined ranges or enter your own start and end date.
5. **Refresh.** Reloads the dashboard with the current settings.
6. **Select language(s).** Limits the analysis to reviews written in one or more of the four analysed languages. In the figure, this control is active and set to a single language; the others show **Any**, meaning no restriction.
7. **Select POI category(ies).** Limits the analysis to highlights or to the wider set of tourism-relevant POIs.
8. **Select POI type(s).** Limits the analysis to one or more of the seven POI groups. As the label reminds you, most POIs carry several types, so a POI can still appear when you have filtered for a different one.

Time slider. Moves a window through the recorded period. The arrows step it forward and back, the play symbol animates it. The slider replaces the time range set in the query bar, so the two never apply at the same time.

The controls in the second row are only available on some dashboards. Where that row is missing, use the query field instead.

A filter stays active until you remove it, including when you leave the dashboard and come back. Before you read a value, check that every control shows the restriction you intend; **Any** means no restriction. **Reset** at the top right of the dashboard returns everything to the default state.

5.2 Filtering directly in a chart

You do not have to use the controls at all. Selecting a bar, a segment, or a legend entry inside a chart adds the corresponding filter to the whole dashboard. Selecting a bar in a chart broken down by POI

type, for example, restricts every panel to that type. This is usually the fastest way to drill into a category, and the dashboard title text points it out as an alternative to the controls above.

Filters created this way behave exactly like the ones you build yourself: they appear as labels below the query field, and you remove them there.

u^b 5.3 Combining filters

Filters apply on top of one another. Each one you add narrows the set of reviews behind every figure on the dashboard, which makes the result more specific and, at the same time, less stable.

Figure 9 shows the same dashboard twice. In (A) no filter is set. In (B) the analysis is restricted to POI category A and to reviews written in English.


Three things change, and all three are worth understanding.

The values move. The average Google rating drops from 4.189 to 3.986, and the average sentiment from 4.272 to 4.048. Neither figure is more correct than the other. They answer different questions: the first covers all recorded POIs and all four languages, the second only highlights those described by English-speaking guests.

Only category A is still shown. The panel comparing categories A and B now displays a single value, because the filter has removed category B from the data entirely. Wherever a dashboard breaks a figure down by something you have filtered on, expect the breakdown to collapse to what remains.

Bars disappear. In the chart by POI type, several bars are gone. This does not mean those POI types are missing from your destination. It means that after filtering, no reviews are left that match all conditions at once: no English-language review of a highlight POI of that type exists in the data. The chart only shows categories for which reviews remain.

That last point is the most common source of confusion. An empty or thinly populated chart after filtering is usually a statement about the data, not about the destination. Before you conclude that a POI type performs poorly, check whether it is present at all, and how many reviews stand behind the bars that remain.






Figure 9: The same panel, unfiltered (A) and filtered by category A and English reviews (B)

5.4 Resetting filters

Filters stay active until you remove them, so it is worth knowing where each type of filter lives.

Controls in the second row. An active control shows the selected value instead of **Any**, together with a small counter indicating how many values are selected. There are two ways to clear it. Select the eraser symbol that appears above the row when a control is active; it clears the selections in one step. Or open the control itself and deselect the values you no longer want, which is the better option when several values are selected, and you want to keep some of them.

Filters you built yourself. These appear as labels below the query field. Select a label to open its menu, where you can temporarily disable it or delete it. Deleting removes it; disabling keeps it visible but inactive, which is useful when you want to compare a view with and without one condition.

Filters added by clicking in a chart. These behave the same way: they appear as labels below the query field, and you remove them there.

Conditions typed into the query field. Clear the field and select **Refresh**.

Time range. This is never empty; it always holds a period. To widen it again, set a longer range rather than trying to remove it.

Reset. At the top right of the dashboard, **Reset** returns everything to the state the dashboard was saved in.




Figure 10: How to remove a filter

Before you read a value, take a moment to check the controls and the labels below the query field. A control showing anything other than **Any** is restricting every figure on the dashboard, and a filter left over from an earlier question is the most common reason for figures that do not match what a colleague sees.

5.5 What to watch out for

- Narrow filters leave few reviews behind a figure. Always check the review count before you draw a conclusion from a filtered value.
- Comparing periods of unequal length compares unequal amounts of data, not just different times.
- The language filter changes the composition of the sample, not only its size: different guest groups visit different types of businesses.

u^b

6. Dashboard Reference

Every dashboard follows the same structure: purpose, the visuals explained, and what to watch out for. The dashboards are grouped into three blocks: understanding your data basis, your destination in detail, and benchmarking against other destinations.

u^b

Group A - Understanding your data basis

6.1 Data Basis

Shows the collected data of the destination and meta-data distributions.

6.1.1 Purpose

Data Basis answers the question that should come before every other one: how much data is actually behind the figures you are about to read? It shows how many POIs are recorded for your destination, how many reviews they have attracted, when those reviews arrived, how many of them contain text at all, and which languages they were written in.

You will use it in two situations. Before interpreting a value anywhere else in the platform, to check whether it rests on enough reviews to be meaningful. And whenever a figure looks surprising, to see whether the explanation lies in the data rather than in the destination.

Key questions it answers:

- How many POIs and how many reviews does my destination contribute, split into highlights (A) and the wider set of tourism-relevant POIs (B)?
- How has the volume of incoming reviews developed, and how current is my data?
- What share of reviews contains text, and is therefore available for the sentiment, hospitableness and hardware metrics?
- Which languages do my guests write in, and how much of my data falls outside the four analysed languages?

The dashboard is laid out in three columns: Points of Interest at the top left, Reviews below it, and Descriptions on the right, where each note explains the panel beside it. Unlike the Metrics dashboards, Data Basis has no control row.

You can still filter it. Selecting a segment in one of the pie charts, for example a language or the share of reviews with text, adds the corresponding filter to the whole dashboard, and it appears as a label below the query field. Section 5.2 explains how filtering by clicking works and how to remove such a filter again. The exception is the POI panel at the top: selecting a POI figure does not filter the reviews, because POIs and reviews are held in separate data pools. To restrict the analysis by POI category, use the controls shown with the reviews, or the query field.




