

Hospitality-VQA: Decision-Oriented Informativeness Evaluation for Vision–Language Models

Jeongwoo Lee^{1,*}, Duhyeong Baek¹, Eungyeol Han¹, Soyeon Shin¹,
Gukin Han², Seungduk Kim², Jaehyun Jeon^{1,†}, Taewoo Jeong^{2,†}

¹{leejeongwoo9941, glzeng99, condense, shin020810, jaehyun.jeon}@yonsei.ac.kr

²{bryan.han, seungduk.kim, taewoo.jeong}@yanolja.com

¹Yonsei University, ²Yanolja NEXT

Abstract

Recent advances in Vision–Language Models (VLMs) have demonstrated impressive multi-modal understanding in general domains. However, their applicability to decision-oriented domains such as hospitality remains largely unexplored. In this work, we investigate how well VLMs can perform visual question answering (VQA) about hotel and facility images that are central to consumer decision-making. While many existing VQA benchmarks focus on factual correctness, they rarely capture what information users actually find useful. To address this, we first introduce *Informativeness* as a formal framework to quantify how much hospitality-relevant information an image–question pair provides. Guided by this framework, we construct a new hospitality-specific VQA dataset that covers various facility types, where questions are specifically designed to reflect key user information needs. Using this benchmark, we conduct experiments with several state-of-the-art VLMs, revealing that VLMs are not intrinsically decision-aware—key visual signals remain underutilized, and reliable informativeness reasoning emerges only after modest domain-specific fine-tuning.

1 Introduction

Images play a central role in the hospitality industry, serving as the primary medium through which guests evaluate and compare accommodations (Zhang et al., 2022). When consumers choose where to stay, they often rely more on visual impressions—such as room layout, view, lighting, and cleanliness—than on textual descriptions. These images convey both factual and atmospheric cues that shape user decisions, making visual understanding a crucial aspect of hospitality intelligence (Cuesta-Valiño et al., 2023).

* Main contributor.

† Corresponding author.

Existing VQA Dataset

: General Questions



Q: What color is the slide?

A: Orange

Ours: Hospitality-VQA

: Decision-Oriented Questions



Q1: What **type** of this facility is this space?

A1: Room Interior

Q2: What **activities** do this space support?

A2 : Sleeping, Sitting

Figure 1: Comparison between general VQA (top) and decision-oriented Hospitality-VQA (bottom).

Previous studies in the hospitality domain have predominantly relied on text-based analytics of online reviews to model customer satisfaction, preferences, and demand patterns (Li et al., 2013; Xiang et al., 2015). In parallel, a growing body of work has examined visual information in accommodation images by extracting predefined or low-level features—such as aesthetics, composition, or object categories—and relating them to outcomes such as booking decisions, user intentions, or perceived accommodation quality (Ren et al., 2021; He et al., 2023). More recently, while some studies have leveraged Large Language Models (LLMs)

for hospitality analysis, they remain primarily focused on textual inputs such as reviews or descriptions (Guidotti et al., 2025). Despite these advancements, existing methods—whether text-centric or feature-based—remain limited in their ability to perform integrated multimodal reasoning. Specifically, they often fail to capture the interplay between higher-level spatial organization and functional semantics in images, factors that are central to how humans evaluate hospitality environments.

Meanwhile, recent advances in Vision-Language Models (VLMs) have significantly improved multimodal reasoning across general domains. Modern models (Comanici et al., 2025; Hurst et al., 2024; Bai et al., 2025) can generate contextualized image descriptions and answer open-ended questions that go beyond traditional visual recognition, suggesting strong potential for domain-specific applications. While these models have demonstrated promising results in specialized fields such as e-commerce (Trabelsi et al., 2025) and medical imaging (Tu et al., 2024), their use in the hospitality domain has been relatively limited with respect to decision-oriented evaluation settings.

When examining these domain-specific applications, one important insight emerges: the performance of VLMs often depends on *how information needs are framed*. Generic prompts (e.g., "What is in this image?") yield vague descriptions that are insufficient for hospitality purposes. As illustrated in Figure 1, appearance-level questions alone provide limited insight into whether a space meaningfully supports guest activities or experiences. Meaningful evaluation requires domain-specific questions that elicit decision-relevant insights—*not just whether a room contains furniture, but how its layout supports guest activities; not merely whether a window exists, but what type of view it provides*. This raises a key design challenge: how to formalize the kinds of visual evidence that actually support user decisions.

To address this challenge, we introduce **Hospitality Informativeness**, a domain-grounded framework that quantifies how much decision-relevant information a hospitality image-question pair provides. Because user information needs vary across facility types—such as spatial clarity in rooms, amenity completeness in bathrooms, or functional elements in shared facilities (Wakefield and Blodgett, 1996)—we first identify the facility type and design domain-specific questions accordingly. Although these needs appear diverse, we observe that

the visual cues influencing booking decisions consistently fall into a small set of structural, functional, and view-related dimensions. Building on this observation, we define four fundamental visual axes (spatial legibility, activity affordance, contextual openness, and geometric completeness). Together, these axes capture the dominant cues that shape user perception and decision-making in hospitality imagery, providing a principled basis for evaluating VLM responses (Greene et al., 2016). We use these axes to construct **Hospitality-VQA**, a new VQA benchmark aligned with decision-centric evaluation rather than generic scene description.

