

From Scenes to Elements: Multi-Granularity Evidence Retrieval for Verifiable Multimodal RAG

Guanhua Chen*, Chuyue Huang*, Yutong Yao, Shudong Liu, Xueqing Song,
Lidia S. Chao, Derek F. Wong†

NLP²CT Lab, Department of Computer and Information Science, University of Macau
{nlp2ct.guanhua, nlp2ct.chuyue, nlp2ct.yutong, nlp2ct.shudong, xqsongangie}@gmail.com
{derekfw, lidiasc}@um.edu.mo

Abstract

Multimodal Retrieval-Augmented Generation (RAG) systems retrieve evidence at coarse granularities (entire images or scenes), creating a mismatch with fine-grained user queries and making failures unverifiable. We introduce **GranuVistaVQA**, a multimodal benchmark featuring real-world landmarks with element-level annotations across multiple viewpoints, capturing the partial observation challenge where individual images contain only subsets of entities. We further propose **GranuRAG**, a multi-granularity framework that treats visual elements as first-class retrieval units through three stages: element-level detection and classification, multi-granularity cross-modal alignment for evidence retrieval, and attribution-constrained generation. By grounding retrieval at the element level rather than relying on implicit attention, our approach enables transparent error diagnosis. Experiments demonstrate that GranuRAG achieves up to 29.2% improvement over six strong baselines for this task.

1 Introduction

Multimodal Large Language Models (MLLMs) have substantially advanced visual understanding (Alayrac et al., 2022; Li et al., 2023). However, their reasoning remains opaque and prone to hallucination (Bai et al., 2024). While Retrieval-Augmented Generation (RAG) mitigates this by conditioning generation on external evidence (Lewis et al., 2020; Chen et al., 2025a), current multimodal extensions operate at coarse granularities, retrieving entire images, scenes, or pages (Fang et al., 2024; Yu et al., 2025a). This creates a fundamental attribution gap. A user asking about a Baroque pediment may receive a building photograph, yet neither party can verify whether the pediment appears in the image, whether the

retrieved knowledge is relevant, or whether the answer faithfully reflects these inputs. Detection failures, retrieval errors, and generation hallucinations collapse into an opaque black box.

In this work, we argue that verifiable multimodal RAG requires treating visual elements as first-class retrieval targets, not merely as regions to be implicitly attended within retrieved scenes. Although recent grounding models can localize objects (Peng et al., 2023; Guo et al., 2024), they rely on parametric knowledge and do not retrieve external evidence. Conversely, existing multimodal RAG systems retrieve evidence but lack explicit element-level grounding, even recent fine-grained approaches (Liu et al., 2025) prioritize representational expressiveness over transparent attribution. We bridge this gap with a detect-then-retrieve approach: we first detect candidate visual elements via open-vocabulary detection, then assemble hierarchical evidence spanning element-level descriptions and global context, and finally generate answers attributable to specific visual spans and retrieved passages. This design transforms evaluation from black-box answer assessment into transparent evidence auditing: we can separately diagnose whether the correct elements were detected, whether relevant knowledge was retrieved, and whether generation remained faithful to both.

However, existing benchmarks inadequately assess multi-granularity alignment as shown in Table 1. Critically, all overlook the partial observation challenge inherent to real-world imagery: photographs capture scenes from varying distances and angles, such that a single image depicts only a subset of elements present at a location. This challenge pervades domains from architectural photography to medical imaging and satellite sensing, yet no existing benchmark provides the supervision needed to diagnose element-level detection and retrieval under partial visibility. To address this, we introduce **GranuVistaVQA**, a benchmark centered on

*Equal contribution

†Corresponding authors.

Dataset	Multi-granularity	Fine-grained alignment	Partial entities	RAG
REAL-MM-RAG (Wasserman et al., 2025)	✗	✓	✗	✓
MMDocIR (Dong et al., 2025)	✓	✗	✗	✗
M3DocRAG (Cho et al., 2024)	✗	✗	✗	✓
MMLongBench-Doc (Ma et al., 2024b)	✗	✗	✗	✗
SPIQA (Pramanick et al., 2024)	✗	✓	✗	✗
GranuVistaVQA	✓	✓	✓	✓

Table 1: Comparison of our dataset with five similar multimodal datasets.

architectural heritage landmarks, a domain where elements have well-defined visual semantics, authoritative knowledge sources exist, and multi-view partial observation naturally arises from real-world photography. The dataset comprises 1,422 images across 71 landmarks, where each view covers only 34% of annotated elements on average. Crucially, we provide human-verified element visibility labels that enable fine-grained error diagnosis unavailable in prior benchmarks. We further propose GranuRAG, a detect-then-retrieve framework that grounds visible elements via open-vocabulary detection, retrieves hierarchical evidence, and generates attribution-constrained answers.

Experiments show GranuRAG outperforms strong baselines. Moreover, LLMs fine-tuned on our pipeline’s reasoning traces surpass both direct fine-tuning and self-generated chain-of-thought (CoT) (Wei et al., 2022), demonstrating that explicit multi-granularity alignment provides a more effective supervision signal.

2 Related Work

Benchmarks for Multimodal RAG Knowledge-intensive multimodal QA has evolved from answer-only evaluation toward attribution-aware assessment requiring verifiable evidence and localized failure analysis. Early benchmarks (Marino et al., 2019; Schwenk et al., 2022) established the need for external knowledge without explicit grounding. Follow-up datasets added structured supervision: VIQUAE (Lerner et al., 2022) identifies retrieval as the primary bottleneck; INFOS-EEK (Chen et al., 2023) and ENCYCLOPEDIA-VQA (Mensink et al., 2023) provide section-level evidence, revealing gaps due to unreliable entity-section linking. Document-centric benchmarks (Yu et al., 2025b; Xu et al., 2025b) evaluate at page, region, and document level granularities. Recent work incorporates explicit spatial supervision: BBOX-DOCVQA (Yu et al., 2025c) grounds answers to semantically coherent regions; TOLOKA

VQA (Ustalov et al., 2023) requires bounding boxes for answer-supporting objects; VISA (Ma et al., 2024a) mandates visual source attribution during generation. However, existing benchmarks lack explicit mappings between visual elements and knowledge entries. GranuVistaVQA addresses this by treating visual elements as core evidence units with element-level knowledge alignment for fine-grained, verifiable attribution.