Figure 11: Data Basis - full dashboard view

6.1.2 The visuals explained

The panel numbers refer to Figure 11 in section 6.1.1.

Number POIs per Category (1) shows how many POIs are recorded for your destination, split into highlights (A) and the wider set of tourism-relevant POIs (B). In the example, 28 highlights and 1,343 other POIs. This is the foundation of everything else: every metric in the platform is built from reviews of exactly these POIs.

Number Reviews per Category (2) shows how many reviews those POIs have attracted, again split into A and B. Note how differently the two categories are populated: in the example, 28 highlights account for 20,699 reviews, while 1,343 other POIs account for 108,931. A handful of highlights therefore carries a considerable share of your data.

Reviews Created Count (3) plots the number of reviews written per day. Use it to see how your review volume has developed and to check how current your data is. The line ends where your most recent data ends. A drop at the far right is usually a collection effect rather than a real decline, because the most recent reviews have not all been retrieved yet.

Proportion of Reviews with/without Text (4) shows what share of reviews contains text. Only those reviews can be analysed for sentiment, hospitableness and hardware; the rest consist of stars only and count towards the Google rating alone. In the example, 42.83 per cent of reviews have no text, so the text-based metrics rest on roughly half the data behind the Google rating.

Language Proportions (5) shows which languages guests write in. Only German, English, French and Italian are analysed. In the example these account for around half of all reviews, with 48 per cent falling into Other. Those reviews are counted here but contribute to none of the metrics, so the figures on this dashboard are consistently larger than the basis of anything you read elsewhere.

Selecting a segment in either pie chart filters the whole dashboard, for example to a single language.

6.1.3 What to watch out for

- The figures on this dashboard count reviews, not the reviews behind a specific metric. A destination with 100,000 reviews may still have far fewer statements about hospitableness, because only text reviews in the four analysed languages contribute, and only some sentences are about people.
- The share shown as Other is not lost data; those reviews still count towards the Google rating. But the more your guests write in other languages, the less the text-based metrics represent your visitors as a whole.
- Reviews are retrieved per POI at intervals of eight hours or one week, depending on how quickly it accumulates new reviews, so the right-hand end of the time series is always slightly incomplete.

6.2 Points of Interest Map

Visualises the POIs of the destination on a map.

6.2.1 Purpose

The Points of Interest Map shows every POI collected for your destination on a map. It is the only place on the platform where you see the values of an individual business rather than an aggregate.

Use it to check which businesses are actually recorded and where they sit, and to look up the figures for one specific POI, for example before a meeting with that business.




Figure 12: Points of Interest Map - full dashboard view

6.2.2 The visuals explained

The map. Every collected POI appears as a dot. The colour follows the category shown in the legend at the top right: green for B, the wider set of tourism-relevant POIs, and blue for A, the highlights. Because highlights are few, blue dots are rare. Dense clusters of dots mark the parts of your destination where tourism activity concentrates.

POI details. Move the mouse over a dot to open its details. The popup shows the name of the POI, the Google place types assigned to it, its category, the collection frequency, the Google place ID as a link, and the four metrics for that POI: Average Hospitableness Metric, Average Sentiment Metric,

Supported by

Average Google Rating, and Average Hardware Metric. Where dots overlap, the popup shows several POIs one after another and indicates how many, for example, 1 of 4.

Controls. The buttons on the left zoom in and out, centre the map, and switch to full screen. Hold the Ctrl key while scrolling to zoom with the mouse wheel. The Layers panel at the top right lets you hide the POI layer or the base map.

Filters. The control row offers three restrictions specific to this dashboard. Category filters by the relevance of the POI, A or B. Types filters by the Google place types the POI owner has assigned. Frequency filters by how often the POI is queried for new reviews: every eight hours for POIs that accumulate reviews quickly, once a week for the rest. The system moves a POI between the two intervals as its review activity changes, so frequency is a reasonable proxy for how active a business is.

6.2.3 What to watch out for

- The values in the popup refer to that single POI, with the current filter settings, so they may rest on very few reviews. Treat them as a starting point for a conversation, not as an assessment. Section 8.8 explains how to handle this.
- A POI carries several Google place types at once, so it will appear even when you filter for a type you would not associate with it.
- The map shows the POIs that are collected, not every business in your destination. Businesses without tourism relevance are not recorded.
- Frequency describes how often a POI is checked, not how good it is. A low frequency simply means few new reviews arrive, which is normal for smaller and seasonal businesses.

Group B - Your destination in detail

6.3 Metrics

Displays Google ratings and sentiment metrics, calculated over the full review text.

6.3.1 Purpose

Metrics gives you the general picture of your destination: the star rating guests award and the sentiment of what they write, without separating the people from the offering. Open it first when you want an overall reading and move on to the hospitableness and hardware dashboards as soon as you want to know why a value looks the way it does.

Use it to establish a baseline before a conversation, and to check whether the stars and the written feedback point in the same direction.

Key questions it answers:

- How are we rated overall, and how positively do guests write about us?
- Do the star rating and the written feedback tell the same story?
- Which POI categories stand out, and how has the picture developed over time?




Figure 13: Metrics - full dashboard view

6.3.2 The visuals explained

Panel 1: Google Rating




Figure 14: Metrics - panel 1

Four views of the same metric: a gauge with the destination average on the scale from 0 to 5, the same average split into categories A and B, a bar chart by POI category, and a line chart over time as a moving average over seven days.

In the example the destination averages 4.189. Highlights (A) reach 4.076 and the wider set of POIs (B) 4.205, so the highlights are rated slightly lower than the destination as a whole. The bar chart shows which POI categories drive that average.

The Google rating is the only metric that includes reviews consisting of stars only, so it rests on more reviews than anything in the Sentiment Score column beside it.

