


Hotel Video InsightHub

Website User Manual

A practical guide to upload, prediction, multimodal diagnosis, and action

For hotel content, brand, and social-media teams; covers both standard-user and administrator workflows.



Current home and sign-in interface (on-screen wording may vary by deployment)

Version	v1.1
Date	3 August 2026
Coverage	Website front end, prediction flow, analysis reports, and administration

GET STARTED

Start here

The 30-second workflow

Sign in → upload the MP4 and content context → let the system complete the analysis → read the Local/Global direction → review AI Diagnosis and P0–P2 actions → edit and resubmit for comparison.

01 Prepare	02 Submit	03 Interpret	04 Improve
----------------------	---------------------	------------------------	----------------------

Roles and page access

Role	Main pages	Permission summary
Standard user (USER)	Home, Overview, Data Upload, History, Profile	Create analysis tasks and view only their own history and reports.
Administrator (ADMIN)	Home, Overview, Data Upload, History, Admin	Use analysis functions and review/delete users and submissions; view summary analytics.

How to use this guide

- First visit: follow Sections 1–7 to complete one upload and interpret its report.
- Need a quick read: go to Section 4, “Read the Score Snapshot,” then Section 5, “Read AI Diagnosis.”
- Ready to revise: use Section 7 and Section 10 to turn the findings into a controlled next version.
- Access or processing problem: use Section 11, “Accounts, Admin, and Troubleshooting.”

Companion videos

The project’s video-guide/ folder contains Chinese and English walkthrough videos plus SRT captions. Use the videos for a quick demonstration and this manual for field definitions, interpretation, and limitations.

1 | ACCESS

Sign in and enter the workspace




Figure 1. Home provides User/Admin sign-in and a language switch

1. Select User Login or Admin Login. This choice sets the landing page only; the account's assigned role controls actual permissions.
2. Enter the credentials issued by your administrator or organization, select Sign in, then Go to Workspace.
3. Use the top-right 中文 or EN control to switch the interface. The button shows the language you can switch to.

Request access

In the current build, Request access opens a demo-account form for Email, Hotel name, Employee ID, Password, Confirm password, and a five-character Captcha. Whether self-registration is allowed in production depends on your administrator's policy.

Language and generated text

The Qwen diagnosis automatically uses Chinese or English according to the primary language of the submitted title and text content; tags are used only when the title and text content do not provide enough language evidence. Switching the interface language does not translate an existing diagnosis. The top-right language control changes the website interface only.

2 | CREATE A TASK

Prepare and upload the media




Figure 2. Top of Data Upload: video, cover, title, and text content

Pre-submission checklist

Item	Requirement	Practical guidance
Video	Required; MP4	Use the actual intended edit, not a proxy file or silent rough cut.
Cover image	Optional; JPG/PNG	Upload the real candidate cover. Automatic frame extraction may not match the publishing-platform cover.
Title	Required; Chinese or English	Use the real post title; placeholder copy weakens consistency analysis.
Text content	Optional	Paste the intended caption/description so text–audio and text–video signals are meaningful.

Data responsibility

Upload the actual candidate cover whenever possible. If no cover is provided, the available preview image depends on the current deployment and may not match the publishing-platform cover.

2 | CREATE A TASK

Complete content and account context

The screenshot shows a dark-themed form for data upload. At the top, there are two placeholder boxes for video thumbnails. Below them are the following fields:

- Title:** A text input field with the placeholder text "Enter a Chinese or English title".
- Text content:** A large text area with the placeholder text "Enter video description...".
- Tags:** A row of three text input fields labeled "Tag 1", "Tag 2", and "Tag 3", followed by a "+" button to add more tags.
- Account follower count:** A text input field.
- Account following count:** A text input field.
- Account total likes and bookmarks:** A text input field.
- Hotel province / region:** A dropdown menu with the placeholder text "Select province or region".
- Submit:** A large blue button at the bottom of the form.