LLM Methods for Multimodal RAG Multimodal RAG has advanced through richer retrieval representations, stronger reranking, and controllable attribution. Unified dense retrievers (Liu et al., 2023; Zhou et al., 2024b,a) use joint text-image embeddings but lack explicit local visual-textual connections. Finer-grained methods (Lin et al., 2024; Yang et al., 2025) operate at paragraph/section levels, leaving element-level grounding underexplored. End-to-end pipelines (Chen et al., 2022; Zhang et al., 2024b) improve recall through iterative retrieval but amplify noise. Robustness-focused approaches (Cui et al., 2024; Yan and Xie, 2024; Tian et al., 2025) address cross-source reconciliation but reason over coarse units. Recent spatial control work targets specific evidence chain components: LOCATE-THEN-GENERATE (Zhu et al., 2023) separates localization from generation in scene-text VQA; HULIRAG (Xi et al., 2025) decouples retrieval from attention via segmentation but lacks grounding-knowledge connection; GROUNDHOG (Zhang et al., 2024c) achieves pixel-level alignment without retrieval integration; VISRAG 2.0 (Sun et al., 2025) improves multi-image reasoning but treats evidence as disconnected region sets; FERRET-V2 (Zhang et al., 2024a) enables fine-grained region-language alignment for in-image comprehension without external knowledge integration. In contrast, GranuRAG treats individual elements as verifiable units, grounding each factual claim to both a detected region and a retrieved snippet, which enables fine-grained alignment and supports systematic er-

ror diagnosis across detection, retrieval, and generation stages.

3 GranuVistaVQA Benchmark

To enable verifiable multimodal RAG with fine-grained attribution, we construct a knowledge-intensive benchmark centered on urban architectural heritage. Unlike prior datasets that treat images as atomic units, our design establishes visual elements as first-class retrieval targets and explicitly models the partial observation challenge: real-world photographs capture landmarks from varying viewpoints, each depicting only a subset of architecturally significant elements.

3.1 Task Formulation

Given a query image I depicting a landmark from an arbitrary viewpoint, the task is to generate a comprehensive description covering all visible architectural elements while avoiding hallucination about occluded or absent components.

We associate each landmark with three components: metadata (name, summary, and style) that provides high-level context; an element inventory $E = \{e_1, \dots, e_k\}$ listing architecturally significant components; and element descriptions $ED : E \rightarrow Paragraphs$ that map each element to expert-written text. For each image I , the ground-truth visible set $E^{gt}(I) \subseteq E$ contains elements visually identifiable in that view. This formulation enables modular evaluation: systems must (i) predict visible elements $\hat{E}(I) \approx E^{gt}(I)$, (ii) retrieve relevant descriptions from ED, and (iii) generate outputs faithful to both visual and textual evidence. This decomposition allows us to isolate failures at each stage, distinguishing detection errors from retrieval mistakes and generation hallucinations.

3.2 Data Collection and Annotation

Domain Selection We focus on architectural heritage for three methodological reasons: (1) elements have well-defined visual semantics amenable to detection, (2) authoritative knowledge sources enable reliable ground truth, and (3) tourist photography naturally exhibits multi-view partial observation. We curate 71 landmarks from official cultural heritage databases, spanning religious buildings, temples, fortifications, and cultural institutions across diverse architectural traditions.

Knowledge Corpus Construction For each landmark, we compile a structured JSON document




Figure 1: Examples of Multi-Perspective Image.

following a two-level schema (full specification in Appendix B.1):

$$x_{\text{landmark}} = (\text{meta}, E, ED) \quad (1)$$

Textual content is sourced from official tourism portals and encyclopedic references, then structured through: (i) element phrase extraction from authoritative descriptions, (ii) cross-landmark normalization to ensure consistent terminology (e.g., “bell tower” \equiv “campanile”), and (iii) LLM-assisted description generation with human validation (details in Appendix B.2). Specifically, we focus on heritage sites in Macau, with all description content in Chinese.

Image Collection We collect 1,422 photographs ensuring viewpoint diversity: panoramic shots capturing overall structure, close-ups revealing fine ornamentation, and oblique partial views (Figure 1). After collection, we perform comprehensive data sanitization to remove privacy-sensitive content, including watermarks, visible human faces, and personally identifiable information. We also apply quality filtering to retain only images with resolution $\geq 512\text{px}$ and no visible artifacts. The full screening protocol is described in Appendix B.2.

Visibility Annotation For each image I , annotators identify $E^{gt}(I)$ following strict visibility criteria:

- **Visual identifiability:** Elements qualify only if recognizable from pixels alone, without relying on prior landmark knowledge
- **Partial occlusion:** Included only when discriminative visual cues remain (e.g., a half-visible scroll counts if its characteristic shape is apparent)
- **Ambiguity handling:** Uncertain cases are excluded to avoid false positives

We employ a human-in-the-loop workflow: an LLM proposes candidate elements, which annotators refine by adding missed elements, remov-

Metric	Value
#Landmark (L)	71
#Img (N)	1422
Avg img/landmark (N/L)	20.03
#Unique elements (U_E)	221
Avg elements per landmark	3.59

Table 2: The statistic of our proposed GranuVistaVQA.

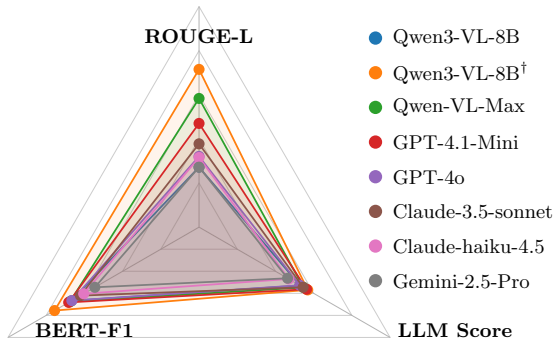


Figure 2: The results of evaluating MLLMs on GranuVistaVQA. † means the fine-tuned LLM.

ing hallucinated ones, and resolving synonyms to canonical forms (protocol in Appendix B.6).