Panel 2: Sentiment Score




Figure 15: Metrics - panel 2

The same four views for the sentiment score, which is calculated from the full text of a review.

In the example the destination averages 4.272, with 4.153 for highlights and 4.289 for the wider set of POIs. Read this column bar for bar against the Google Rating column: where sentiment sits below the rating, guests write more critically than they score, and where it sits above, the text is warmer than the stars suggest.

Only reviews containing text contribute here, so the two columns never rest on the same set of reviews. Both use the scale from 0 to 5.

6.3.3 What to watch out for

- The two columns are calculated from different material. Small differences between the Google rating and the sentiment score are normal and are not an error.

- Most POIs carry several Google place types, so one POI can contribute to more than one bar in the charts by POI category.
- The time series is a moving average over seven days. Use it for trends rather than for single days, and treat the right-hand end as incomplete, because the most recent reviews have not all been collected yet.

u^b

6.4 Hospitableness Metrics

Displays sentiment metrics, calculated over the sentences concerning hospitableness.

6.4.1 Purpose

Hospitableness Metrics is where you look once you know that something is happening with the human side of the guest experience and want to know what exactly. It reports the hospitableness score for your destination, breaks it down into topics, shows how both develop over time, and tells you how many statements each topic rests on.

Use it to decide where staff-related quality work would pay off, and to prepare conversations in which you need more than an overall score.

Key questions it answers:

- How positively do guests describe encounters in my destination, and is that changing?
- Which aspects of hospitableness do guests praise, and which do they criticise?
- Do guests writing in different languages emphasise different things?
- Which POI categories carry the destination, and which drag it down?

Note the sentence under the title: all metrics on this dashboard are based on English, German, French and Italian reviews only.




Figure 16: Hospitableness Metrics - full dashboard view

6.4.2 The visuals explained

Panel 1: Hospitableness Sentiment




Figure 17: Hospitableness Metrics - panel 1

A gauge with the destination average on the scale from 0 to 5, the same value split into categories A and B, and a line chart over time as a moving average per month.

In the example the destination averages 4.311, with 3.384 for highlights and 4.369 for the wider set of POIs. That gap is large and worth understanding: the few highlights are described noticeably less warmly than the destination as a whole.

Panel 2: Hospitableness Topics Sentiments




Figure 18: Hospitableness Metrics - panel 2

One bar chart per review language, with one bar per topic. In the example, English reviews give 4.873 for Attentiveness and Helpfulness, 4.391 for Expert Experience Guidance and 3.745 for Tone and Demeanor in Guest Contact. The German chart shows different topics, because a separate model runs per language: 4.815 for Attentiveness and Helpfulness, 4.818 for Warm and Inviting Atmosphere and 2.267 for Tone and Demeanor in Guest Contact.

Read these bars as a ranking within one language, not as a verdict on the destination. Topics differ in how much criticism they attract by nature, so a low value tells you where guests are most critical, not that the destination fails there.

Panel 3: Hospitableness Topic Sentiments over Time




Figure 19: Hospitableness Metrics - panel 3

The same topics as a monthly time series, one line per topic and one chart per language. Use it to check whether a low topic value is a persistent pattern or a recent movement. Lines that break off indicate months in which no statements on that topic were recorded.

Panel 4: Hospitableness Topic Count




Figure 20: Hospitableness Metrics - panel 4

How many statements each topic rests on, per language. In the example, German reviews contribute 2,435 statements to Attention and Helpfulness, English 1,635, French 386 and Italian 254. This is the panel that tells you how much weight to give a value in panel 2. A topic with a few dozen statements moves easily; one with thousands does not.

Panel 5: Hospitableness by POI type




Figure 21: Hospitableness Metrics - panel 5

A bar chart of the hospitable score per POI category, plus two time series: one comparing categories A and B, one showing each POI category separately. This is where you see which parts of the service chain differ from the destination average.

All bars in the first chart share the same colour, so read the label beneath each bar rather than relying on colour to tell the categories apart.

6.4.3 What to watch out for

- Topics are not comparable with one another. Compare a topic across destinations or over time, not against the topic beside it.
- Topics differ between languages in number and in content. The German set here contains Warm and Inviting Atmosphere, the English set contains Expert Experience Guidance instead.

Supported by

innovation
tourismus

- Always read panel 2 together with panel 4. A striking value resting on few statements is an indication, not a finding.
- The gap between categories A and B can be substantial. Check which of the two you are looking at before quoting a number.

u^b 6.5 Hardware Metrics

Displays sentiment metrics, calculated over the sentences concerning hardware.

6.5.1 Purpose

Hardware Metrics is the counterpart to Hospitableness Metrics and is built identically. It reports how positively guests describe the physical and material side of the offering: rooms, facilities, cleanliness, location, views, food and drink.

Use it to separate an infrastructure problem from a people problem, and to see which themes drive the difference.

Key questions it answers:

- How do guests judge the offering itself, and is that changing?
- Which hardware themes are praised, and which are criticised?
- Where do hospitableness and hardware diverge, and what does that imply for where to act?




Figure 22: Hardware Metrics - full dashboard view

6.5.2 The visuals explained

Panel 1: Hardware Sentiment




Figure 23: Hardware Metrics - panel 1

Gauge, split by category and a monthly time series, exactly as on the hospitableness dashboard. In the example the destination averages 4.269, with 4.271 for highlights and 4.268 for the wider set.

Note the contrast with hospitableness, where the same two categories were 3.384 and 4.369: the highlights are judged normally on their infrastructure and offering, but distinctly less warmly on encounters. That is precisely the kind of finding this platform exists to surface.

Panel 2: Hardware Topics Sentiments




Figure 24: Hardware Metrics - panel 2

One chart per language, one bar per topic. English reviews in the example give 4.851 for Swiss Specialties, 4.844 for Location, Views and Accessibility, 4.549 for Thematic Visitor Experiences, 4.429 for Culinary Offering and 3.798 for Comfort, Amenities and Cleanliness. The German chart contains Value for Money at 3.1 instead, a topic that does not exist in the English model.