Figure 3. Bottom of Data Upload: tags, account counts, and hotel region

Field	What to enter	Why it matters
Tags	Use topics directly supported by the content; empty entries are ignored.	Supports semantic checks among title, tags, and visual content.
Follower count	A non-negative integer captured at submission time.	Adds source reach and social-proof context.
Following count	A non-negative integer; follow the interface label.	Adds account-activity context; it is not a subscriber field.
Total likes & bookmarks	The account's cumulative total, not a single-post value.	A visible source-popularity cue in the research model.
Province / region	Required; select from the list.	Adds deployment-specific geographic context and analytics.

4. Confirm that the video, title, and region are complete.
5. Select Submit. The page will prepare the submission, upload the video and any provided cover image, and create the task.
6. A successful submission opens the prediction-status page automatically. Do not submit the same task again.

Accuracy is better than completeness

If account counts, tags, or caption text are unknown, leave optional fields empty rather than inventing values. Incorrect context can distort the benchmark and diagnosis.

3 | TRACK A TASK

Wait for analysis and reopen reports from History

PENDING Queued	PROCESSING Analysing	COMPLETED Ready	FAILED Failed
-------------------	-------------------------	--------------------	------------------

The prediction page polls for status about every two seconds. When processing is complete, it redirects to the analysis report automatically. The two score cards on the transition page may appear only briefly; you do not need to remain there.