3.3 Data Statistic and Evaluation

Statistic Table 2 summarizes the dataset statistics. The multi-view design creates natural partial observations. close-ups capture fine details but miss broader structures, while panoramic shots show overall layouts but lose granular information (Figure 1). On average, individual images cover only 34% of their landmark’s element inventory. This validates our core premise that answering element-level questions requires aggregating evidence across complementary viewpoints. The statistic of image distribution is explained in Appendix B.7.

Evaluation Metrics To evaluate our framework, we report three evaluation metrics: ROUGE-L (Lin, 2004) and BERT-F1 (Zhang et al., 2020) computed against the gold-standard reference, and an LLM-as-a-judge score (Zheng et al., 2023). For the LLM-Score, we use an ensemble of three strong LLMs (GPT-4.1, Gemini-2.5-Pro, claude-haiku-4-5 (OpenAI, 2024; Team and Google, 2025; Anthropic, 2025)) to rate each explanation on a 0 – 100 scale under a weighted rubric covering Coverage (40%), Faithfulness (40%), and Cohesion (20%); the final LLM Score is the mean across judges to reduce

single-model variance (Chia et al., 2024).

Benchmark Evaluation We evaluated state-of-the-art (SOTA) different powerful MLLMs on our benchmark. Figure 2 shows that current models struggle across all metrics as shown in Section 5.1. Even the best-performing model achieves limited success on ROUGE-L and BERT-F1, while all models perform poorly on LLM-as-Judge evaluation, which measures factual accuracy and hallucination control. These results highlight the challenge of multi-granularity information alignment, motivating our GranuRAG framework.

4 Methodology

Based on the above analysis, we introduce our GranuRAG framework as shown in Figure 3: given a query image I and candidate element set E , we (1) localize and match visual regions to elements in E , yielding visible subset $\hat{E}(I)$; (2) retrieve hierarchical evidence for $\hat{E}(I)$; and (3) generate attribution-constrained output grounded in the retrieved evidence and global description.

4.1 Visual Region Detection and Filtering

The first stage localizes salient architectural regions in image I without requiring prior knowledge of which elements are present. We employ YOLO-World (Cheng et al., 2024), an open-vocabulary object detector, to identify candidate regions based on generic architectural primitives such as columns, carvings, and decorative motifs. This detection strategy offers broad coverage across diverse heritage landmarks without domain-specific fine-tuning.

Raw detections often contain redundant bounding boxes that capture the same visual content at multiple scales. To address this issue, we apply overlap-based filtering: when two boxes overlap by more than 80%, we retain the smaller one to preserve fine-grained architectural details. The sensitive analysis of this overlap percentage will be shown in Appendix A.2. This denoising step produces a refined set of cropped regions:

$$\mathcal{B}(I) = \{b_1, b_2, \dots, b_K\} \quad (2)$$

where each b_k represents a distinct visual region likely to contain meaningful architectural content. The filtered crops serve as visual queries for the subsequent matching stage.

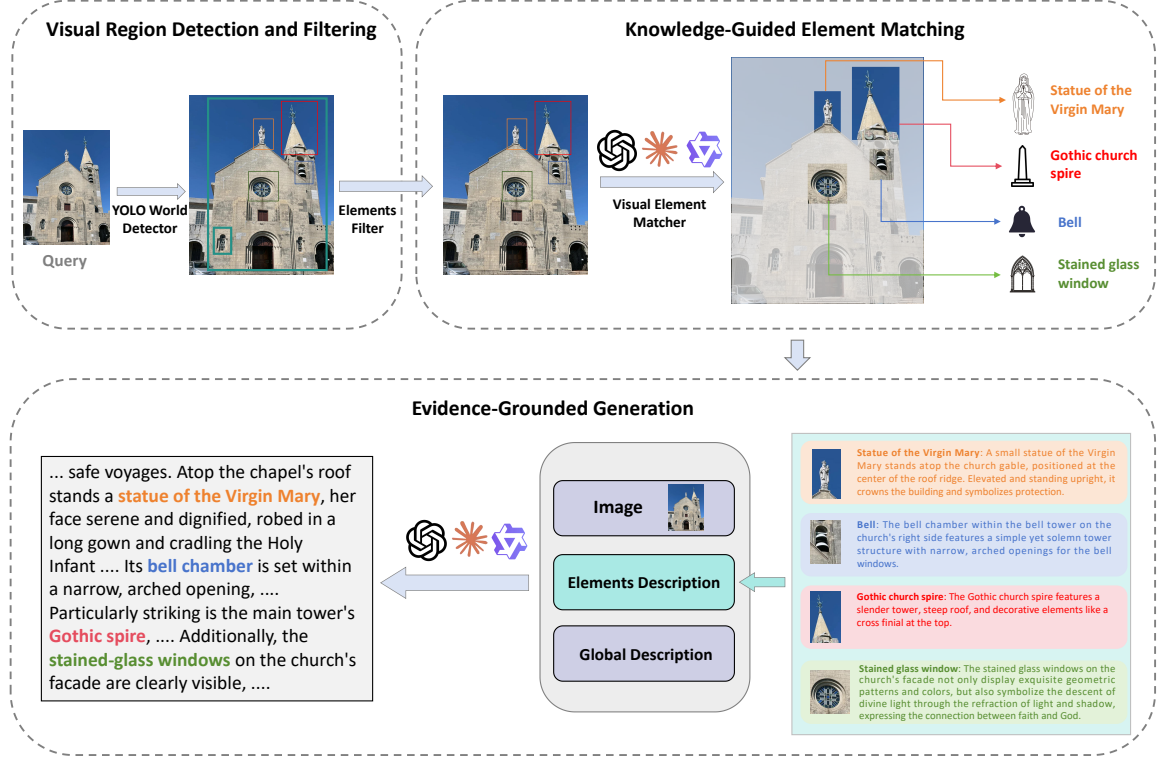


Figure 3: The overview of our proposed GranuRAG framework.