Our main contributions are:

- We formalize *Informativeness* in the hospitality domain as a set of four interpretable axes that capture decision-relevant visual cues in hotel and facility imagery.
- We build **Hospitality-VQA**, a VQA dataset whose questions and labels are derived from these axes and tailored to diverse facility types.
- We benchmark eight general-purpose VLMs and show that they struggle with fine-grained hospitality informativeness. Our dataset enables measurable performance gains through lightweight domain adaptation, highlighting its value as a foundation for future model development on hospitality domain.

2 Related Works

2.1 Visual Analysis in Hospitality

Research in hospitality AI has largely focused on structured prediction tasks such as room-type classification and price estimation using CNN-based frameworks, treating images as static inputs and overlooking richer semantic cues relevant to user assessment. In parallel, a growing line of work extracts computable visual descriptors—ranging from low-level color statistics to mid-level attributes such as aesthetics and composition—and relates them to outcomes like booking intentions or demand (Zhang et al., 2022; He et al., 2023; Cuesta-Valiño et al., 2023). However, these approaches are not designed to evaluate whether models can answer *decision-relevant questions* about an image. Recent work has also explored multimodal hotel retrieval and preference matching (Askari



Figure 2: Bad vs. Good examples for each informativeness dimension. Bad images lack decision-relevant visual cues—resulting in low spatial legibility, weak activity affordance, obstructed or unbalanced contextual openness, or incomplete geometric completeness. Good images exhibit high spatial legibility, clear activity affordances, well-balanced contextual openness, and strong geometric completeness, enabling more reliable assessment of hospitality informativeness.

et al., 2025), but focuses on similarity or relevance rather than explicit question answering and decision-oriented evaluation.

Existing approaches rarely model how guests simulate a potential stay experience from visual evidence. Although the presentation of accommodation photographs can sway selection behavior (Sánchez-Torres et al., 2024), the notion of visual utility—how visual elements convey functional and spatial affordances—remains underspecified. Consequently, evaluation typically centers on prediction accuracy or correlational signals rather than decision-oriented reasoning. Our work addresses this gap by shifting the focus to the systematic evaluation of decision-relevant information through a VQA benchmark grounded in *Hospitality Informativeness*.

2.2 Vision–Language Models and Domain Adaptation

Recent general-purpose VLMs, including GPT-4o (Hurst et al., 2024) and Gemini 2.5 Pro (Comanici et al., 2025), demonstrate impressive capabilities in image captioning and open-ended QA. However, these models are trained primarily on web-scale, caption-style data to describe “what exists,” often lacking the specialized reasoning required to evaluate “how useful it is” in a vertical domain. In hospitality, visual understanding goes beyond object recognition; it requires inferring spatial habitability and functional affordance. Since standard VLMs are not inherently optimized for

such evaluative reasoning, it remains unclear to what extent they can interpret the nuanced visual evidence essential for consumers—motivating the need for a domain-grounded benchmark.

2.3 From Factuality to Decision-Centric VQA

Standard VQA benchmarks (e.g., VQA v2 (Goyal et al., 2017), GQA (Hudson and Manning, 2019)) have driven progress in multimodal reasoning but primarily evaluate factual correctness or common-sense knowledge. While recent goal-oriented VQA tasks explore navigation or physical manipulation (Das et al., 2018; Gurari et al., 2018), they rarely address consumer-facing decisions in which the goal is to assess the suitability of a space or service. Existing benchmarks are not designed to measure whether an image provides the type of evidence needed to support informed accommodation choices (Cuesta-Valiño et al., 2023). Addressing this limitation, we introduce *Informativeness* as a metric to quantify the specific visual signals—such as layout clarity and functional completeness—that facilitate reliable accommodation assessment.

3 Quantifying Informativeness in Hospitality

We argue that true understanding in the hospitality domain requires quantifying the visual evidence that supports user decision-making. While general VQA benchmarks focus on factual correctness (e.g., “is there a window?”), hospitality users rely on images to envision their stay—judging layout,

Facility Type	SL	AA	CO	GC
Room Interior	•	•		
Indoor Facility	•	•		
Outdoor Facility		•	•	
Accommodation Exterior			•	•

Table 1: Facility types and applicable informativeness dimensions (SL: Spatial Legibility; AA: Activity Affordance; CO: Contextual Openness; GC: Geometric Completeness).

affordance, and atmosphere. Because these subjective assessments directly drive booking decisions, mere descriptions are insufficient (Cuesta-Valiño et al., 2023). To address this, we formalize *Informativeness* as a measurable metric. We propose that the vague notion of “a useful hotel image” can be decomposed into specific, quantifiable axes that act as proxies for the user’s envisioned stay experience (Greene et al., 2016).

3.1 Facility Taxonomy and Informativeness Dimensions

Hospitality imagery encompasses diverse scenes, ranging from critical facility views to irrelevant content. To structure our analysis, we categorize images into five main facility types: *Room Interior*, *Indoor Facility*, *Outdoor Facility*, *Accommodation Exterior*, and *Irrelevant*. To capture more specific functional contexts, images are additionally annotated with finer-grained sub-facility labels; the full taxonomy is provided in the Appendix A.

We define an image as *informative* if it provides quantifiable visual cues along four fundamental dimensions: **Spatial Legibility (SL)**, **Activity Affordance (AA)**, **Contextual Openness (CO)**, and **Geometric Completeness (GC)**. Figure 2 illustrates the characteristic visual patterns corresponding to each dimension. The applicability of these dimensions depends on the facility type, and Table 1 specifies which dimensions serve as valid evaluative criteria for each category.