Panel 3: Hardware Topic Count




Figure 25: Hardware Metrics - panel 3

The number of statements per topic and language. In the example, German reviews contribute 14,106 statements to Culinary Offering and English 7,279, while Swiss Specialties rests on 642 English statements. The volume differences are large; read the sentiment values against them.

Panel 4: Hardware by POI type




Figure 26: Hardware Metrics - panel 4

The hardware score per POI category, with time series by category and by POI type. As on the hospitableness dashboard, the bars share one colour, so read the labels beneath them.

6.5.3 What to watch out for

- The most useful reading of this dashboard is comparative. Open it next to Hospitableness Metrics and look for categories where the two diverge; that is where a decision follows.
- Value for Money appears as a hardware topic in some languages. Statements mentioning only price, without referring to anything concrete, are not counted at all, so this topic captures price in relation to something specific rather than price perception in general.
- The topic sets differ per language here as well.

Group C - Benchmarking against other destinations

6.6 Overview

Compares the four key metrics across destinations: sentiment, hospitableness, hardware and Google rating.

6.6.1 Purpose

Overview is the entry point to the benchmarking dashboards. It compares your destination with all other participating destinations across four metrics: sentiment, hospitableness, hardware and Google rating.

Start here. From the navigation list on the left you then move to the more detailed benchmarking views: Hospitableness Topic Benchmarking, Hardware Topic Benchmarking and POI Type Benchmarking.

Overview tells you whether your destination differs from the others at all, and on which metric; the other three tell you where that difference comes from.




Figure 27: Overview - full dashboard view

6.6.2 The visuals explained

Navigation. The list at the top left links to the other three benchmarking dashboards. The entry you are currently on is highlighted. These links open a different dashboard, so any filter or time range you have set is not carried over.

Metrics shown in this section. The box below the navigation lists the four metrics presented on this dashboard: Sentiment, Hospitableness, Hardware and Google Rating.

Average Metrics per Destination. Each metric has its own panel, and within a panel each destination is one bar, coloured according to the legend on the right. The panel shown in the figure is the sentiment metric; scroll down for the remaining three. The legend is interactive: select a destination to focus on it, or hide destinations to reduce the chart to the comparison you actually want.

The note above the charts. Read it before drawing conclusions. All metrics on this dashboard rest on reviews written in German, English, French, or Italian; reviews in other languages are not included in

any of them. Within that set, the Google rating additionally covers reviews consisting of stars only, which the text-based metrics cannot use.

6.6.3 What to watch out for

u^b

- The vertical axis adjusts to the values shown and does not necessarily start at zero. All metrics use a scale from 0 to 5, so check the axis before you judge a difference by the height of a bar. Use the value labels or the tooltips instead.
- Differences between destinations are small, usually in the second decimal place. Treat the chart as a way to spot patterns, not as a ranking.
- Destinations joined the project at different times, so the bars do not all rest on periods of equal length. Section 8.2 explains how to deal with this.
- A high Google rating alongside a lower hospitableness value is not a contradiction. Both cover the same four languages, but the Google rating also counts reviews without text, and hospitableness is calculated only from what guests actually wrote about people.

6.7 POI Type Benchmarking

Compares average sentiment, hospitableness and Google rating scores by POI category across destinations.

6.7.1 Purpose

POI Type Benchmarking compares your destination against all other participating destinations, separately for each of the seven POI categories. It answers the question that Overview leaves open: not whether you differ, but in which part of the service chain.

Use it before a conversation with a sector association, or when you need to know whether a weakness is destination-wide or concentrated in one type of business.

Key questions it answers:

- How do my hotels, my gastronomy, my transport compare with the same category elsewhere?
- Is a difference in the overall figure driven by one category, or spread evenly?




Figure 28: POI Type Benchmarking - full dashboard view

6.7.2 The visuals explained

The POI grouping table. The top of the dashboard lists which Google place types belong to which of the seven categories. Consult it whenever a result surprises you; the categories are broad, and Food, Drink and Nightlife, for example, includes every Google restaurant subtype.

u^b Panel 1: Average Metric Sentiment Score




Figure 29: POI Type Benchmarking - panel 1

One panel per POI category, one bar per destination, coloured according to the legend. In the example panel for Transportation and Mobility, values range from 3.7 to 5, with your own destination among them. Like every metric in the platform, these charts cover reviews in the four analysed languages.

Panel 2: Average Hospitableness Metric




Figure 30: POI Type Benchmarking - panel 2

The same structure for hospitableness, covering the same four review languages as the panels above and below.

Panel 3: Average Google Rating




Figure 31: POI Type Benchmarking - panel 3

The same structure again for the star rating.

6.7.3 What to watch out for

- A bar reaching exactly 5, as one destination does for Transportation and Mobility in the example, almost always signals a very small number of reviews rather than perfection. Check the count before you draw a conclusion.
- Missing bars mean no data for that destination in that category, not a score of zero.
- All three panels rest on the same four languages. They still differ in one respect: the Google rating includes reviews consisting of stars only, while sentiment and hospitableness need text.
- Some categories contain very few POIs in a small destination. The narrower the category, the more careful the reading.

These dashboards exist in four versions, one per review language, because a separate topic model runs for each language. Choose the version matching the guest language you want to understand, and read the topic list shown at the top of the dashboard.

6.8 Hospitableness Topic Benchmarking

Displays average hospitableness metric scores across multiple topics and destinations.

6.8.1 Purpose

Hospitableness Topic Benchmarking compares destinations topic by topic. Where Overview tells you that your hospitableness score differs, and POI Type Benchmarking tells you in which part of the service chain, this dashboard tells you which aspect of the encounter guests are actually talking about.

It exists in four versions, one per review language. The links above the chart switch between them; they open a different dashboard, so any filter or time range you have set is not carried over.