4.2 Knowledge-Guided Element Matching

Given the detected regions $\mathcal{B}(I)$ and a candidate element set E with associated appearance descriptions $\{a_e\}_{e \in E}$, this stage determines which architectural elements are actually visible in the image. We formulate this as a multimodal matching problem where each cropped region is compared against element descriptions from the knowledge corpus.

An MLLM receives each annotated bounding box specifically for an image alongside all appearance descriptions, which specify visual attributes such as shape, material, and stylistic features. For each detected region b_k , the model identifies the best-matching element by comparing observed visual characteristics against documented descriptions:

$$e_k = M_\phi(I, b_k, \{(e, a_e)\}_{e \in E}) \in E \cup \{\emptyset\} \quad (3)$$

where M_ϕ is the MLLM parameterized by ϕ , and the output \emptyset indicates that no candidate description is sufficiently consistent with the region, in which case the region is discarded. This design combines the spatial localization capability of the detector with the fine-grained semantic discrimination ability of the MLLM, enabling reliable identification even among visually similar elements.

The matching process yields a grounded element

set $\hat{E}(I) = \{e_k \mid e_k \neq \emptyset\}$ containing only elements with confirmed visual evidence. By explicitly filtering out undetected elements, we establish a principled boundary between what the system observes and what it knows, reducing the risk of hallucinating information about absent components.

4.3 Evidence-Grounded Generation

The final stage synthesizes a coherent interpretation by conditioning on both the annotated image and retrieved knowledge for matched elements. For each $e \in \hat{E}(I)$, we retrieve its expert-written description d_e from the knowledge corpus, which provides factual details including historical background, symbolic meaning, and architectural significance. We prepend global metadata m containing landmark name, architectural style, and historical period to contextualize the element-level information:

$$\mathcal{C}(I) = [m] \oplus [(e_i, d_{e_i})]_{e_i \in \hat{E}(I)} \quad (4)$$

The generator then produces the final output conditioned on this hierarchical evidence:

$$y = G_\theta(I, \mathcal{C}(I) \mid \Omega(\hat{E}(I))) \quad (5)$$

where the generation prompt Ω instructs the model to describe only elements confirmed in $\hat{E}(I)$ and ground all factual claims to retrieved descriptions.

Model	ROUGE-L	BERT-F1	LLM
Qwen3-VL-8B (A)	10.88	40.04	52.90
Qwen3-VL-8B (B)	12.23	44.46	61.17
Qwen3-VL-8B (C)	18.82	46.49	65.00
Qwen3-VL-8B [†] (A)	28.61	45.38	56.90
Qwen3-VL-8B [†] (B)	31.96	46.05	63.20
Qwen3-VL-8B [†] (C)	35.74	46.96	70.24
Qwen-VL-Max (A)	23.34	40.53	54.70
Qwen-VL-Max (B)	30.29	48.39	74.10
Qwen-VL-Max (C)	32.01	51.60	83.90
GPT-4.1-Mini (A)	18.81	40.94	56.30
GPT-4.1-Mini (B)	19.55	42.79	68.90
GPT-4.1-Mini (C)	22.72	44.04	75.10
GPT-4o (A)	12.83	39.95	52.70
GPT-4o (B)	18.41	42.09	63.97
GPT-4o (C)	19.16	43.43	75.40
Claude-3.5-Sonnet (A)	15.11	37.26	54.50
Claude-3.5-Sonnet (B)	16.39	39.54	66.70
Claude-3.5-Sonnet (C)	16.36	41.27	80.20

Table 3: The main results of our method on four different LLMs under three settings: (A) Baseline, (B) CoT, and (C) GranuRAG. [†] means the fine-tuned LLM. **Bold** means the best results for each LLM.

Such a design ensures that each claim in the output traces back to a detected visual region and a retrieved knowledge snippet. Such traceability facilitates systematic error diagnosis: missing information indicates detection failure, incorrect facts suggest retrieval error, and unsupported claims reveal generation hallucination.

5 Experiment

5.1 Setup

Experiments are conducted on the GranuVistaVQA dataset as described in section 3. We use the official APIs of **Qwen3-VL-8B**, **Qwen-VL-Max**, **GPT-4o**, **GPT-4.1-Mini**, and **claude-3.5-sonnet** (Bai et al., 2023; OpenAI, 2024; Anthropic, 2024). For the element detector, we employ the open-vocabulary YOLO-World-XL¹ D_ψ with a fixed confidence threshold. We use the evaluation metrics following Section 3.3. More implementation details will be shown in Appendix A.1.

5.2 Main Results

We evaluate GranuRAG across six state-of-the-art MLLMs under three settings: (A) Baseline, where the generator observes the image and the full noisy candidate set E_{all} ; (B) CoT, which augments (A) with chain-of-thought prompting (Wei et al., 2022)

¹<https://replicate.com/franz-biz/yolo-world-xl>




Figure 4: Ablation on visual presentation and element filtering. T3 and T4 use grounded subset $\hat{E}(I)$; T1 and T2 use full candidate set E_{all} .

Variant	ROUGE-L	BERT-F1	LLM
Text Only (Gold E_{gold})	32.05	46.32	72.10
Image + All E_{all}	29.47	45.57	68.50
Image + Chosen $\hat{E}(I)$	31.01	48.60	74.60

Table 4: Ablation on visual modality and knowledge relevance. **Bold** means the best results.

to encourage structured reasoning; and (C) GranuRAG, where the generator receives only the grounded element subset $\hat{E}(I)$ from our two-stage pipeline. For both the fine-tuned models (denoted by [†]) and Setting B, we synthesize the thinking process using powerful LLM and validate them through manual inspection (see Appendix A.3 for details).