Beyond these dimensions, *Room Interior* images are additionally annotated with two semantic attributes—view type and room type—to capture preferences not fully represented by geometric or functional cues. Conversely, the *Irrelevant* category contains images lacking decision-relevant visual evidence and is excluded from further evaluation.

Hospitality-VQA Annotation Schema
(1) Hierarchical Labels main: <i>Primary facility category</i> sub : <i>Fine-grained sub-facility</i>
(2) Informativeness Axes SL, AA, CO, GC <i>*Mapped based on Table 1</i>

Figure 3: The formal annotation schema used in Hospitality-VQA. We record hierarchical facility labels and quantify visual utility across the four informativeness dimensions.

3.2 Axis Definitions

We define the four axes as quantifiable prediction targets to measure visual utility.

Spatial Legibility. Defined as the count of distinct planar surfaces (floor, walls, ceiling), this metric serves as a proxy for **spatial comprehension**, distinguishing ambiguous close-ups from structural views that reveal room volume (Oliva and Torralba, 2001).

Activity Affordance. We quantify *meaningful components*—functional objects that explicitly afford guest activities (e.g., desks, seating, storage surfaces)—to capture the space’s **functional habitability** while filtering out purely decorative elements (Greene et al., 2016).

Contextual Openness. Measured as the ratio of non-facility elements (sky, nature, background structures), this metric assesses **contextual balance**, identifying overly occluded views or excessively distant shots that hinder environmental interpretation (Cuesta-Valiño et al., 2023).

Geometric Completeness. Approximating a building as a dominant cuboid, we assess the visibility of its three canonical faces—front, side, and roof—to evaluate **geometric integrity** and the perceptibility of its 3D form (Sánchez-Torres et al., 2024).

For *Room Interior* images, we supplement these structural axes with two semantic attributes—**View Type** and **Room Type**—which capture domain-specific preferences essential for booking decisions.

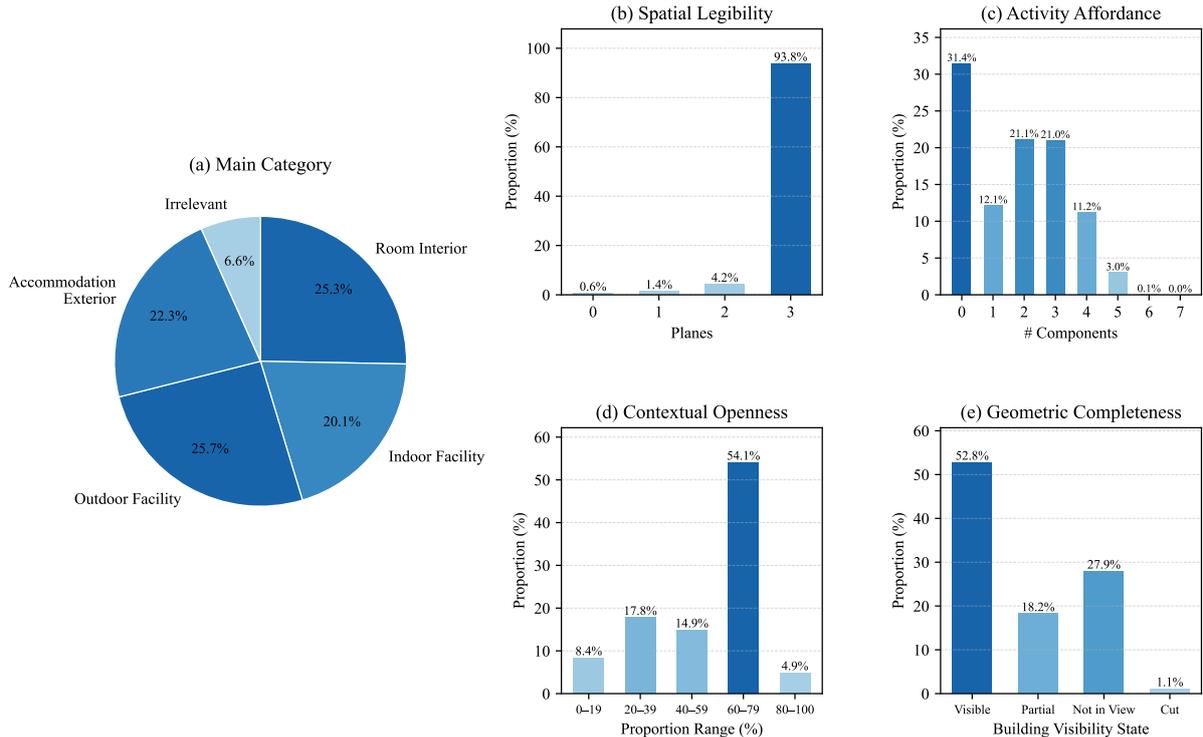


Figure 4: Dataset statistics of Hospitality-VQA. (a) Distribution of main facility categories. (b–e) Distributions of the four informativeness axes, reflecting characteristic properties of professionally curated hospitality listing images.

4 Hospitality-VQA Dataset

To translate the informativeness framework into a measurable benchmark, we introduce **Hospitality-VQA**. As existing VQA datasets lack the hospitality-specific imagery and annotations aligned with the four informativeness dimensions, they are ill-suited for evaluating decision-oriented visual reasoning. To address this gap, our dataset provides expert-annotated supervision explicitly tailored to these axes. The following subsections detail our pipeline for image collection, hierarchical annotation, and the derivation of instruction–answer pairs.