Figure 32: Hospitableness Topic Benchmarking - full dashboard view

6.8.2 The visuals explained

Language and covered topics. The subtitle names the review language, and the box below lists the topics covered in that language. In the English version these are Tone and Demeanor in Guest Contact, Expert Experience Guidance and Attentiveness and Helpfulness; the German version has Warm and Inviting Atmosphere in place of Expert Experience Guidance. Always check this box first, because the panels below depend on it.

Panel 1: Average Hospitableness Metric, English reviews




Figure 33: Hospitableness Topic Benchmarking - panel 1

One panel per topic, one bar per destination, with a How to read this chart box above. In the example for Topic 1, Tone and Demeanor in Guest Contact, English reviews put the destinations between 3.055 and 3.856.

Use the legend to hide destinations and reduce the chart to the comparison you actually want.

Panel 2: The same view in another language




Figure 34: Hospitableness Topic Benchmarking - panel 2

The German version of the same dashboard, reached through the language links. Here the same topic, Tone and Demeanor in Guest Contact, places the destinations between 4.6 and 4.95. That reflects a different model and a different set of statements, not a different reality.

6.8.3 What to watch out for

- The vertical axis adjusts to the values shown. In the example the English Topic 1 panel ends at 4, not at 5. Compare values, not bar heights.
- Never compare across language versions. A topic scored 3.3 in the English dashboard and 4.7 in the German one says nothing about a difference between guest groups; the models and the topic definitions differ.
- Differences between destinations within one panel are typically a few tenths. Look for destinations that sit consistently high or low across several topics rather than for a single gap.

6.9 Hardware Topic Benchmarking

Displays average hardware metric scores across multiple topics and destinations.

6.9.1 Purpose

Hardware Topic Benchmarking is the counterpart to the previous chapter and works identically. It compares destinations for each hardware topic, in four language versions.

Use it to see whether a weakness on the offering side is specific, for example to comfort and cleanliness, or spread across everything guests mention.




Figure 35: Hardware Topic Benchmarking - full dashboard view

6.9.2 The visuals explained

Language and covered topics. The English version covers Culinary Offering, Comfort Amenities and Cleanliness, Location Views and Accessibility, Swiss Specialties and Thematic Visitor Experiences. The French version covers four topics, among them Views Ambience and Experience Spaces, which does not exist in the English set.

Panel 1: Average Hardware Metric, English reviews




Figure 36: Hardware Topic Benchmarking - panel 1

One panel per topic, one bar per destination. In the example, Topic 1 Culinary Offering places the destinations between 4.421 and 4.715, a narrow band. Topic 2 Comfort, Amenities and Cleanliness spreads more widely, from 3.822 to 4.5. The wider the spread within a topic, the more that topic distinguishes destinations, and the more useful it is as a starting point for action.

Panel 2: The same view in another language




Figure 37: Hardware Topic Benchmarking - panel 2

The French version, reached through the language links. Note that the covered topics listed at the top differ from the English set.

6.9.3 What to watch out for

- The same three cautions as in the previous chapter apply: an adjusting axis, no comparison across languages, and small differences between destinations.
- Culinary Offering carries by far the most statements in every language, so it dominates the overall hardware score. A destination can look average overall while differing sharply on a smaller topic.

7. Typical Workflows

The three workflows below correspond to the three levels at which DMOs use the Interaction Quality Dashboard: a conversation with a single service provider, an exchange with a sector association, and a strategic review of the destination as a whole.

u^b

7.1 Preparing a conversation with a single service provider

The situation: you have a meeting with an important mountain railway, hotel or attraction. The business matters for how the destination is perceived, and you would rather not argue from gut feeling or from a handful of individual complaints.

The question: how do guests actually perceive this business, and do their judgements rest on hospitableness, on hardware, or on other features of the offering?

7.1.1 How to get there

1. Open the Points of Interest Map. It is the only dashboard that reports figures for a single business.
2. Set the time range before you read anything. A period that all participants cover, for example the last two years, is a reasonable default.
3. Narrow the map if the business is hard to find. Types filters by the Google place types, for example to show only accommodation or only transport. Alternatively zoom in on the part of the destination you are interested in, holding Ctrl while scrolling.
4. Move the mouse over the POI. Its details open in a popup.
5. Note the values together with the filter settings, then open Hospitableness Metrics and Hardware Metrics to see the destination averages you will compare them against.

7.1.2 What you see




Figure 38: Preparing a conversation with a service provider, key view

The popup names the business, lists the Google place types its owner has assigned, and states the category, the collection frequency and the Google place ID. Below that come the four metrics for this POI: Average Hospitableness Metric, Average Sentiment Metric, Average Google Rating and Average Hardware Metric.

In the example the business is a category B POI carrying four place types, among them train station and transit station. It scores 4.064 on the Google rating, 3.676 on sentiment, 3.607 on hardware and 3 on hospitableness. The counter at the bottom reads 1 of 3, meaning three POIs sit close enough together at this zoom level to share the popup; step through them before assuming you are looking at the right one.

7.1.3 Reading the four values together

The value of this view lies in the differences between the metrics, not in any one of them.

The **Google rating** is what the business itself sees and usually already knows. It rests on every review, including the many that consist of stars only, and it tends to be the most generous of the four figures.

The **sentiment score** covers the full text of those reviews that have text. When it sits clearly below the star rating, guests write more critically than they score. That gap is common and not alarming in itself, but it tells you that the written feedback contains more than the stars suggest.

The **hospitableness** and **hardware** scores are the pair that makes the conversation concrete, because they separate the people from the offering. Three constellations occur regularly:

u^b

- Hardware clearly below hospitableness: guests like the people and criticise the rooms, the facilities or the food. The lever is investment or maintenance, and the staff are carrying the business.
- Hospitableness clearly below hardware, as in the example, where 3 stands against 3.607: the offering is judged acceptably, and the encounters are not. The lever is staffing, training or workload, and the business will often recognise the pattern once it is named.
- Both close together and both below the destination average: the difference is not specific, and the figures alone will not tell you where to start. Use the topic views on the Hospitableness Metrics and Hardware Metrics dashboards to see which themes dominate and treat the meeting as an opportunity to ask rather than to conclude.