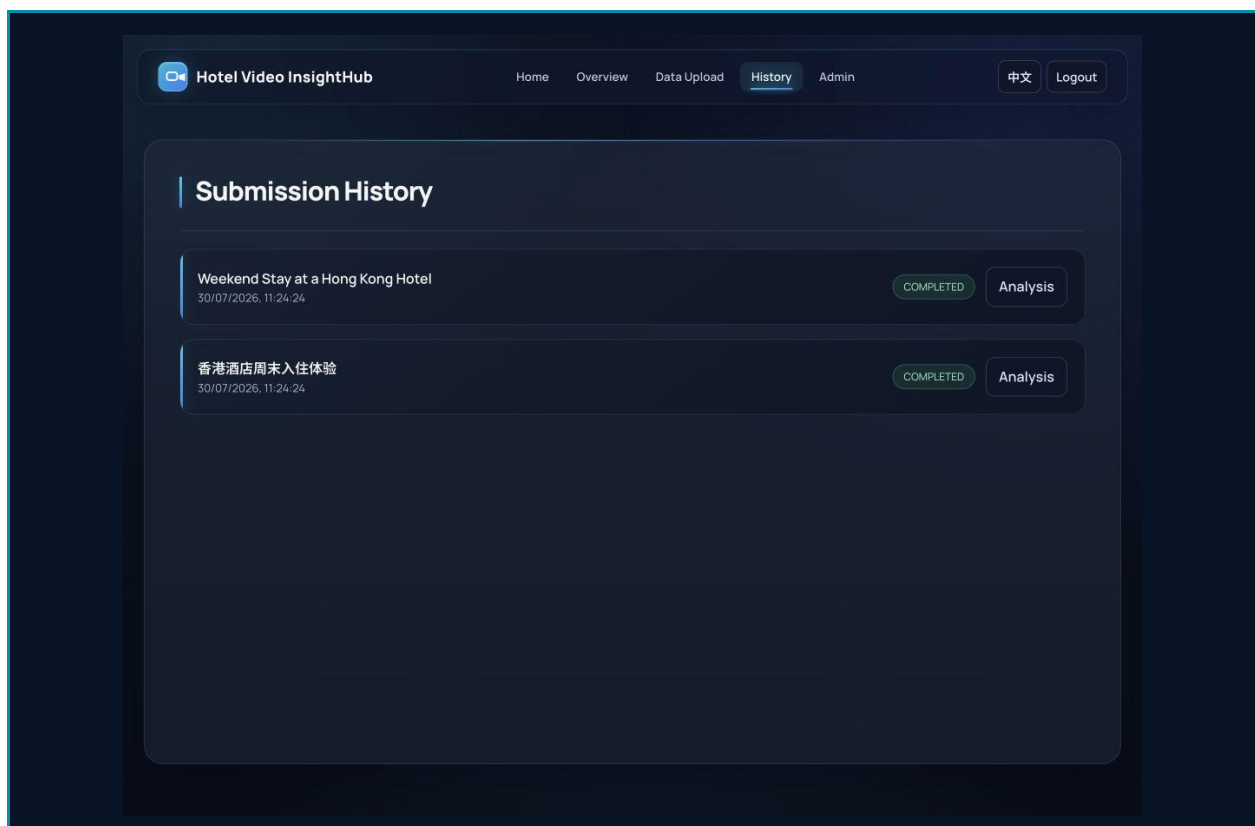


Figure 4. History is newest-first; Analysis is available only for completed tasks

Use History correctly

- A standard user sees only their own submissions; administrators have a separate all-submission view.
- For COMPLETED tasks, select Analysis to reopen the report.
- For PENDING or PROCESSING tasks, wait and refresh rather than creating duplicates.
- For FAILED tasks, check the MP4, title, and region, then create a new submission from Data Upload.

4 | INTERPRET RESULTS

Read the Score Snapshot first

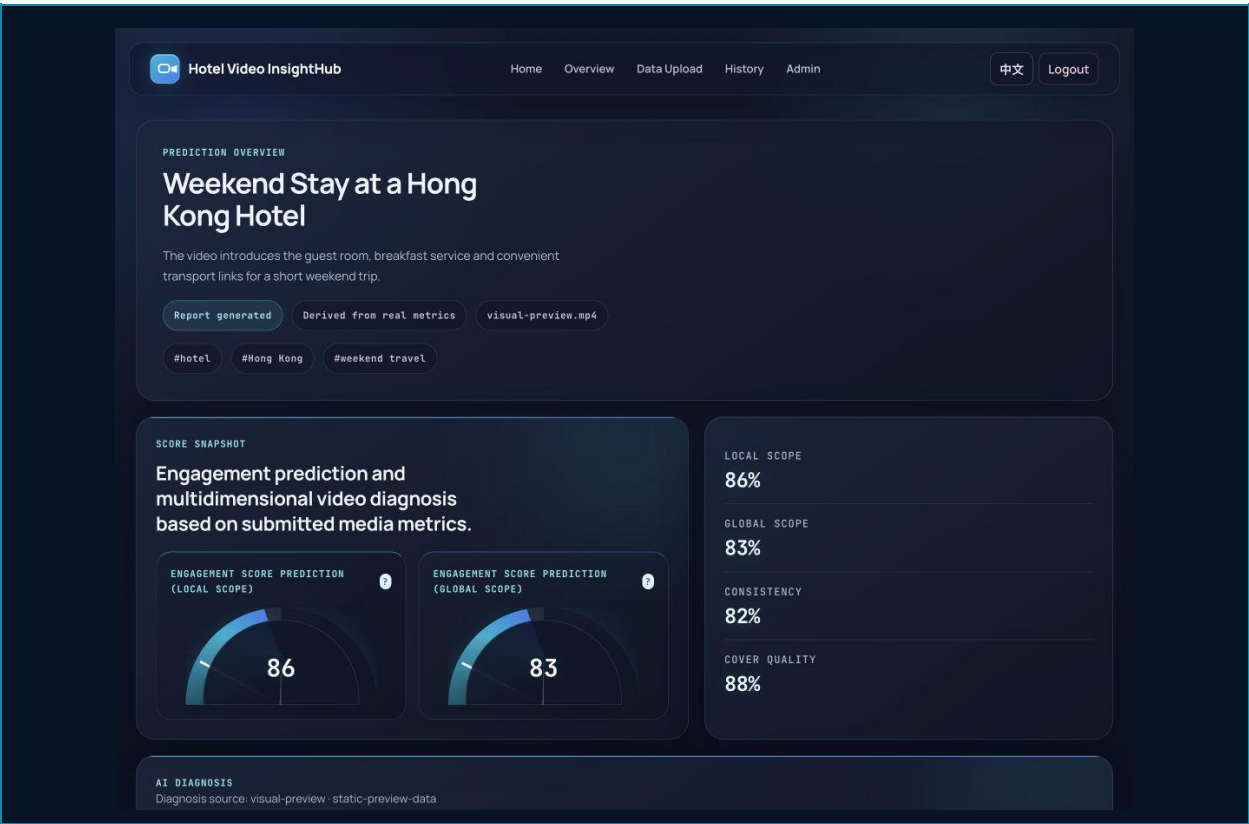


Figure 5. Report header: task context, Local/Global scores, and key snapshot metrics