4.1 Data Collection

A total of 5,000 hospitality images were collected from `nol.yanolja.com` through random sampling of listing pages. Because the pool was not pre-filtered by facility type, categorization into facility types and relevance labels was performed during annotation (Section 3).

4.2 Data Annotation

Annotation was conducted by five annotators who were instructed in the Informativeness Framework defined in Section 3. Figure 3 summarizes the anno-

tation schema used throughout dataset construction. A pilot round was conducted prior to the main annotation phase to calibrate labeling practices and align annotators’ interpretations. All 5,000 images were then independently labeled by all annotators.

For quality control, we adopted a strict consensus protocol. Labels with high agreement (at least 4 out of 5 annotators)—covering 86.4% of all annotations—were accepted as ground truth. Cases with lower agreement were flagged and resolved through consensus discussion, producing finalized facility-type and axis annotations for every image.

For VLM assessment, each label is converted into a concise instruction–answer pair using fixed templates specifically defined for each facility type and axis. These templates are designed to ensure consistent and scalable evaluation by mapping classification targets into a VQA format while controlling variation in question phrasing. Only the templates applicable to an image’s facility type are instantiated, and example templates are shown in Appendix B.

4.3 Dataset Analysis

Across the 5,000 collected images, a total of 19,729 QA pairs are generated by applying the fixed templates to the facility-type and axis-level annotations.

Figure 4 summarizes the overall distributions of facility categories and informativeness annotations. As shown in Fig. 4a, the main facility categories are relatively well balanced, each accounting for roughly a quarter of the dataset.

Figures 4b–e report the distributions of the four informativeness dimensions. Spatial Legibility (Fig. 4b) and Activity Affordance (Fig. 4c) are summarized as discrete counts reflecting visible planes and meaningful components, respectively, while Contextual Openness (Fig. 4d) and Geometric Completeness (Fig. 4e) are reported using pre-defined categorical bins.

Across these informativeness axes, we observe skewed distributions toward higher levels of visual informativeness. Such tendencies are characteristic of official listing imagery provided by hospitality platforms, which typically relies on professional photography to enhance spatial clarity, contextual visibility, and overall visual appeal for potential guests. Unlike user-generated review photos, these images are intentionally composed to reveal room structure, spatial volume, and surrounding context. As a result, the observed distributions reflect the visual evidence that users commonly encounter during actual booking decisions, supporting the ecological validity of our benchmark. This structured coverage enables models to be evaluated not only on generic scene understanding, but also on the decision-relevant visual properties that matter in hospitality settings.

In Section 5, we use this dataset to benchmark several state-of-the-art VLMs and analyze their performance across facility types and informativeness axes.

5 Experiments

We evaluate a range of general-purpose Vision–Language Models (VLMs) on Hospitality-VQA to examine how well they capture the domain-specific informativeness axes introduced in Section 3.

5.1 Experimental Setup

Data split. Hospitality-VQA contains 5,000 labeled accommodation images. We reserve 300 images for evaluation. The remaining 4,700 images are used for training. The evaluation split is sampled to preserve the overall distribution of facility types and informativeness factors, with class proportions matched within a 5% margin relative to the full dataset.

Models. We evaluate eight vision–language models that span both commercial APIs and open-weight systems: GPT-5 (OpenAI, 2025), GPT-4o-mini (Hurst et al., 2024), Gemini 2.5 Pro (Comanici et al., 2025), GLM-4.1V-9B-Thinking (Hong et al., 2025), Qwen2.5-VL-3B and Qwen2.5-VL-7B (Bai et al., 2025), LLaVA-NeXT-7B (Li et al., 2024), and Gemma-3-12B (Team et al., 2025). The proprietary models (GPT-5, GPT-4o-mini, Gemini 2.5 Pro, and GLM-4.1V-9B-Thinking) serve as strong general-purpose assistants that have been optimized for broad, web-scale multimodal use, whereas the open-weight models (Qwen2.5-VL-3B, Qwen2.5-VL-7B, LLaVA-NeXT-7B, and Gemma-3-12B) provide instruction-tuned checkpoints with varying capacities and training pipelines that are accessible for research and adaptation. This combination allows us to examine how both deployment setting and model family affect performance on hospitality-oriented VQA.

Beyond zero-shot evaluation, we also derive task-adapted variants of Qwen2.5-VL-3B and Qwen2.5-VL-7B by applying LoRA fine-tuning (Hu et al., 2022) on Hospitality-VQA. In this configuration, the models are trained to predict the discrete axis labels in our framework from an image–prompt pair, aligning their outputs with our informativeness-oriented, classification-style supervision rather than generic captioning or open-ended generation.

Tasks and metrics. To align with real-world hospitality applications (e.g., booking platforms) that require discrete, interpretable attributes rather than free-form text, we formulate all tasks as classification problems. We evaluate six core tasks—main facility type, main+sub facility type, visible plane count, meaningful component count, discretized scenery proportion, and building-face visibility—along with two auxiliary interior attributes: room and view type.

Models are prompted with a natural-language instruction template and must output a single categorical label. We report exact-match accuracy, reflecting the binary nature of practical decision-making; predictions that fail to map to a valid label are strictly penalized, mirroring real-world failure modes in attribute extraction systems. For API-based models, we use deterministic decoding (temperature = 0).