Table 3 presents the results. Setting C consistently outperforms both baselines across almost all LLMs and metrics, demonstrating that explicit element alignment yields robust improvements regardless of backbone capacity. While CoT prompting (Setting B) moderately improves over the noisy baseline, it does not close the gap to our full pipeline, suggesting that reasoning-time prompting alone cannot fully compensate for irrelevant knowledge. Fine-tuning further amplifies these gains: Qwen3-VL-8B[†] (C) achieves 35.74 ROUGE-L, 46.96 BERT-F1, and 70.24 LLM score, substantially outperforming its zero-shot counterpart. Notably, even closed-source LLMs benefit markedly from grounded evidence, with Setting C improving LLM scores by 22.7, 18.8, and 25.7 points over Setting A, respectively. These results confirm that our GranuRAG framework provides complementary value beyond model scale or reasoning prompts.

5.3 Ablation Studies

We conduct ablation studies to isolate the contributions of visual evidence presentation and knowl-

Variant	ROUGE-L	BERT-F1	LLM
Baseline	23.34	40.53	54.70
No detector (LLM-only)	11.11	38.13	57.65
Grounding DINO	27.56	41.83	72.15
YOLO-World (Ours)	32.01	51.60	83.90

Table 5: Ablation on the detector component. **Bold** means the best performance.

edge selection. First, we examine four configurations varying visual input and element sets: (T1) raw image with all candidates E_{all} , (T2) box-annotated image with E_{all} , (T3) raw image with grounded subset $\hat{E}(I)$, and (T4) box-annotated image with $\hat{E}(I)$. As shown in Figure 4, T3 and T4 consistently outperform T1 and T2, confirming that filtering noisy candidates is crucial. T4 achieves the highest LLM score, indicating that combining localized visual cues and grounded textual evidence yields optimal performance.

Table 4 further contrasts three variants: text-only with gold elements E_{gold} , image with all candidates, and image with grounded subset. While text-only provides a strong upper bound, adding the image without filtering (Image + All) paradoxically degrades LLM score, suggesting that irrelevant visual-textual mismatches introduce noise. In contrast, pairing the image with the grounded subset substantially recovers text-only performance and exceeds it on most metrics, demonstrating that selective retrieval enables effective multimodal integration.

We further ablate the detectors to verify that our design choices are well-motivated rather than arbitrary. As shown in Table 5, replacing YOLO-World with Grounding DINO (Liu et al., 2024) still outperforms both the baseline and the approach of directly extracting relevant elements with LLMs, showing that the detect-then-match paradigm is robust across detectors.

6 Analysis

6.1 Generalization Analysis

To assess generalization beyond the fine-tuning distribution, we compare performance on in-domain (ID) and out-of-domain (OOD) test data. ID samples appear in training with different viewpoints, while OOD samples are entirely unseen during training. As shown in Figure 5, ID samples consistently achieve higher scores, reflecting expected distributional advantages. Importantly, the absolute improvements from GranuRAG on OOD samples




Figure 5: Performance comparison on in-domain and OOD data across three methods.

Method	ROUGE-L	BERT-F1	LLM
Baseline	23.79	40.83	56.40
Embedding Retrieval	29.47	45.57	63.45
RAVQA(PreFLMR)	21.27	42.60	69.24
VisRAG	24.06	43.35	68.06
GranuRAG (Ours)	32.27	52.19	79.30

Table 6: The results of comparing our method with different RAG strategies. **Bold** means the best result.

are comparable to, and sometimes exceed, those on ID samples across all three metrics. Moreover, the performance gap between ID and OOD data does not widen as we progress from Base to CoT to GranuRAG. This suggests that our grounding mechanism improves reasoning quality rather than exploiting memorized training instances. The consistent gains on OOD data confirm that GranuRAG generalizes well to novel visual content.

6.2 Comparison with Retrieval Baselines

We evaluate the effectiveness of our GranuRAG framework by fixing the generator backbone (Qwen-VL-Max) and comparing three strategies for constructing the evidence that conditions generation. The global baseline provides the generator with descriptions of the full candidate set E without visual grounding, requiring it to implicitly filter irrelevant elements. Embedding Retrieval replaces our second-stage matching with CLIP-based dense retrieval² (Radford et al., 2021): for each detected region, we retrieve the top-1 most similar element description from a vector database. Besides, we also compare GranuRAG with two strong multimodal RAG frameworks, which is RAVQA (Lin et al., 2023)³ and VisRAG 2.0 (Sun et al., 2025)⁴,

²<https://huggingface.co/openai/clip-vit-large-patch14>

³<https://github.com/LinWeizheDragon/Retrieval-Augmented-Visual-Question-Answering>

⁴<https://github.com/openbmb/visrag>

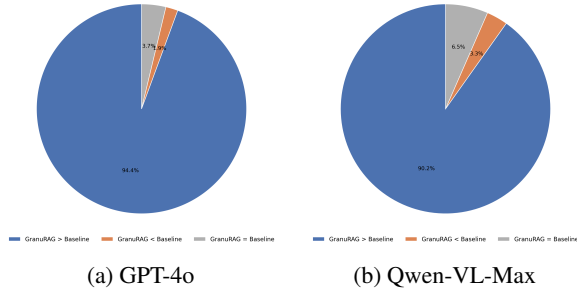


Figure 6: Distribution of answer quality when both LLMs successfully extract relevant elements.

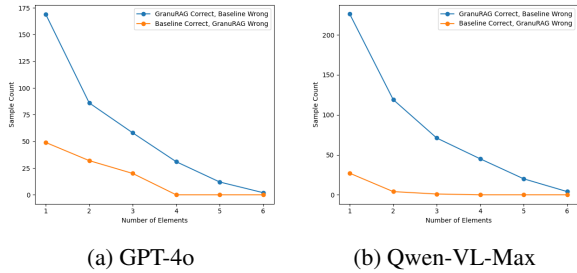


Figure 7: Extraction accuracy comparison across images with different numbers of elements.

using Qwen-VL-Max model.

Table 6 shows that GranuRAG substantially outperforms both baselines across all metrics. The global baseline achieves the lowest scores, struggling to identify relevant elements from the full candidate set. Embedding Retrieval improves performance but still lags behind our approach. GranuRAG achieves the best results with particularly large gains in LLM Score (+15.85 over CLIP), demonstrating that LLM-based semantic matching is more effective than embedding similarity for fine-grained element recognition.