On screen	Meaning	How to interpret it
Local scope	Account-context benchmark	Relative to the configured account-history baseline. The preview build demonstrates Hotel ICON history; confirm the production baseline with your administrator.
Global scope	Cross-account benchmark	Relative to the configured luxury-hotel social-media benchmark; use it as a broader competitiveness signal.
Consistency	Title–video consistency	The snapshot uses this one metric, not an average of all consistency metrics. See Underlying Metrics for all six.
Cover quality	Cover feature score	Reflects technical and visual properties of the current cover. It is not click-through rate and does not guarantee performance.

A score is a direction, not traffic

The 0–100 display is a deployment-specific relative prediction/feature scale. It is not a forecast of views, media return, bookings, guaranteed exposure, or a statistical confidence interval for one upload.

4 | INTERPRET RESULTS

Read Local and Global together

The system does not define one pass/fail threshold for every account. Compare within the same deployment, against recent versions and similar content types, then use the matrix below.

Relative pattern	Possible signal	Recommended next move
Local higher + Global higher	Good fit with the account and broader competitive potential.	Preserve the core promise; test small cover, opening, or pacing variants.
Local higher + Global lower	Improved against account history, but limited cross-account distinctiveness.	Strengthen the hotel’s unique value, recognizable setting, and cover focal point.
Local lower + Global higher	The concept may travel, but it is not fully aligned with this account’s voice or audience.	Keep the concept; adapt title, tone, human presence, or publishing context.
Local lower + Global lower	The content promise, visual appeal, or multimodal consistency may all need work.	Fix the P0 issue before minor tag edits; then resubmit a controlled revision.

Recommended reading order

01	02	03	04	05
Score direction	Core problem	Dimension evidence	P0–P2 actions	Base metrics

- Verify the title, filename, tags, and status so you are reading the intended version.
- Use relative score direction, not a one-off absolute value, as the starting point.
- Treat generated diagnosis as a testable hypothesis and check the evidence below it.
- Keep account-context fields constant when resubmitting a revised edit.

Local preview reports

Fixed reports such as visual-preview-cn and visual-preview-en are for interface review only; they are not production inference results.

