The Hotel Prompt Library
AI Visibility Across the Hotel Journey

A 3,000-prompt test set that shows where in the hotel journey, from choosing an area to checking out, a hotel or booking brand appears in AI answers, and where it disappears.

Research method · AI search visibility

FOUR DIMENSIONS, ONE PROMPT SET

20	18	10	32
cities	traveller groups	journey steps	question topics

H0001-T1 · What kinds of places to stay in Paris would suit us? We are two adults with children aged 6 and 10. We are staying one night.

3,000	64	4
prompts	prompt first wave	AI tools

Ameneh Saeednia
Senior SEO and Generative Search Strategist

01 · The problem

AI visibility tools ask the wrong questions

Most AI visibility tracking runs short, generic prompts such as **"best hotels in Paris"**. Real travellers don't ask like that. They describe their situation, and their questions change as they move from an idea to a booked room, and through the stay itself.

How tracking usually works

A few hundred **topic prompts** per brand.
Results grouped by **broad topic**: luxury hotels, family hotels, cheap hotels.
One number: how often the brand is mentioned.

What it can't tell you

Which traveller the brand fails: wheelchair users, light sleepers, travellers with pets?
Which journey step it drops out of: inspiration, comparison, booking?
What to fix: which page, which proof, which question is unanswered?

The idea

Build prompts the way travellers think: **a real situation, a real constraint, a step in the journey**. Then score every AI answer, so gaps can be traced to a specific traveller, step, topic or city.

The library answers	For example
Where in the journey does the brand disappear?	Is the brand there when people ask *which area should we stay in?*, but gone when they *compare hotels* or *check the fine print*?
Which travellers are underserved?	Do answers for **wheelchair users** name any hotel or booking site at all?
Which sources does AI trust instead?	Booking platforms, review sites, forums, travel blogs or official tourism sites?
Does the AI name brands at all?	Does an advice-only answer still **cite** a brand's page as a source?

02 · Design

Four dimensions

Four dimensions combine into **3,000 prompts**, each with two follow-up questions.

1 · Cities — 20

Paris · Rome · Florence · London · Barcelona · Lisbon · Amsterdam · Prague · Vienna · Athens · Dubrovnik · Naples · Istanbul · Kyoto · Tokyo · Singapore · Vancouver · Cape Town · Reykjavik · Marrakech

2 · Traveller groups — 18

Family with kids · Family with teens · Baby and cot · Mixed-age family · Couple · Friends · Solo traveller · Business traveller · Wheelchair access · No stairs · Light sleeper · Dietary needs · Travelling with a pet · Early arrival, late departure · First overseas trip · Returning visitors · Limited budget · Celebration

3 · Journey steps — 10

1 Explore ideas · 2 Choose an area · 3 Hotel or alternative? · 4 Plan the stay · 5 Compare hotels · 6 Check the hotel · 7 Book the stay · 8 Before arrival · 9 During the stay · 10 After the stay

4 · Question topics — 32

Hotel or alternative · Neighbourhood and location · Hotel type · Room type and beds · Family rooms and cots · Accessibility · Quiet rooms and noise · Breakfast and dining · Dietary needs · Pets · Work and Wi-Fi · Check-in and check-out times · Luggage storage · Airport and station transfers · Parking · Price and taxes · What is included · Reviews and trust · Photos versus reality · Booking direct or platform · Loyalty and member rates · Payment and deposits · Cancellation and refunds · Dates and availability · Safety and security · Amenities · Cleanliness · Sustainability and local impact · Special occasions · Room problems and overbooking · Late arrival · Complaints and feedback

Why these four

City tests local relevance. **Traveller group** tests real needs. **Journey step** shows where intent turns commercial. **Topic** links each gap to a page or proof to publish.

Anatomy of a prompt

Each prompt reads like a real traveller's message. The situation and length of stay change the answer, which is exactly what generic prompts miss.

PROMPT H0541-T1 · JOURNEY STEP 6 · TOPIC: ACCESSIBILITY

How can we verify a **Paris hotel's accessibility claims?** **One of us uses a wheelchair.** **We are staying a week.**

- **Question template** (from the topic)
- **Traveller situation** (1 of 18 groups)
- **City** (1 of 20)
- **Length of stay** (varies the need)

Follow-up 1: Which measurements or photos should we ask for?
Follow-up 2: What is the backup plan if the room is not accessible?

Prompt ID	Traveller	Step	Main question
H0607-T1	Light sleeper	2 Choose an area	Which areas of Amsterdam are quiet at night but still convenient? One of us is a very light sleeper. We are staying three nights.
H1541-T1	Limited budget	5 Compare hotels	How can we compare total hotel prices in Paris, including taxes and fees? We want to keep the total cost down. We are staying a week.
H0964-T1	Travelling with a pet	7 Book the stay	What should we confirm about pets before booking a hotel in London? We are travelling with a small dog. We are staying three nights.