To judge whether a value is actually low, compare it with the average for the same POI category on the POI Type Benchmarking dashboard and with your destination average on the Hospitableness Metrics and Hardware Metrics dashboards; a hotel scoring below the destination average but in line with other accommodation is a different conversation from one that trails both.

What the values cannot tell you is why. A low hospitableness score may reflect a staffing gap during one season, a single member of staff, or a guest mix that arrives with different expectations. Bring the figure, and let the business supply the explanation.

7.1.4 Before you use it

The four values reflect whatever filters and time range are currently set. Narrow the period, or restrict the analysis to one language, and the same POI will show different figures. Check the time range at the top right and the controls before you write anything down, and note the settings alongside the values so that you can reproduce them later.

A hospitableness value of exactly 3 is a warning sign in itself. Round figures usually mean very few statements stand behind them, sometimes only one or two, and a narrow filter makes that more likely. The map does not show how many, so treat a single-POI value as a reason to ask a question, never as a finding to present. Section 8.8 sets out how to discuss such figures with a business.

7.2 Preparing the annual meeting with the hoteliers association

The situation: you are preparing the annual exchange with the hoteliers association and want a shared understanding of how guests perceive hotels, and whether action is needed on the quality of encounters, service attitude, or hardware.

The question: what patterns show up across hotels as a whole, and where do overall rating, sentiment, hospitableness and hardware diverge?

7.2.1 How to get there

1. Open Hospitableness Metrics and set the time range you want to discuss. Scroll to the chart Average Hospitableness sentiment metric per POI Type and read the value for Accommodation.
2. Open Hardware Metrics and read the Accommodation value in the equivalent chart. You now know how hotels are judged on both dimensions inside your own destination.
3. Open POI Type Benchmarking. Each of the three metric sections, sentiment, hospitableness, and Google rating, contains one panel per POI category. Scroll to the panel labeled 2. Accommodation in each section.

4. Identify your own destination in the legend and hover over the bars to read the exact values. Hiding the other destinations through the legend makes your own bar easier to find.
5. Note the values together with the time range and any filters, so that the figures you bring to the meeting can be reproduced.

u^b

7.2.2

What you see




Figure 39: Preparing a sector meeting, key view

The first two charts answer the internal question: how does accommodation compare with the other parts of the service chain in your own destination? In the example, hotels reach roughly 4.5 on hospitableness, above transport at 3.9 and close to recreation at 4.6. On hardware, they reach roughly 4.2, below attractions at 4.6 and recreation at 4.5, but above transport at 3.7.

The benchmarking panels answer the external question: how does accommodation here compare with accommodation elsewhere? In the example the Accommodation panel spans 4.251 to 4.682 for

sentiment and 4.339 to 4.71 for hospitableness. Both bands are narrow, which is typical: at this level of aggregation destinations differ by a few tenths, not by whole points.

7.2.3 How to read it

u^b

Take the two questions in that order. First, are hotels a strong or a weak part of your own service chain? Second, are your hotels strong or weak compared with hotels in other destinations? The answers can diverge, and the combination is what makes the meeting productive.

Hotels above your destination average but at the lower end of the benchmarking band means the association is doing well relative to the rest of your destination, while the wider field has moved further. That is a conversation about ambition, not about deficits.

Hotels below your destination average but in the middle of the benchmarking band means the pattern is not specific to your hotels. It is normal for accommodation to score differently from attractions, because the interaction is longer and expectations are higher. Present it as context rather than as a finding.

If hotels fall short in both of these aspects, this is a reason to take action, and distinguishing between hospitableness and hardware then provides insight into what kind of measures are needed. Hardware below hospitableness points to the substance of the offering, rooms, facilities, breakfast.

Hospitableness below hardware points to encounters: staffing levels, training, workload at peak times. If you want to go one step further before the meeting, the topic panels on the Hospitableness Metrics and Hardware Metrics dashboards show which themes drive the value.

7.2.4 Before you use it

Missing bars in the benchmarking panels mean that no data is available for that destination in that category, not that it scored zero. Do not read a gap as a result.

The bars all start at zero and run to five, so visually the differences look tiny. Work from the value labels, not from bar heights.

All panels on POI Type Benchmarking cover the same four review languages, so sentiment, hospitableness and Google rating can be quoted side by side. The one difference worth mentioning is that the Google rating also counts reviews without text.

And the figures describe accommodation as a group. They cannot identify which hotels drive the average, and the platform deliberately does not rank named businesses. If the association asks who is pulling the number down, that is the point to say that the data does not answer it, and that individual conversations are the right place for that question.

7.3 Positioning the destination for a strategy or board meeting

The situation: you are preparing a strategy paper, a budget, or a board meeting and want to know how your destination compares with similar destinations.

The question: are we strong or weak relative to others, and does that concern overall perception, hospitableness, or rather hardware and the offering?

7.3.1 How to get there

1. Open Overview. Set the time range first and choose a period that all destinations cover; destinations joined the project at different times.

u^b

2. Read the four panels in turn: Sentiment, Hospitableness, Hardware, and Google Rating. Each shows one bar per destination.
3. Find your own destination in the legend. Selecting entries in the legend hides the other destinations, which makes your own bar easy to locate and the chart easier to read.
4. Note your value and the range of the field for each of the four metrics.
5. If a metric stands out, follow it further: POI Type Benchmarking shows whether the difference sits in one part of the service chain, and the two topic benchmarking dashboards show which theme guests are talking about.

7.3.2 The three levels of comparison

The benchmarking dashboards are built to be used in this order, from broad to specific.

Overview compares whole destinations on four metrics. It answers: do we differ at all, and on which metric?

POI Type Benchmarking compares the same destinations within one POI category at a time. It answers: is the difference spread across the destination, or concentrated in accommodation, gastronomy, transport?

Topic benchmarking, hospitableness, and hardware compare destinations within one theme. It answers: what exactly are guests talking about when they judge us differently?

For a board discussion, the first level usually carries the message and the other two supply the explanation.