5 | INTERPRET RESULTS

Read AI Diagnosis

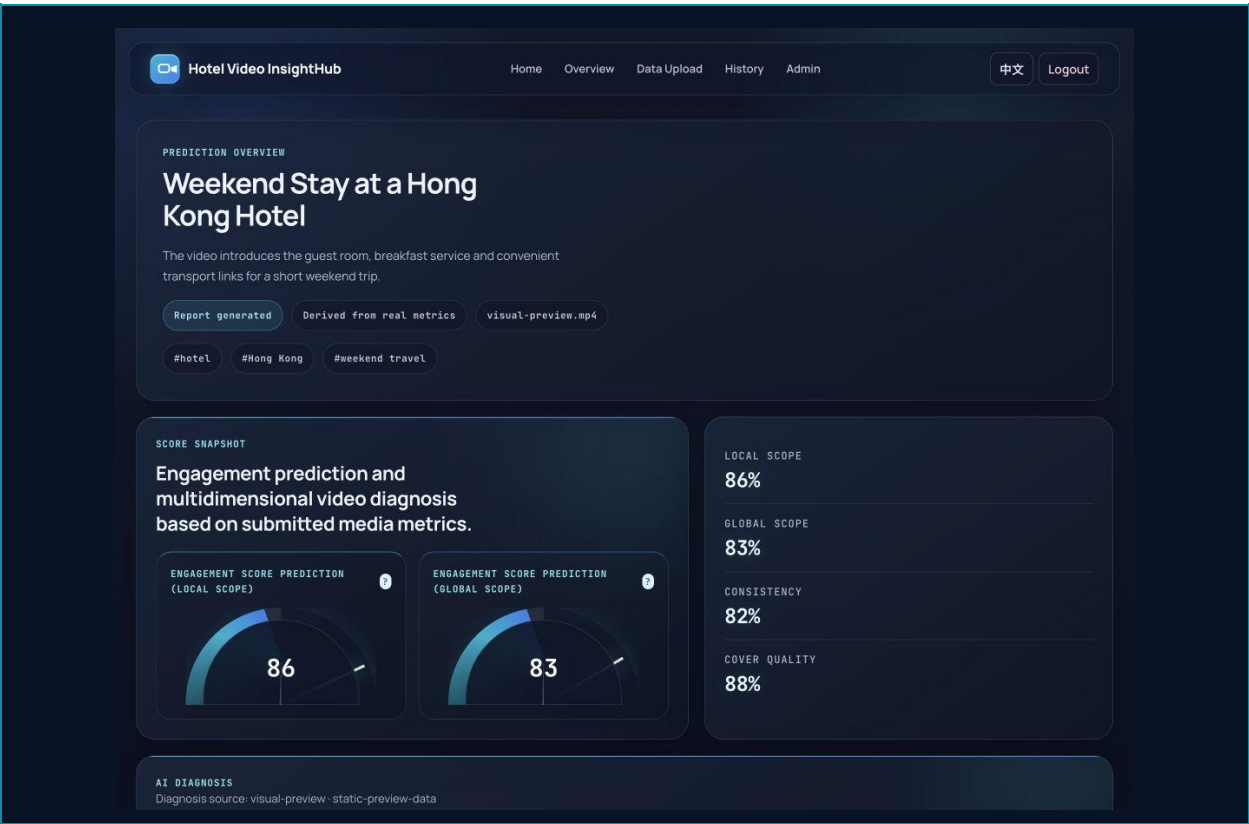


Figure 6. AI Diagnosis includes a summary, Core Problem, Audience Insight, Strengths, and Risks

Module	Question it answers	How to use it
Summary	What does this version communicate overall?	Confirm that the system’s reading matches the creative intent.
Core Problem	What is the highest-priority friction?	Treat it as the first testable hypothesis; pair it with P0 and a base metric.
Audience Insight	Who might respond, and why?	Use as a positioning hypothesis, not measured demographics or personal attributes.
Strengths / Risks	What should stay, and what may hold performance back?	Preserve the strengths while changing weak variables to maintain comparability.
Expected ROI / uplift	What direction might a change unlock?	Treat as a scenario estimate, not causal evidence, guaranteed ROI, or promised uplift.

Missing Qwen diagnosis text does not necessarily mean task failure

If the Qwen diagnosis service is disabled, temporarily unavailable, times out, or returns output that does not pass validation, the base analysis task may still complete. The report may use a locally derived fallback diagnosis instead. Use the task status and Underlying Metrics to assess the result, and contact us if the issue persists.