6.3 Error Analysis

To understand the source of improvements, we conduct error analysis on two LLMs from two aspects. First, we compare answer quality when both methods extract the correct elements. As shown in Figure 6, GranuRAG produces better answers in 94.4% and 90.2% of cases for the two models respectively, while the baseline wins in less than 5%. This indicates that our fine-grained element representations provide more accurate information for generation, even when extraction is equally successful.

Second, we analyze extraction accuracy by examining cases where only one method identifies the correct elements. Figure 7 shows that GranuRAG succeeds where the baseline fails far more often

	win	loss
GPT-4o	82.22%	17.78%
Qwen-VL-Max	91.11%	8.89%

Table 7: The average win and loss rate of human evaluation. **Bold** means the best results.

than the reverse, especially for images with fewer elements. This gap narrows as element count increases, reflecting the growing difficulty of fine-grained reasoning in complex images. Overall, these results confirm that GranuRAG improves both extraction accuracy and generation quality.

6.4 Human Evaluation

To complement automatic metrics, we conduct human evaluation comparing our method against the baseline. We sample 20 questions for two models individually (GPT-4o and Qwen-VL-Max), yielding 40 pairwise comparisons. Three graduate students in computer science independently assess baseline and GranuRAG outputs according to four criteria: *coverage* of relevant scenic elements, *accuracy* of factual details, *fluency* of language, and *acceptability* as tourist guidance. Annotators select the superior output in each pair without knowing which system produced it.

Table 7 reports the aggregated win rates. GranuRAG achieves 82.22% preference over the baseline for GPT-4o and 91.11% for Qwen-VL-Max, confirming that grounded retrieval yields explanations that human evaluators judge more comprehensive and trustworthy. To assess annotation reliability, we compute Fleiss’ Kappa across the three raters, obtaining $\kappa = 0.712$ (SE = 0.105, $z = 6.75$, $p < 0.001$), which indicates substantial agreement (Landis and Koch, 1977) and validates the consistency of our evaluation protocol.

6.5 Visualization

To examine how GranuRAG influences visual focus, we fine-tune three adapters (Base, CoT, and GranuRAG) on Qwen3-VL-8B and compare their attention weights on the same samples. Specifically, we compute the pixel-level difference $\Delta A = A_{\text{GranuRAG}} - A_{\text{baseline}}$ by averaging attention weights across all generated tokens, then visualize the top-10 regions with the largest positive shifts (red + markers) and negative shifts (blue - markers). As shown in Figures 8a and 8b, red markers concentrate on semantically relevant scenic ele-

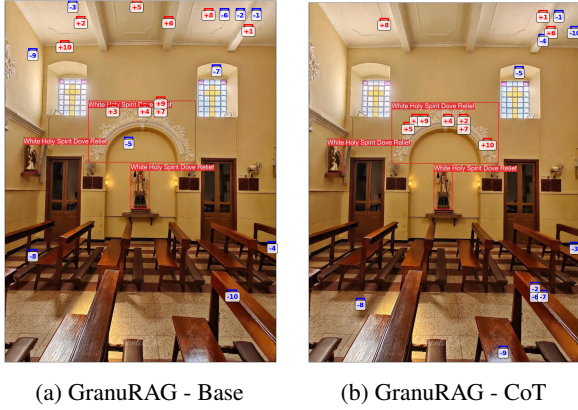


Figure 8: Top 10 attention difference regions. **Red**: higher attention in our method; **Blue**: higher attention in the base/CoT method.

ments, such as the “White Holy Spirit Dove Relief” sculptures and decorative arches, while blue markers scatter across generic background regions like ceilings and walls. This pattern confirms that our grounding mechanism systematically reallocates attention toward knowledge-relevant regions and away from distractors. More case studies are provided in Appendix A.5.

7 Conclusion

In this paper, we introduced GranuVistaVQA, a benchmark that reflects real-world visual reasoning challenges. Our GranuRAG framework explicitly grounds visual elements before retrieving multi-granularity evidence, enabling transparent attribution and consistently outperforming strong baselines. Experiments demonstrate that organizing evidence hierarchically from atomic visual elements, rather than relying on implicit attention within coarse retrievals, is essential for verifiable multimodal reasoning.

Limitations

While GranuVistaVQA advances multimodal RAG evaluation, several limitations should be noted. First, our benchmark focuses on landmark images where visual elements are typically buildings or monuments. This may not cover all visual reasoning scenarios, such as understanding abstract diagrams or everyday indoor scenes. Second, regarding efficiency, our method takes approximately 3.5 seconds per sample compared to 2 seconds for baseline approaches, nearly twice as long, due to the additional detection and multi-granularity retrieval steps, which may limit applicability in

time-sensitive scenarios. Third, our current implementation for the multi-granularity RAG scenario relies solely on traditional supervised fine-tuning (SFT). Future work will explore a broader range of parameter-efficient fine-tuning methods, including diverse LoRA variants (Zhang et al., 2023; Chen et al., 2025b) as well as recently popular neuron-level fine-tuning techniques (Xu et al., 2025a; Chen et al., 2026).

Acknowledgments

This work was supported in part by the Science and Technology Development Fund of Macau SAR (Grant Nos. FDCT/0007/2024/AKP, EF2024-00185-FST), the UM and UMDF (Grant Nos. MYRG-GRG2024-00165-FST-UMDF, MYRG-GRG2025-00236-FST), the Tencent AI Lab Rhino-Bird Research Program (Grant No. EF2023-00151-FST), the Stanley Ho Medical Development Foundation (Grant No. SHMDF-AI/2026/001), and the National Natural Science Foundation of China (Grant No. 62266013).