7.3.3 What you see




Figure 40: Strategic positioning, key view

The box at the top left lists the four metrics presented here. The note above the charts states an important restriction: every metric here is based on reviews written in English, German, French or Italian, and reviews in other languages do not enter any of them. In the example the sentiment panel places the destinations between 4.276 and 4.595. Nineteen destinations sit within about three tenths of one another, which is the pattern you should expect across all four panels.

7.3.4 How to read it

Start with your position in the field rather than with the absolute number. A value of 4.4 means little on its own; being at the bottom of a band that runs from 4.28 to 4.60 means something, and so does being at the top. Do this against a peer group you have chosen deliberately, not against all destinations at once.

u^b

Then compare the four panels. The interesting result is rarely a low value; it is an inconsistent one. A destination in the upper half on the Google rating and the lower half on hospitableness tells you that guests award stars generously but describe encounters less warmly than elsewhere. That is a finding a board can act on. A destination that sits mid-field on all four metrics has no story to tell from this dashboard, and saying so plainly is a legitimate result.

Finally, decide whether the difference is worth pursuing. A gap of a few hundredths is not a position in a ranking. Look for a destination that sits consistently at one end across several metrics, or for a gap that has persisted over more than one period.

7.3.5 Before you use it

Not every destination is a fair comparison. The field includes cities, alpine resorts, and small rural destinations, and they differ in the mix of businesses they contain, the segments they attract, and what guests expect of them. A city with a dense transport network and a large share of business travel is judged on different things than a mountain resort where guests stay a week. The data basis differs as well: destinations joined the project at different times, so the periods and the review volumes behind the bars are not identical.

None of that makes the comparison worthless, but it changes how you use it. Rather than reading the whole field as one league table, pick the destinations that genuinely resemble yours in size, type, and guest structure, hide the rest through the legend, and compare against those. Where your position differs sharply between comparable and non-comparable destinations, the structural difference is part of the explanation, not a flaw in the data.

The vertical axis adjusts to the values shown, and the bars all start at zero, so the visual impression understates real differences and overstates trivial ones. Work from the value labels.

Destinations are named. Anything you take out of this dashboard shows how your peers perform, and they see you under the same conditions. Section **8.9** covers what that means for circulating figures outside your organisation.

The Google rating and the text-based metrics do not rest on quite the same reviews: the rating also counts reviews consisting of stars only. If you place them side by side in a board paper, a footnote saves you the question.

And the data shows what guests wrote, not why. Overview can tell you that your position on hospitableness has slipped relative to others. It cannot tell you whether that is staffing, a change in the guest mix, or a single large business with many reviews. Present it as the starting point of an investigation, not as its conclusion.

8. Interpreting Results Responsibly

The Interaction Quality Dashboard makes previously invisible patterns visible. It does not make them certain. This chapter lists the points where results are most easily overinterpreted.

u^b

8.1 How many reviews stand behind a figure

There is no universal minimum. What matters is that a value based on few statements is far less stable than a value based on many: with a small number of statements, a single unusual review moves the average visibly, and repeated measurements of the same thing scatter more widely. Values based on large numbers settle down and change slowly.

Before you act on a figure, open the Data Basis dashboard, or narrow your filters step by step and watch how much the value moves. If it jumps around, treat it as an indication rather than as a measurement.

8.2 Destinations have different data histories

Destinations joined the project at different times, and data collection started when each was regionalised. One destination may therefore have several years of reviews while another has considerably less. When you compare destinations, set an identical time range for all of them, and prefer periods that all participants cover.

8.3 Star ratings and text analysis measure different things

A large share of Google reviews consists of stars only, with no text at all. Those reviews count towards the Google rating, but they contain nothing to analyse, so they cannot contribute to sentiment, hospitableness or hardware.

As a result, the Google rating and the text-based metrics rest on different sets of reviews. A destination can look strong on stars and less strong on the text metrics without either figure being wrong. Where the two diverge, the text metrics usually carry the more specific information, because guests who take the time to write tend to say what they mean.

8.4 Not everything guests complain about is measured

Statements about price, waiting times, booking and organisation, along with unspecific praise or criticism, are not assigned to either hospitableness or hardware; see section 4.5. A destination can therefore have a real, recurring problem that neither metric captures. Use the topic views and your own knowledge together, not the scores alone.

8.5 Language effects

Guest groups differ in where they go, what they notice and how they write. Filtering by review language therefore changes the composition of the sample, not just its size. In addition, a separate topic model runs for each language, so the topics themselves differ. Compare languages with care, and compare topics only within one language.

Culture plays into this as well. What counts as warm, attentive, or correct service differs between guest groups, and so does how praise and criticism are expressed: some guests write effusively, while others state a shortcoming plainly and leave the rest unsaid. A lower value in one language therefore does not automatically mean a poorer experience.

English reviews are a special case. English is written by guests from many countries, so an English-language result reflects an international mix rather than one culture. Read it as the view of your international guests, not as the view of guests from English-speaking countries.

8.6 Differences are small, and that is expected

u^b

Across participating destinations the overall level is high. Most differences appear at the second decimal place. That does not make them meaningless, but it does mean that a gap of a few hundredths is not a ranking. Look for patterns that persist over time and across POI categories rather than for a single decisive number.

8.7 No causal conclusions

The data shows what guests wrote, not why they wrote it. A falling hospitableness value may reflect a staffing change, a difficult season, a shift in the guest mix, or a single business with many reviews. What the platform does show is a development, and where that development is worth a closer look: it tells you which questions to ask, and of whom, rather than supplying the answer.

8.8 Communicating results to businesses

Figures about a business, discussed with that business, are useful only as a starting point for a conversation. We recommend the following:

- Present findings as a basis for discussion, not as an assessment or a verdict.
- State the period the figures cover and how many reviews they rest on whenever you quote them. A value without a time range invites comparison with a value from a different one.
- Show hospitableness and hardware side by side, so that the conversation is about where to act rather than about whether the score is fair.
- Avoid rankings of named businesses. The platform deliberately does not produce them.
- Let the business add its own context. The data cannot see staffing, refurbishment or seasonality.