6 | INTERPRET RESULTS

Use the three dimensions to locate the issue

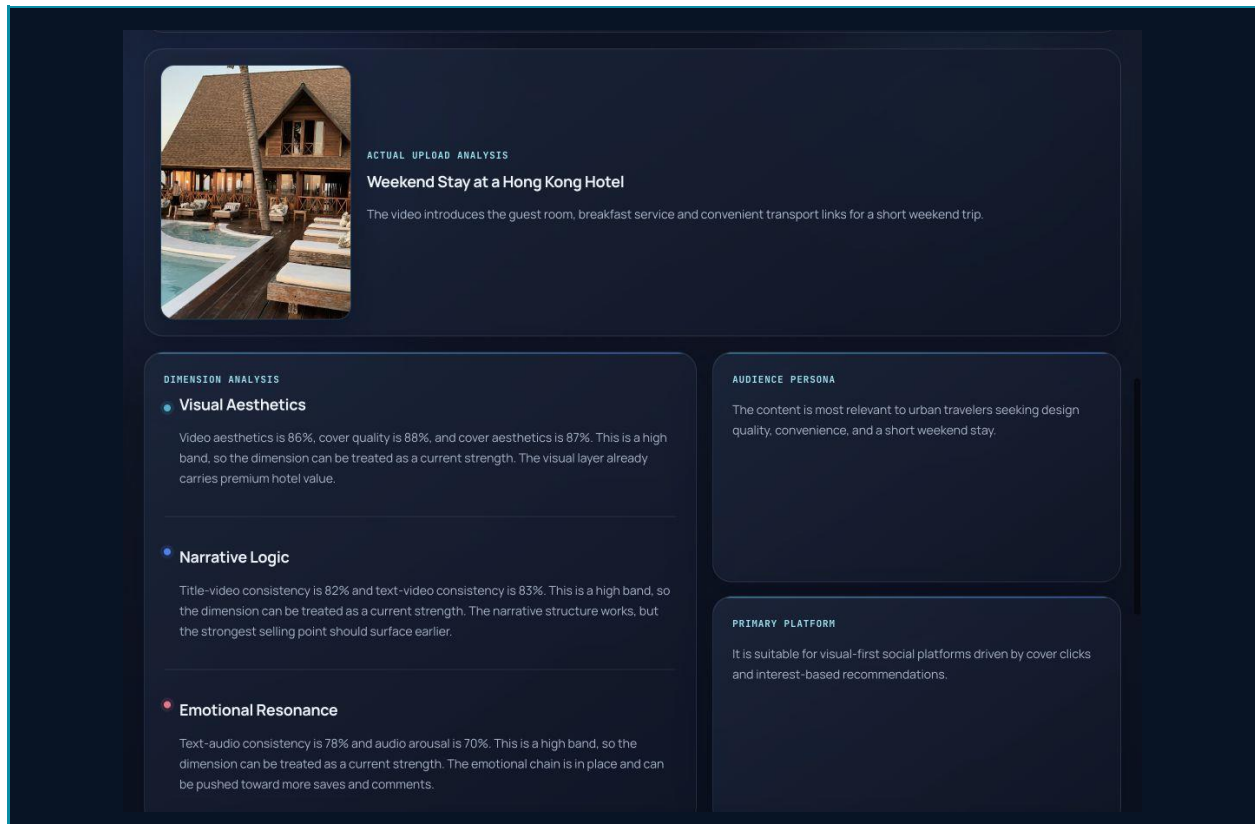


Figure 7. Dimension Analysis, Audience Persona, and Primary Platform

Dimension	Main evidence	Typical edit
Visual Aesthetics	Video aesthetics, cover quality/aesthetics, composition, and visual appeal.	Change focal point, light, clarity, human/space balance, or opening shot.
Narrative Logic	Title-video, text-video, title-tags, and cross-modal consistency.	Align the title promise with footage, remove unsupported tags, and reorder shots around one message.
Emotional Resonance	Text/audio sentiment and arousal, voice presence, text-audio, and video-audio alignment.	Adjust narration, ambient sound, music intensity, rhythm, and emotional direction.

Audience and platform are strategy hypotheses

Audience Persona and Primary Platform are inferred from content features. They are not audience research, media-testing results, or person-level tracking. Validate them against brand strategy, campaign data, and local regulation.

7 | TAKE ACTION

Edit in P0 → P1 → P2 order, then verify with metrics



Figure 8. Action Plan cards and Underlying Metrics tabs

Priority	Meaning	Execution rule
P0	The core issue most likely to block comprehension or engagement	Complete first; change only one or two core variables to preserve comparability.
P1	A second-layer optimization that can amplify the concept	After P0, refine title, pace, sound, or a cover variant.
P2	Incremental polish and long-tail adaptation	Finish with tags, detail copy, platform adaptation, and test variants.