Model	Facility		Informativeness					
	Main	Main&Sub	SL	AA	CO	GC	Room	View
Gemini 2.5 Pro	90.66	75.00	11.51	9.43	46.81	7.35	80.65	50.00
GPT-5	<u>92.33</u>	82.55	46.76	18.87	31.91	20.59	<u>83.87</u>	64.10
GPT-4o-mini	<u>92.33</u>	<u>84.91</u>	97.12	38.21	56.03	8.82	70.97	79.49
GLM-4.1V-9B-Thinking	93.66	79.25	89.21	35.85	<u>57.45</u>	16.18	61.29	57.69
LLaVA-NeXT-7B	73.33	53.77	94.24	8.02	19.86	5.88	25.81	79.49
Gemma-3-12B	92.00	82.08	86.33	15.09	55.32	22.06	72.19	43.59
Qwen2.5-VL-3B	64.66	44.34	68.35	19.34	39.72	1.47	41.94	70.51
Qwen2.5-VL-3B Finetuned	86.66	81.13	<u>94.96</u>	<u>42.92</u>	<u>57.45</u>	<u>26.47</u>	80.65	<u>76.92</u>
Qwen2.5-VL-7B	78.66	64.15	43.88	25.94	48.94	5.88	25.81	69.23
Qwen2.5-VL-7B Finetuned	92.00	85.37	97.12	44.34	67.37	32.35	87.10	74.36

Table 2: Comparison of VLM performance across facility types and informativeness categories. Best in each column is in **bold** and second-best is underlined.

5.2 Overall Results

Table 2 reports accuracy across facility recognition and all informativeness-related tasks. We summarize the results by (i) task difficulty across axes, (ii) model-family trends, and (iii) the effect of domain adaptation.

5.2.1 Task Difficulty Across Axes

Table 2 shows that *main* facility classification transfers well across most evaluated VLMs, with several models exceeding 90% accuracy. In contrast, *main&sub* recognition is consistently lower, indicating that fine-grained sub-category prediction is more demanding than coarse scene categorization under the same prompting and evaluation protocol.

Axis-level tasks exhibit a sharper drop in performance than facility recognition. Across models, Spatial Legibility (SL) is generally more stable than the other informativeness axes, whereas Activity Affordance (AA) and Geometric Completeness (GC) are notably weaker for many models. Contextual Openness (CO) falls between these extremes but still remains substantially below facility recognition performance, suggesting that decision-relevant attributes are not reliably recovered from generic multimodal capabilities alone.

Room and view attributes for *Room Interior* show additional variability across models. While some models achieve strong accuracy on these auxiliary tasks, others lag despite high facility recognition, reinforcing that success on global categorization does not guarantee robust prediction of hospitality-relevant fine-grained attributes.

5.2.2 Model Family Trends

Model families show broadly similar behavior on coarse facility recognition but diverge more on axis-level prediction. Several proprietary models achieve high accuracy on *main* facility classification, and some open-weight models also reach comparable levels, indicating that recognizing the overall facility category is not the primary bottleneck in this benchmark.

Differences become more pronounced for informativeness axes. For instance, GPT-4o-mini attains very high SL accuracy (97.12), yet AA and GC remain much lower (38.21 and 8.82). A similar pattern appears in multiple open-weight baselines (e.g., Qwen2.5-VL-7B: SL 43.88 vs. AA 25.94 and GC 5.88), where axis-level prediction does not track facility recognition. These results suggest that axis performance reflects additional reasoning requirements beyond generic scene labeling.

We avoid attributing these gaps to a single cause, as controlled ablations over training data, vision encoders, and instruction-tuning procedures are outside the scope of this work. Nonetheless, the consistent separation between facility recognition and axis-level performance across both proprietary and open-weight systems motivates explicit domain-grounded supervision for decision-oriented attributes.

5.2.3 Effect of Domain Adaptation

Domain adaptation via LoRA fine-tuning (Hu et al., 2022) consistently improves Qwen2.5-VL models across all evaluated tasks. Table 3 reports absolute gains (Finetuned–Base) computed from Table 2. Improvements are observed for both coarse facil-

Task	3B (Δ Acc)	7B (Δ Acc)
Main	+22.00	+13.34
Main&Sub	+36.79	+21.22
SL	+26.61	+53.24
AA	+23.58	+18.40
CO	+17.73	+18.43
GC	+25.00	+26.47
Room	+38.71	+61.29
View	+6.41	+5.13

Table 3: Absolute accuracy gains (%) from domain adaptation via LoRA fine-tuning for Qwen2.5-VL models (Finetuned–Base).

ity recognition and fine-grained facility prediction, with particularly large gains on *main&sub* classification.

Gains are also evident on informativeness axes, which are challenging in the zero-shot setting. Notably, both model sizes improve on AA, CO, and GC, while the 7B model shows a pronounced increase on SL. Interior attributes benefit as well: room type accuracy increases substantially for both models, whereas view type shows smaller but consistent gains. Overall, these results indicate that axis-aligned supervision in Hospitality-VQA provides an effective signal for aligning VLM outputs with decision-oriented hospitality attributes under a strict label-matching evaluation.

6 Conclusion

This work addressed the gap between general-purpose visual understanding and the kinds of fine-grained, decision-relevant reasoning required in the hospitality domain. While images play a central role in shaping guest expectations and booking decisions, existing multimodal systems lack the structured grounding necessary to interpret the spatial, functional, and view-related cues that matter in real domain use cases, just interpreting surface-level visual scenes.