8.9 Data protection and confidential figures

The Interaction Quality Dashboard processes publicly available reviews. No guest is identified, and no personal data about guests is stored for the platform. Reviews are analysed at sentence level and reported only in aggregated form.

Note: the platform never displays the reviews themselves. You see aggregated values, topic results and counts, but not the text of an individual review.

The detailed figures for your destination are intended for your own quality work and for conversations with your partners. Benchmarking values that name other destinations should not be published or passed on outside your organisation without agreement, since every participant sees the others under the same expectation of restraint.

9. Glossary

Term	Definition
Benchmarking	A comparison of the same metric across participating destinations or POI categories.
Destination	The area covered by one participating DMO, defined by the set of POIs assigned to it during regionalisation.
Google rating	The star rating guests award on Google, from 1 to 5. The only metric that also includes reviews without text.
Hardware	The tangible, non-human side of the offering: rooms, facilities, cleanliness, location, surroundings, views, food and drink.
Hospitableness	The human, interaction-based side of the guest experience: how staff, hosts and other guests behave, and how guests perceive that behaviour.
POI, Point of Interest	A single place recorded in the system, for example a hotel, restaurant, mountain railway or museum.
POI category	One of seven groups into which the Google place types are combined for comparison. The full list is shown on the POI Type Benchmarking dashboard.
Regionalisation	The process of deciding, together with the DMO, which POIs belong to a destination.
Review	A publicly available guest rating on Google, consisting of stars and, in many cases, text.
Sentence classification	The step that decides, for each sentence, whether it refers to hospitableness, to hardware, to both or to neither.
Sentiment	How positive or negative a text reads, expressed on the 0 to 5 scale.
Topic	A theme within hospitableness or hardware, derived from what guests wrote and named by the project team. Topics differ by review language.

10. FAQ and Support

10.1 Frequently asked questions

Why is my Google rating higher than my hospitableness score?

u^b

Because the two are calculated from different reviews. The Google rating counts every review, including the many that consist of stars only and contain nothing to analyse. Hospitableness is calculated only from sentences about people, in reviews that contain text. A generous star rating alongside a more moderate hospitableness value is the normal pattern, not an error. Section 4.4 sets out what each metric is calculated from.

How do I know whether a value is reliable?

Check how much data stands behind it. Open the Data Basis dashboard for the destination as a whole, or the topic count panels on the Hospitableness Metrics and Hardware Metrics dashboards for a specific theme. There is no fixed minimum, but a value resting on a few dozen statements moves as soon as one unusual review arrives, while a value resting on thousands changes slowly. If narrowing your filters step by step makes a figure jump around, treat it as an indication rather than a measurement.

Why does a colleague see different figures than I do?

Almost always because of filters. A control showing anything other than Any, a filter label below the query field, or a different time range changes every value on the dashboard at once. Compare settings before comparing numbers. Section 5.4 explains where each type of filter lives and how to remove it.

Why have bars disappeared from a chart after I filtered?

Because no reviews are left that match all your conditions at once. The chart shows only the categories for which data remains, so a missing bar means no data, not a score of zero. This says something about your data basis, not about the quality of that POI category.

Can I see the figures for a single business?

Only on the Points of Interest Map. Move the mouse over a POI and the popup shows its Google rating, sentiment, hospitableness and hardware values. Every other dashboard reports aggregates, and the platform deliberately produces no ranking of named businesses. Values for a single POI usually rest on very few statements; section 8.8 explains how to use them in a conversation.

Why do the topics differ between the language versions?

Because a separate model runs for each review language, and each set of topics is derived from what guests actually wrote in that language. English reviews produce a topic on Expert Experience Guidance, German reviews produce one on Warm and Inviting Atmosphere instead. Compare a topic across destinations or over time, never across languages. Section 4.6 explains how the topics were formed.

Guests complain about waiting times and prices. Why does that not show up anywhere?

Because those statements are not assigned to either metric. A sentence is only counted when it clearly refers to people or to the offering. Price on its own, waiting times, booking and organisation, and unspecific praise or criticism carry no such reference and therefore lower neither the hospitableness nor the hardware score. Your destination can have a real, recurring problem that the scores do not capture; section 4.5 lists examples.

Which businesses are included, and why is a particular one missing?

The platform covers the POIs that were assigned to your destination during regionalisation: the highlights, category A, and the wider set of tourism-relevant POIs, category B. Businesses without tourism relevance are not recorded at all. If a POI you expect is missing, or one appears that should not be there, report it to the support address in section **10.3**.

How current is the data?

Reviews are collected automatically, per POI. Busy POIs are queried every eight hours, quieter ones once a week, and the system moves a POI between the two as its review activity changes. The dashboards do not display a data-as-of date, so to see how recent your data is, open the Data Basis dashboard and look at how many reviews arrived in the most recent weeks.

The differences between destinations look tiny. Is the comparison worth anything?

Yes, but not as a ranking. Across participating destinations the overall level is high and most differences appear in the second decimal place. What carries meaning is a destination that sits consistently at one end across several metrics or several periods, and a comparison against destinations that genuinely resemble yours in size, type and guest structure. Section **8.6** and the workflow in section **7.3** explain how to work with this.

May I share these figures outside my organisation?

Treat the detailed figures for your destination as material for your own quality work and for conversations with your partners. Benchmarking views name the other destinations, and they see you under the same conditions, so agree before publishing or circulating values that identify others. When you use a chart externally, credit the source as shown in section **3.5**.

10.2 Reporting a problem

When you report a problem, please include the dashboard name, the filter settings and time range you applied, when the problem occurred, and a screenshot. This is usually enough to reproduce the issue without further questions.

10.3 Contact

For now, questions about access, technical problems, and content all go to the same address. There are no guaranteed response times.

General inquiries: adrian.mueller@unibe.ch

Technical inquiries: mascha.kurpicz@bfh.ch