Four display conventions in Underlying Metrics

Tab	Display	Important note
Content quality	One decimal for four scores; Yes/No for voice; Likely/Unknown for face	The first four are not shown as percentages; compare within the same build and deployment.
Sentiment signals	Percent (%)	Signals are model estimates, not a guest-sentiment survey.
Consistency	Percent (%)	High alignment does not guarantee engagement; the content still needs value and distinction.
Eastern aesthetics	0.00–1.00	Do not add or directly compare these values with 0–100 scores.

8 | METRIC DICTIONARY

What the base metrics can tell you

Group	Metric	Question it can answer
Content quality	Video aesthetics / Cover quality / Cover aesthetics	Does the video or cover show relatively clear, appealing, and usable visual properties?
Content quality	Text readability	Is the title/caption easy to process? It does not verify truth or brand compliance.
Content quality	Human voice	Was recognizable human speech detected? It does not assess delivery, meaning, or audio quality.
Content quality	Human face	Was a human face detected in the uploaded cover image? Unknown means the input or analysis was insufficient for a reliable result.
Sentiment signals	Title / Text / Audio sentiment	Do emotional directions across modalities support or contradict one another?
Sentiment signals	Text / Audio arousal	Does content energy match the intended brand and campaign purpose?
Consistency	Title–tags / Title–cover / Title–video	Does the preview promise match the visible evidence?
Consistency	Text–audio / Text–video / Video–audio	Do copy, footage, and sound support the same experience during viewing?
Eastern aesthetics	Richness, Harmony, Adaption, Modern, Eastern, Western	How are visual-style signals distributed? Use for creative comparison, not cultural-value judgement.

Presence is not quality

Voice and face indicators answer whether something may be present, not whether it feels natural, credible, well produced, or appropriate. Likewise, high consistency means greater coordination, not guaranteed originality or persuasion.

If the snapshot and base metrics seem inconsistent

- Confirm which metric the snapshot uses: the current Consistency snapshot is Title–video consistency.
- Check whether the upload included a caption, cover, and usable audio; missing modalities narrow interpretation.
- Use the specific base metric to guide the edit rather than trying to maximize every number.

9 | RESEARCH INTERPRETATION

Engagement is not one behaviour

The website currently summarizes direction through Local and Global scores. The VMNNA research manuscript instead models likes, shares, bookmarks, and comments separately. Use the distinctions below to define an objective, but do not treat the research feature importance as a causal explanation of an individual report.

Outcome	Behavioural meaning	Most sensitive cue in the study	Creative implication
Like	Low-effort endorsement	Cumulative account likes/bookmarks (source social proof)	Maintain recognizability and credibility; account popularity is not content quality.
Share	Propagation to others	Cover-image quality	Make the cover's focal point, value, and setting immediately legible.
Bookmark	Retention for later use	Cover-image quality	Surface practical, comparable, or trip-planning value.
Comment	Higher-effort social interaction	Human presence	If conversation is the goal, test natural people, service interactions, or a prompt.

Research basis: main manuscript, Table 5, p. 27. Behaviour labels are literature-informed, not directly observed psychological constructs.

Content type and modality mix

Content type	High-engagement route in the study	A testable creative direction
Surroundings / public areas / food	More often coordinated text + audio + visual signals	Build around one experiential promise and align footage, music/ambience, and copy.
Rooms / bathrooms	Visual signals were more prominent	Prioritize layout, materials, scale, clarity, and recognizable detail.

Cases are illustrations, not rules

An appendix case shows that aligned text and images can still be undermined by loud, noisy sound. Another food case combined coordinated copy and visuals with light, stable music. These are local explanations, not a fixed recipe for every hotel.

Research basis: main manuscript, Table 6, pp. 27–29; Appendix J, p. 25.

10 | TURN INSIGHT INTO AN EDIT

Revise and retest in a comparable way

01	02	03	04	05	06
Set objective	Choose P0	Change 1–2	Retest	Compare	Validate live

1. Choose the primary outcome: endorsement, sharing, saving, or conversation. Different goals need not optimize the same cues.
2. Pair the Core Problem with one base metric—for example, a promise gap with Title–video consistency.
3. Change only one or two controllable variables, such as cover + title or music + opening pace.
4. Resubmit with the same follower, following, and cumulative-engagement context, and name each version clearly.
5. Compare Local/Global direction, whether P0 changed, and the relevant evidence—not just the headline score.
6. After publishing, validate with impressions, completion, platform engagement, and business outcomes.