References

- Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katie Millican, Malcolm Reynolds, Roman Ring, Eliza Rutherford, Serkan Cabi, Tengda Han, Zhitao Gong, Sina Samangooei, Marianne Monteiro, Jacob Menick, Sebastian Borgeaud, and 8 others. 2022. [Flamingo: a visual language model for few-shot learning](#). *Preprint*, arXiv:2204.14198.
- Anthropic. 2024. [Claude 3.5 sonnet model card addendum](#).
- Anthropic. 2025. [Introducing Claude Haiku 4.5](#). Accessed: 2025-12-30.
- Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. 2023. [Qwen-vl: A frontier large vision-language model with versatile abilities](#). *CoRR*, abs/2308.12966.
- Zechen Bai, Pichao Wang, Tianjun Xiao, Tong He, Zongbo Han, Zheng Zhang, and Mike Zheng Shou. 2024. [Hallucination of multimodal large language models: A survey](#). *arXiv preprint arXiv:2404.18930*.
- Guanhua Chen, Yutong Yao, Lidia S. Chao, Xuebo Liu, and Derek F. Wong. 2025a. [SGIC: A self-guided iterative calibration framework for RAG](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, ACL 2025, Vienna, Austria, July 27 - August 1, 2025, pages 28357–28370. Association for Computational Linguistics.

- Guanhua Chen, Yutong Yao, Ci-Jun Gao, Lidia S. Chao, Feng Wan, and Derek F. Wong. 2025b. [Not all lora parameters are essential: Insights on inference necessity](#). *CoRR*, abs/2503.23360.
- Wenhu Chen, Hexiang Hu, Xi Chen, Pat Verga, and William W. Cohen. 2022. [Murag: Multimodal retrieval-augmented generator for open question answering over images and text](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7-11, 2022*, pages 5558–5570. Association for Computational Linguistics.
- Xin Chen, Junchao Wu, Shu Yang, Runzhe Zhan, Zeyu Wu, Min Yang, Shujian Huang, Lidia S. Chao, and Derek F. Wong. 2026. [Neuron-aware data selection in instruction tuning for large language models](#). *CoRR*, abs/2603.13201.
- Yang Chen, Hexiang Hu, Yi Luan, Haitian Sun, Soravit Changpinyo, Alan Ritter, and Ming-Wei Chang. 2023. [Can pre-trained vision and language models answer visual information-seeking questions?](#) *arXiv preprint arXiv:2302.11713*.
- Tianheng Cheng, Lin Song, Yixiao Ge, Wenyu Liu, Xinggang Wang, and Ying Shan. 2024. [Yolo-world: Real-time open-vocabulary object detection](#). In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16901–16911.
- Yew Ken Chia, Liying Cheng, Hou Pong Chan, Chaoyun Liu, Maojia Song, Sharifah Mahani Aljunied, Soujanya Poria, and Lidong Bing. 2024. [M-longdoc: A benchmark for multimodal super-long document understanding and a retrieval-aware tuning framework](#). *arXiv preprint arXiv:2411.06176*.
- Jaemin Cho, Debanjan Mahata, Ozan Irsoy, Yujie He, and Mohit Bansal. 2024. [M3docrag: Multi-modal retrieval is what you need for multi-page multi-document understanding](#). *arXiv preprint arXiv:2411.04952*.
- Wanqing Cui, Keping Bi, Jiafeng Guo, and Xueqi Cheng. 2024. [More: Multi-modal retrieval augmented generative commonsense reasoning](#). *Preprint*, arXiv:2402.13625.
- Kuicai Dong, Yujing Chang, Xin Deik Goh, Dexun Li, Ruiming Tang, and Yong Liu. 2025. [Mmdocir: Benchmarking multi-modal retrieval for long documents](#). *arXiv preprint arXiv:2501.08828*.
- Junfeng Fang, Zongze Bi, Ruipeng Wang, Houcheng Jiang, Yuan Gao, Kun Wang, An Zhang, Jie Shi, Xiang Wang, and Tat-Seng Chua. 2024. [Towards neuron attributions in multimodal large language models](#). In *Advances in Neural Information Processing Systems*.
- Qiushan Guo, Shalini De Mello, Hongxu Yin, Wonmin Byeon, Ka Chun Cheung, Yizhou Yu, Ping Luo, and Sifei Liu. 2024. [Regionopt: Towards region understanding vision language model](#). *arXiv preprint arXiv:2403.02330*.
- 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.
- J Richard Landis and Gary G Koch. 1977. [The measurement of observer agreement for categorical data](#). *Biometrics*, 33(1):159–174.
- Paul Lerner, Olivier Ferret, Camille Guinaudeau, Hervé Le Borgne, Romaric Besançon, Jose G. Moreno, and Jesús Lovón Melgarejo. 2022. [Viquae, a dataset for knowledge-based visual question answering about named entities](#). In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '22*, page 3108–3120, New York, NY, USA. Association for Computing Machinery.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, and 1 others. 2020. [Retrieval-augmented generation for knowledge-intensive nlp tasks](#). *Advances in neural information processing systems*, 33:9459–9474.
- Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. 2023. [Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models](#). *Preprint*, arXiv:2301.12597.
- Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Weizhe Lin, Jinghong Chen, Jingbiao Mei, Alexandru Coca, and Bill Byrne. 2023. [Fine-grained late-interaction multi-modal retrieval for retrieval augmented visual question answering](#). *Advances in Neural Information Processing Systems*, 36:22820–22840.
- Weizhe Lin, Jingbiao Mei, Jinghong Chen, and Bill Byrne. 2024. [Peflrmr: Scaling up fine-grained late-interaction multi-modal retrievers](#). *Preprint*, arXiv:2402.08327.
- Pei Liu, Xin Liu, Ruoyu Yao, Junming Liu, Siyuan Meng, Ding Wang, and Jun Ma. 2025. [Hm-rag: Hierarchical multi-agent multimodal retrieval augmented generation](#). In *Proceedings of the 33rd ACM International Conference on Multimedia*, pages 2781–2790.
- Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Qing Jiang, Chunyuan Li, Jianwei Yang, Hang Su, Jun Zhu, and Lei Zhang. 2024. [Grounding DINO: marrying DINO with grounded pre-training for open-set object detection](#). In *Computer Vision - ECCV 2024 - 18th European Conference, Milan, Italy, September 29-October 4, 2024*,