To bridge this gap, we introduced *Hospitality Informativeness*, a domain-grounded framework that formalizes four fundamental visual axes—spatial legibility, activity affordance, contextual openness, and geometric completeness, whom are interpretable and measurable. Building on this framework, we constructed **Hospitality-VQA**, a decision-centric VQA benchmark designed to elicit and evaluate the kinds of visual evidence that influence guest perception across diverse facility types.

E.g., whether models capture layout, functional components, scenery, and exterior visibility that matter for booking decisions. Together, these contributions provide the first structured basis for measuring how well VLMs interpret hospitality imagery beyond generic scene recognition.

Our empirical study revealed that state-of-the-art general-purpose VLMs struggle with the fine-grained informativeness reasoning that the hospitality domain demands. However, we also showed that lightweight domain adaptation using our dataset leads to consistent and measurable improvements, highlighting both the challenge of the task and the value of the benchmark as a foundation for future model development.

Future Directions Looking ahead, Hospitality-VQA and the hospitality informativeness framework open several research directions, including domain-aware representation learning, prompt optimization, and test-time reasoning strategies. A particularly promising extension is modeling *human-preferred accommodation attractiveness*, as user impressions are often shaped by images. This line of work carries clear practical value: **B2C** applications include displaying more appealing images to improve user experience and booking rates, while **B2B** applications involve curating and ranking property images based on user appeal. We hope our benchmark provides a foundation for future advances in hospitality-aware multimodal intelligence that benefits both users and service providers.

Limitations

This work has several limitations. First, Hospitality-VQA focuses on static images collected from a specific set of hotels and platforms, and the proposed informativeness axes represent a pragmatic but necessarily incomplete abstraction of real-world user information needs. In particular, while our framework emphasizes functional, spatial, and contextual visual cues, it does not explicitly capture aesthetic qualities such as visual style, ambiance, or emotional appeal, which can also influence user preferences in hospitality settings.

Second, our study does not model additional modalities or contextual factors commonly involved in accommodation decisions, such as textual reviews, pricing information, temporal media (e.g., videos), or personalized user preferences. As a result, the evaluation is limited to image-based visual

reasoning under a controlled decision setting.

Third, all model evaluations are conducted under a single annotation protocol and question formulation. We do not assess the robustness of the reported results under alternative labeling schemes, prompt designs, or downstream task definitions.

Finally, although the dataset contains 5,000 annotated images in total, quantitative evaluation is performed on a held-out subset of 300 images. This relatively small evaluation set may limit statistical power and reduce sensitivity to rare or long-tail cases.

Acknowledgments

The views and conclusions expressed in this paper are those of the authors and should not be interpreted as representing the official views, policies, or products of their affiliated organization.

References

- Arian Askari, Emmanouil Stergiadis, Ilya Gusev, and Moran Beladev. 2025. Hotelmatch-llm: Joint multi-task training of small and large language models for efficient multimodal hotel retrieval. *arXiv preprint arXiv:2506.07296*.
- Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, and 1 others. 2025. Qwen2.5-VL Technical Report. *arXiv preprint arXiv:2502.13923*.
- Gheorghe Comanici, Eric Bieber, Mike Schaekermann, Ice Pasupat, Noveen Sachdeva, Inderjit Dhillon, Marcel Blistein, Ori Ram, Dan Zhang, Evan Rosen, and 1 others. 2025. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *arXiv preprint arXiv:2507.06261*.
- Pedro Cuesta-Valiño, Sergey Kazakov, Pablo Gutiérrez-Rodríguez, and Orlando Lima Rua. 2023. [The effects of the aesthetics and composition of hotels' digital photo images on online booking decisions](#). *Humanities and Social Sciences Communications*, 10:59.
- Abhishek Das, Samyak Datta, Georgia Gkioxari, Stefan Lee, Devi Parikh, and Dhruv Batra. 2018. Embodied question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. 2017. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Michelle R. Greene, Christopher Baldassano, Andre Esteva, Diane M. Beck, and Li Fei-Fei. 2016. [Visual scenes are categorized by function](#). *Journal of Experimental Psychology: General*, 145(1):82–94.
- Dario Guidotti, Laura Pandolfo, and Luca Pulina. 2025. Discovering sentiment insights: streamlining tourism review analysis with large language models. *Information Technology & Tourism*, 27(1):227–261.
- Danna Gurari, Quchen Li, Anthony J Stangl, Yongsun Guo, Chuan-He Lin, Kristen Grauman, Jiebo Luo, and Jeffrey P Bigham. 2018. Vizwiz: Nearly real-time answers to visual questions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Jiaxiu He, Bingqing Li, and Xin Shane Wang. 2023. Image features and demand in the sharing economy: A study of airbnb. *International Journal of Research in Marketing*, 40(4):760–780.
- Wenyi Hong, Wenmeng Yu, Xiaotao Gu, Guo Wang, Guobing Gan, Haomiao Tang, Jiale Cheng, Ji Qi, Junhui Ji, Lihang Pan, and 1 others. 2025. Glm-4.1 v-thinking: Towards versatile multimodal reasoning with scalable reinforcement learning. *arXiv preprint arXiv:2507.01006*.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, and 1 others. 2022. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2):3.
- Drew A Hudson and Christopher D Manning. 2019. Gqa: A new dataset for real-world visual reasoning and compositional question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Aaron Hurst and 1 others. 2024. GPT-4o System Card. *arXiv preprint arXiv:2410.21276*.
- Bo Li, Kaichen Zhang, Hao Zhang, Dong Guo, Renrui Zhang, Feng Li, Yuanhan Zhang, Ziwei Liu, and Chunyuan Li. 2024. [Llava-next: Stronger llms supercharge multimodal capabilities in the wild](#).
- Huiying Li, Qiang Ye, and Rob Law. 2013. Determinants of customer satisfaction in the hotel industry: An application of online review analysis. *Asia Pacific journal of tourism research*, 18(7):784–802.
- Aude Oliva and Antonio Torralba. 2001. Modeling the shape of the scene: A holistic representation of the spatial envelope. *International journal of computer vision*, 42(3):145–175.
- OpenAI. 2025. [Gpt-5 system card](#).
- Meng Ren, Huy Quan Vu, Gang Li, and Rob Law. 2021. Large-scale comparative analyses of hotel photo content posted by managers and customers to review platforms based on deep learning: implications for hospitality marketers. *Journal of Hospitality Marketing & Management*, 30(1):96–119.
- Javier A. Sánchez-Torres, Sandra-Milena Palacio-López, Yuri Hernandez-Fernandez, Francisco J. Arroyo-Cañada, and Ana Argila-Irurita. 2024. [Visual photography's influences on hotel selection: an analysis using e-booking as a comparative platform](#). *International Journal of Electronic Customer Relationship Management*, 14(2):128–142.
- Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivière, and 1 others. 2025. Gemma 3 technical report. *arXiv preprint arXiv:2503.19786*.
- Ameni Trabelsi, Maria Zontak, Yiming Qian, Brian Jackson, Suleiman Khan, and Umit Batur. 2025. What matters when building vision language models for product image analysis? In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision Workshops (WACV Workshops)*.
- Tao Tu, Shekoofeh Azizi, Danny Driess, Mike Schaekermann, Mohamed Amin, Pi-Chuan Chang, Andrew Carroll, Charles Lau, Ryutaro Tanno, Ira Ktena, and 1 others. 2024. Towards generalist biomedical ai. *Nejm Ai*, 1(3):AIoa2300138.