Minimum comparison log

Version	What changed	Local / Global	Core Problem	Observed platform result
A baseline	Baseline	___ / ___	_____	_____
B revision	_____	___ / ___	_____	_____
C revision	_____	___ / ___	_____	_____

Managerial judgement remains the final gate

Use the model for pre-publication audit, option comparison, and issue localization. Brand fit, legal compliance, production feasibility, media budget, and truthful guest experience still require human judgement.

11 | ACCOUNTS & SUPPORT

Profile, Admin, and troubleshooting

Page	What it provides	Important note
Profile (USER)	Account/role, submission totals, completed/processing/failed counts, and latest submission.	Session time is not a security-audit login timestamp; admins are redirected to Admin.
Admin (ADMIN)	View/delete users and submissions; view counts, locations, and activity summaries when available.	Deleting a submission also removes uploaded files. The current build may not display a success message after confirmation.

Common issues

Symptom	Check first	Resolution
Cannot sign in	Credentials, account role, network	Re-enter the details. The login card cannot change the assigned role; contact an administrator if needed.

Symptom	Check first	Resolution
Submission fails	MP4, title, region	Complete required fields, reselect the file, and submit again.
Processing takes too long	Existing duplicate task, service status	Wait and refresh History; avoid duplicates and contact the administrator if it persists.
History is empty	Signed-in account	A standard user sees only their own tasks; confirm you did not switch accounts.
No Qwen diagnosis text	COMPLETED status, diagnosis source	Review Underlying Metrics first. The base task may still complete when the Qwen service is unavailable or its output fails validation; a local fallback diagnosis may appear. Contact the administrator if the issue persists.
Human face is Unknown	Cover and heuristic limitation	Do not read this as “no face”; Check whether people or faces are too small, blurred, obscured, or absent from the relevant frames.
Missing Admin or Profile	Assigned role	ADMIN opens Admin; USER opens Profile.

12 | RESPONSIBLE USE

What the results do—and do not—mean

Appropriate uses	Not appropriate
Compare versions from the same account and similar content category.	Guarantee views, engagement, bookings, or revenue.
Locate weak cover, title, copy, audio, visual, or cross-modal coordination.	Prove that one feature caused engagement to rise or fall.
Create a shared pre-publication review language for the team.	Infer an individual viewer’s identity, travel stage, booking intent, or mental state.
Prioritize the next creative change.	Replace brand, legal, media-buying, or actual platform-performance judgement.

Research scope and limitations

- The main study used 4,045 public RedNote short videos from high-end hotels in Mainland China and Hong Kong. A robustness dataset used 2,340 public Facebook videos from US hotel brands. Cue weights may shift across platforms, cultures, and time.
- Distribution algorithms, impression volume, and paid promotion affect observed engagement. Complete impressions were unavailable, so the output is not a causal engagement rate under equal exposure.
- Cover quality, sentiment, readability, human presence, and source popularity are computational proxies. They improve interpretability but are not causally validated management rules.
- MAE, MSE, and RMSE are research evaluation statistics—not confidence scores for a website report. The appendix’s median split for high/low engagement is not a universal website threshold.
- Production deployments should periodically test and recalibrate with recent platform, account, and campaign data.

Research basis and traceability

This guide draws on the user-supplied R2 manuscript and appendix, “Decoding customer engagement with hotel short videos: an interpretable multimodal neural network approach with fine-grained user attention.” Operational instructions reflect the website code and local-preview screens updated through 3 August 2026. Research interpretation primarily uses manuscript pp. 21–32 and Tables 5–6, plus Appendix C and E–J (pp. 10–25).

Final rule of thumb

Use the score as navigation, the base metrics as evidence, and live platform outcomes as validation. The goal is not one high number—it is a repeatable, comparable, and learnable content-improvement process.