

HYB26

Technology Edition

AI Everywhere



The Hotel Yearbook

Foresight and innovation in the global hotel industry

HY8



Beyond Bias: AI and the Reconstruction of Perception

Post-Stay

Rita Jusztina Varga
CEO, RAIZUP



Rita Juszina Varga raises an uncomfortable question the industry is only beginning to grapple with: when AI summarizes guest reviews at scale, whose experience actually gets represented? Drawing on her own decade-old stay as a case study, she argues that AI bias in review aggregation systematically amplifies certain voices while quietly erasing others — and that the hospitality industry urgently needs to rethink how it collects, segments, and trusts feedback data.

Guests really start to form their memories of a stay after they leave, and we are not the ones shaping those memories.

Ten years ago, I went to train one of my boutique hotel partners in Switzerland, where I was greeted with a mug that read: "Come as a guest, leave as a friend." That statement set the tone for my whole stay, and they lived up to it, consistently meeting the expectation they'd set. A decade on, it's a business trip I still remember, and I became a "friend for life" — that was my review. On paper, I served as a data point: a solo female business traveler from Hungary, a short length of stay, a direct booking, and a short booking window. For a machine, that is a data point.

For years, I've learned about and advocated for the business case of reviews: how they shape the guest experience cycle, operations, staffing, and revenue. Now we are all trying to navigate how to shape all the above in this new era.

We all know 80%+ of guests read reviews before booking (Tripadvisor), and hospitality businesses have caught up, with almost 90% using some sort of review system, directly or indirectly. So what happens in the business mix when we throw the learning machine in? My own story, like many others, is now read, processed, and summarised by AI during someone else's trip planning. The way we plan our journeys is formed by a system that mixes many feelings into a single, fictional average — a preconceived idea of a stay that no guest actually had.

Today, influencer and peer reviews shape travel decisions faster than marketers can keep up. More and more, a machine now stands between what guests say and what future travellers see. According to a recent Phocuswright study, one-third of U.S. travellers use AI, and the most common use is to analyse reviews and ratings.

Let that sit for a moment: it means people get their information from a machine before reading the actual review. How many reviews have now passed through this lens, leaving us with a version no one really experienced? And how relevant is that average to any specific individual's needs?

UNDERSTANDING THE AI BIAS OF REVIEWS

AI bias is created when a certain model analyses hotel reviews and produces results that are systematically skewed, favouring or disadvantaging certain hotels or guests purely because of how the AI was built and its sources.

AI doesn't understand reviews the way a person does because it learns certain data patterns it was trained on. So, if it carries a certain slant, AI reproduces and amplifies that angle at scale. Exactly that scale is the differentiator from one biased human.

It comes from three different places.

Data: Reviews skew toward able-bodied, straight, Western, native English-speaking guests, toward very happy or very angry guests, and toward hotels or hotel chains with a large number of reviews. AI treats this input as if it were the sole truth.

Labeling: Humans, when they tag a review as positive or negative to train the model, bring their own views and biases, which then shape the outcome. They might be misreading a cultural 5-star guest who gave an "unenthusiastic" review.

Design and aggregation: The lack of segmentation can be the source of it all. The choice to average everything into a single source means that a hotel that is great for couples but not suited for families — and other factors that have their own weight such as recency and volume — decides who wins the battle.

This leaves us at a place where AI can look objective and neutral, while in reality it can disadvantage a boutique hotel like my hotel partner in Switzerland, draw out an accessibility complaint, or anchor a guest's expectations with a clean summary they trust more than they should.

The question we should all ask ourselves is: which voices get amplified, and which are left out in the process?

We can create and use tools that produce a single, average truth, favouring the loudest and most fluent voices. Or we can choose tools that preserve differences, treating quiet and foreign-language reviews as important, and showing us what our real, diverse guests experienced — not just what an average guest might have thought.

I understand the appeal of simplification. A short summary is easy, less stressful, and quicker to act on. But real hospitality is about people, not the average. True hospitality notices what makes each guest unique, not just a face in a crowd.

We are not just hosts anymore; we help shape how people remember and talk about their stay. The systems that now store and share these memories do not always understand our environment or what matters most.

DIFFERENT PERSPECTIVES FOR REVERSE ENGINEERING

Let's start with who is at the core of both the stay and the reviews: the guests. Consider their motivation to review (or not), their cultural rating norms, their specific needs (accessibility, purpose of travel, etc.), and the fact that they read other AI-summarised reviews before booking.

Ask yourself: does our feedback process and public review profile actually serve the variety of guests we get, or just the loudest ones?

Most bias starts at collection time. Hotels often tend to ask only obviously happy guests to review. To make a difference, collect reviews from every guest. The better and more segmented the dataset serving you, the more any AI can get a truer picture.

Create specific questions that identify the segments, background, and further needs that help you understand the context. Capture structured context to reduce the "average" guest, include two or three targeted prompts, keep an open free-text prompt for anything unprompted, and store feedback with its own segment tags so patterns are visible later — like reverse prompting.

The second often-overlooked perspective: the frontline staff. They generate most of what gets reviewed, and they are the ones who must respond and probably fix the root causes. Consider their workload and their morale, because public criticism affects them personally, and they quite often know why a problem is recurring before any data points it out. Bias avoidance fails if staff are pressured to chase scores or cherry-pick happy guests. Management and ownership determine whether the company culture rewards genuine improvement or short-term score-grabbing, as they control the budget for tools and training.

You can close the loop and act, building advocacy and authenticity at the same time by responding to all reviews, setting service-level response times, stating specific fixes, reviewing feedback, flagging, and reporting suspected fake reviews. Most of all, never incentivise, edit, or buy reviews.

Distribution channels, OTAs, GDSs, and metasearch platforms each have their own review systems, ranking algorithms, and biases. Anyone in the game knows that the same hotel can look different across channels, because what factors into recency, volume, verified stays, and AI summaries is outside the actual business's control. It is worth looking at how to manage reputation consistently across channels and understanding how each channel's algorithm treats you. You can't control channel reviews or their AI, but quality and consistency pay off. Respond on every channel, not just your favourite, push for verified-stay reviews where you can, and keep property content consistent across channels.

All three perspectives lead back to the start. Once the guest leaves the property, it is up to them how they share their experience, which is fed back to other bookers. What you can do is map these perspectives against each other — and act accordingly.