- Kirk L. Wakefield and Jeffrey G. Blodgett. 1996. [The effect of the servicescape on customers' behavioral intentions in leisure service settings.](#) *Journal of Services Marketing*, 10(6):45–61.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, and 1 others. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837.
- Zheng Xiang, Zvi Schwartz, John H Gerdes Jr, and Muzaffer Uysal. 2015. What can big data and text analytics tell us about hotel guest experience and satisfaction? *International journal of hospitality management*, 44:120–130.
- Shunyuan Zhang, Dokyun Lee, Param Vir Singh, and Kannan Srinivasan. 2022. What makes a good image? airbnb demand analytics leveraging interpretable image features. *Management Science*, 68(8):5644–5666.

A Detailed Facility Taxonomy

To support consistent annotation and evaluation, we define a hierarchical facility taxonomy with clear operational criteria. Images are first assigned to one of five main facility types based on accessibility and visible structural boundaries: *Room Interior*, *Indoor Facility*, *Outdoor Facility*, *Accommodation Exterior*, and *Irrelevant*.

Room Interior is restricted to private guest spaces. In cases of spatial overlap (e.g., studio-type rooms), a fixed priority order is applied (*Bedroom* > *Kitchen* > *Bathroom* > *Living room*) to ensure unique assignment. **Indoor** and **Outdoor Facilities** are distinguished by whether the space is fully enclosed, with outdoor facilities required to be the primary visual focus rather than part of a general landscape. **Accommodation Exterior** is assigned only when the building itself constitutes the main subject with identifiable accommodation features. Images lacking discernible hospitality context are grouped into the **Irrelevant** category, which functions as a noise class.

For finer-grained functional analysis, each main category (except *Accommodation Exterior* and *Irrelevant*) is further annotated with sub-facility labels, summarized in Table 4. This granularity enables evaluation of whether models can recognize specific functional contexts relevant to hospitality decision-making.

Main Category	Sub-category
Room Interior	Bedroom; Kitchen; Bathroom; Living room
Indoor Facility	Guest lounge; Reception desk; Hallway; Restaurant & Cafe; Indoor pool; Indoor parking lot; Other amenities
Outdoor Facility	Outdoor pool & Spa; Outdoor lounge & BBQ area; Sports & Recreation facility; Outdoor parking lot; Camping area
Accommodation Exterior	—
Irrelevant	—

Table 4: Full taxonomy of hospitality facility classification and sub-category labels.

General Instruction–Answer Template Format
<p>Task Target classification or assessment objective.</p>
<p>Prompt Natural language instruction defining task semantics and decision rules.</p>
<p>Answer Format Strictly constrained output schema (e.g., class ID or fixed key–value pairs).</p>
<p>Answer Example output following the specified format.</p>

Figure 5: General structure of instruction–answer templates shared across all evaluation tasks.

B Instruction-Answer Construction Template

This appendix provides details on how the expert-verified labels are mapped into instruction–answer pairs using our fixed templates. As described in Section 4.3, these templates are designed to ensure consistency across the dataset by formatting classification targets into a standardized VQA format.

To maintain evaluation rigor, each template consists of a task-specific prompt and a constrained answer format. Figure 5 illustrates the general structure of these templates. Representative examples for facility-type classification and informativeness axis evaluation are presented in Figures 6 and 7, respectively.

Task: Facility Type (Main)

Prompt: Your task is to classify given image.

Definitions and specific instructions for each category are as follows:

1.