Why follow-ups

AI tools answer a conversation, not a single query. Follow-ups test whether the brand appears **when the traveller digs deeper**, where comparison and booking decisions happen.

Why neutral wording

No prompt names a brand or asks for "the best company". That keeps the test fair: a brand appears only if the AI **chooses** to mention or cite it.

04 · Method

Sampling and running the prompts

Running 3,000 prompts in four AI tools by hand isn't realistic. The study runs in **waves**, starting with a balanced sample.

Wave 1 prompts per journey step (64 total)



Wave 1: 64 prompts

Every one of the **32 topics twice**.
All **10 journey steps**, all **18 traveller groups**, all **20 cities**.
Chosen so no group or city dominates, so gaps can be attributed fairly.

Later waves

Add prompts where Wave 1 shows a gap, to confirm it, then re-run the same prompts **4 to 8 weeks** after fixes go live.

Run protocol	Rule
Tools	ChatGPT, Google AI Mode, Gemini, Perplexity
Session	New chat each time, **logged out or memory off**, same location
Sequence	Main question, then Follow-up 1 in the same chat
Record	Date, tool, model, full answer link, brands named, domains cited
Repeats	AI answers vary. Run key prompts **twice** and record both

05 · Method

Scoring every answer

Each run is scored on three things. Separating them is the most important design choice in the study.

1

Mentioned

Does the answer **name** the brand in its text? Yes or no.

2

Cited

Does a link to the brand's **website** appear in the sources? Yes or no, plus the URL.

3

Answer type

Does the answer **name any brands**, give **advice only**, or refuse?

Metric	Formula	What it shows
Mention rate	Runs mentioning the brand ÷ all runs	Overall presence
Citation rate	Runs citing the brand's site ÷ all runs	Whether AI **uses the brand's content** as evidence
Mention rate in brand answers	Mentions ÷ runs where the AI named any brand	The fairest measure of **competitive** visibility
Gap	3 or more runs, never mentioned or cited	Where to act first

Why answer type matters

Many travel questions get **advice-only** answers. A brand can't be mentioned when no brand is, so a low mention rate there isn't a loss. But it can still be **cited as the source** of the advice. That is often the first realistic win.

06 · From data to action

Reading the gaps

Results roll up into a dashboard by journey step, funnel stage, AI tool, topic, traveller group and city. Each gap points to a different fix.

Journey step	Runs	Mention rate	Citation rate	Mention rate in brand answers	Flag
2 Choose an area	n	%	%	%	GAP / blank
3 Hotel or alternative?	n	%	%	%	GAP / blank
5 Compare hotels	n	%	%	%	GAP / blank
6 Check the hotel	n	%	%	%	GAP / blank
7 Book the stay	n	%	%	%	GAP / blank

Structure of the dashboard. The same table repeats by funnel stage, AI tool, topic, traveller group and city.

Gap pattern	Likely cause	What to do
Missing at compare and book steps	No clear comparison or booking pages AI can quote	Comparison pages, transparent pricing, cancellation and inclusions pages
Missing for one traveller group	No content or proof for that need	Accessibility, family, pet or quiet-room details, with photos and measurements
Advice answers with no citation	Content exists but isn't liftable	Answer-first sections, clear dated facts, FAQs
Competitors cited from forums and reviews	Off-site reputation gap	Review programmes, community answers, PR
Missing in one AI tool only	Different source preferences	Check that tool's top cited domains and fill them

07 · Status

Where the study is now

Built

3,000 prompts across four dimensions, each with two follow-ups.
A **tracker workbook**: results log, automatic lookups, gap dashboard with flags.
A balanced **64-prompt first wave**.

Next

Run **Wave 1** in four AI tools.
Score the whole category, not one brand: **hotel chains, independent hotels, booking platforms and rental sites**.
Publish the findings as the next **AI Shelf** study.

Reusable for	Swap these dimensions
Private and guided tours	The original Travel Prompt Library (3,000 prompts)
Restaurants and food	Cuisines, dining occasions, dietary needs
Local services	Service types, customer situations, decision steps
E-commerce	Product categories, buyer personas, purchase steps

The principle

AI visibility isn't one number. **It's a map of who you show up for, and when.** A prompt set built from real guest situations makes that map visible, and tells you what to build next.

Method and prompt library by **Ameneh Saeednia**, seogirl.ca. Prepared with AI assistance. No results are reported in this document; findings will follow in a separate study.