Private accommodation room interior space for guest sleeping/living functions. Includes bedrooms, bathrooms, living rooms, kitchens, and photos taken from inside rooms. Shared areas or facilities do not belong to this category.

2.

Shared "indoor" facilities within accommodation (e.g. customer lounges, reception desks, corridors, restaurants/cafes, indoor pools, indoor parking, other amenities (gyms, indoor golf, saunas, convenience stores, seminar rooms etc.))

3.

Specific "outdoor" facilities that falls into following cases: outdoor pools/spas, outdoor lounges/garden/terrace/BBQ areas, outdoor sports/recreation facilities, outdoor parking, outdoor camping areas. Must be clearly identifiable as one of these facility types and be the image's primary focus, not part of general accommodation or landscape views. Exclude: overall building/accommodation views even if outdoor facilities are visible, pure nature shots without specific facilities.

4.

Image showing accommodation building exterior AS THE MAIN SUBJECT. Building must occupy significant portion of image with clear structural elements (walls, windows, roof) and typical accommodation features (guestroom windows, balconies, nearby amenities) to be identifiable as accommodation. Only for photos that do NOT fall into categories 1, 2, 3, or 5 AND where the building itself is the primary focus, not background. Excluded: no visible building, appears to be non-residential buildings due to lack of accommodation features, main accommodation building unclear among multiple scattered buildings, accommodation too small/distant to recognize (e.g. tiny in a wide drone shot).

5.

Image lacking any clues identifying them as prior 4 categories of accommodation. Includes pure nature shots, pet/person-focused shots without spatial clues, notice/text-based images (e.g., posters, receipts), and building exteriors not meeting criteria of 4. For close-up images, prefer classifying to 1-4 over 5 when possible, if there are any clues suggesting accommodation context, even if subtle (e.g., cushion seems to be on bed → 1, food seems to be in restaurant → 2).

ANSWER FORMAT

Output a single number: <1-5>

Do not include any explanation, spaces, or other characters.

Answer: 1

Figure 6: Example of constructed instruction–answer template of main facility.

Task: Geometric Completeness

Prompt: You are an expert image analyst specializing in architectural assessments.

Your task is to analyze the visible faces of the single most plausible and visually prominent lodging building in an image.

For the selected building, output a visibility status code (1–4) for each of these three faces:

- '1' = Front Facade
- '2' = Side Wall
- '3' = Roof

Apply the rules in this order for each face:

1. Assign 1 if the face is absent or not visible at all.
2. Assign 2 if a clear, identifiable portion of the face is cut off by the image's edges.
3. Assign 3 if the face is visible but significantly blocked by an external object.
4. Assign 4 if the face is clearly visible and unobstructed (roof only if distinct and unambiguous).

ANSWER FORMAT

Output exactly in this format, with no spaces or extra text:

'1': <1-4>, '2': <1-4>, '3': <1-4>

Answer: '1': 3, '2': 3, '3': 4

Figure 7: Example of constructed instruction–answer template of geometric completeness.

C Additional Experimental Details

We provide implementation details to support the reproducibility of the results in Section 5. All fine-tuning experiments on open-weight models were conducted on a single NVIDIA RTX 4090 GPU using the unsloth framework.

C.1 Training Setup

Models were fine-tuned for two epochs using supervised learning on image–instruction pairs. We used the AdamW optimizer with a learning rate of 2×10^{-5} , 5% linear warmup, and cosine decay. The effective batch size was 16 (batch size 2 per device with gradient accumulation of 8). Training was performed in bfloat16 precision with a maximum context length of 8,192 tokens. Weight decay and gradient clipping were not applied.

C.2 LoRA Configuration

We adopted Low-Rank Adaptation (LoRA) (Hu et al., 2022) with rank $r = 16$ and scaling factor $\alpha = 32$. Adapters were inserted into the vision encoder and language decoder, covering the attention projections and MLP layers. LoRA dropout was set to 0, and all other model parameters were frozen.

D CoT vs. No CoT

We explicitly investigated the impact of incorporating Chain-of-Thought (CoT) (Wei et al., 2022) reasoning during the supervised fine-tuning process. Table 5 presents a performance comparison between the base models, models fine-tuned with CoT supervision, and models fine-tuned with direct answers (w/o CoT).

In the 3B setting, CoT supervision provides small gains on a few attributes (e.g., *Scenery* and *Building Faces*), but these improvements are neither consistent across tasks nor robust across model scales. Overall, direct-answer supervision without CoT yields more reliable performance for our classification-oriented evaluation.

Model	Facility		Informativeness					
	Main	Main&Sub	SL	AA	CO	GC	Room	View
Qwen2.5-VL-3B	64.66	44.34	68.35	19.34	39.72	1.47	41.94	70.51
Qwen2.5-VL-3B FT (w/o CoT)	86.66	81.13	94.96	42.92	57.45	26.47	80.65	76.92
Qwen2.5-VL-3B FT (w/ CoT)	85.66	73.58	93.53	34.91	63.38	27.94	64.52	70.51
Qwen2.5-VL-7B	78.66	64.15	43.88	25.94	48.94	5.88	25.81	69.23
Qwen2.5-VL-7B FT (w/o CoT)	92.00	85.37	97.12	44.34	67.37	32.35	87.10	74.36
Qwen2.5-VL-7B FT (w/ CoT)	91.33	83.02	94.24	42.45	59.57	26.47	83.87	76.92

Table 5: Comparison of VLM performance with and without CoT. Best in each column is highlighted in **bold**.