- Proceedings, Part XLVII*, Lecture Notes in Computer Science, pages 38–55. Springer.
- Zhenghao Liu, Chenyan Xiong, Yuanhuiyi Lv, Zhiyuan Liu, and Ge Yu. 2023. [Universal vision-language dense retrieval: Learning a unified representation space for multi-modal retrieval](#). *Preprint*, arXiv:2209.00179.
- Xueguang Ma, Shengyao Zhuang, Bevan Koopman, Guido Zuccon, Wenhui Chen, and Jimmy Lin. 2024a. [Visa: Retrieval augmented generation with visual source attribution](#). *Preprint*, arXiv:2412.14457.
- Yubo Ma, Yuhang Zang, Liangyu Chen, Meiqi Chen, Yizhu Jiao, Xinze Li, Xinyuan Lu, Ziyu Liu, Yan Ma, Xiaoyi Dong, and 1 others. 2024b. [Mmlongbench-doc: Benchmarking long-context document understanding with visualizations](#). *Advances in Neural Information Processing Systems*, 37:95963–96010.
- Kenneth Marino, Mohammad Rastegari, Ali Farhadi, and Roozbeh Mottaghi. 2019. [Ok-vqa: A visual question answering benchmark requiring external knowledge](#). *Preprint*, arXiv:1906.00067.
- Thomas Mensink, Jasper Uijlings, Lluís Castrejon, Arushi Goel, Felipe Cadar, Howard Zhou, Fei Sha, André Araujo, and Vittorio Ferrari. 2023. [Encyclopedic vqa: Visual questions about detailed properties of fine-grained categories](#). *Preprint*, arXiv:2306.09224.
- OpenAI. 2024. [Gpt-4o system card](#). *Preprint*, arXiv:2410.21276.
- Zhiliang Peng, Wenhui Wang, Li Dong, Yaru Hao, Shaohan Huang, Shuming Ma, and Furu Wei. 2023. [Kosmos-2: Grounding multimodal large language models to the world](#). *arXiv preprint arXiv:2306.14824*.
- Shraman Pramanick, Rama Chellappa, and Subhashini Venugopalan. 2024. [Spiqa: A dataset for multimodal question answering on scientific papers](#). *Advances in Neural Information Processing Systems*, 37:118807–118833.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, and 1 others. 2021. [Learning transferable visual models from natural language supervision](#). In *International conference on machine learning*, pages 8748–8763. PmLR.
- Dustin Schwenk, Apoorv Khandelwal, Christopher Clark, Kenneth Marino, and Roozbeh Mottaghi. 2022. [A-okvqa: A benchmark for visual question answering using world knowledge](#). *Preprint*, arXiv:2206.01718.
- Yubo Sun, Chunyi Peng, Yukun Yan, Shi Yu, Zhenghao Liu, Chi Chen, Zhiyuan Liu, and Maosong Sun. 2025. [Visrag 2.0: Evidence-guided multi-image reasoning in visual retrieval-augmented generation](#). *CoRR*, abs/2510.09733.
- Gemini Team and Google. 2025. [Gemini 2.5: Pushing the frontier with advanced reasoning, multi-modal capabilities, and long context](#). *arXiv preprint arXiv:2507.06261*.
- Qwen Team. 2025. [Qwen3 technical report](#). *Preprint*, arXiv:2505.09388.
- Yang Tian, Fan Liu, Jingyuan Zhang, Victoria W., Yupeng Hu, and Liqiang Nie. 2025. [Core-mmrag: Cross-source knowledge reconciliation for multi-modal rag](#). *Preprint*, arXiv:2506.02544.
- Dmitry Ustalov, Nikita Pavlichenko, Sergey Koshelev, Daniil Likhobaba, and Alisa Smirnova. 2023. [Toloka visual question answering benchmark](#). *Preprint*, arXiv:2309.16511.
- Navve Wasserman, Roi Pony, Oshri Naparstek, Adi Raz Goldfarb, Eli Schwartz, Udi Barzelay, and Leonid Karlinsky. 2025. [Real-mm-rag: A real-world multi-modal retrieval benchmark](#). *arXiv preprint arXiv:2502.12342*.
- 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 its reasoning in large language models](#). *Advances in neural information processing systems*, 35:24824–24837.
- Suyang Xi, Chenxi Yang, Hong Ding, Yiqing Ni, Catherine C. Liu, Yunhao Liu, and Chengqi Zhang. 2025. [Taming a retrieval framework to read images in humanlike manner for augmenting generation of mllms](#). *Preprint*, arXiv:2510.10426.
- Haoyun Xu, Runzhe Zhan, Yingpeng Ma, Derek F. Wong, and Lidia S. Chao. 2025a. [Let’s focus on neuron: Neuron-level supervised fine-tuning for large language model](#). In *Proceedings of the 31st International Conference on Computational Linguistics, COLING 2025, Abu Dhabi, UAE, January 19-24, 2025*, pages 9393–9406. Association for Computational Linguistics.
- Mingjun Xu, Zehui Wang, Hengxing Cai, and Renxin Zhong. 2025b. [A multi-granularity retrieval framework for visually-rich documents](#). *Preprint*, arXiv:2505.01457.
- Yibin Yan and Weidi Xie. 2024. [Echosight: Advancing visual-language models with wiki knowledge](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, page 1538–1551. Association for Computational Linguistics.
- Wei Yang, Jingjing Fu, Rui Wang, Jinyu Wang, Lei Song, and Jiang Bian. 2025. [Omgm: Orchestrate multiple granularities and modalities for efficient multimodal retrieval](#). *Preprint*, arXiv:2505.07879.
- Shi Yu, Chaoyue Tang, Bokai Xu, Junbo Cui, Junhao Ran, Yukun Yan, Zhenghao Liu, Shuo Wang, Xu Han, Zhiyuan Liu, and Maosong Sun. 2025a. [Visrag: Vision-based retrieval-augmented generation on](#)

- multi-modality documents. In *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*. OpenReview.net.
- Shi Yu, Chaoyue Tang, Bokai Xu, Junbo Cui, Junhao Ran, Yukun Yan, Zhenghao Liu, Shuo Wang, Xu Han, Zhiyuan Liu, and Maosong Sun. 2025b. *Visrag: Vision-based retrieval-augmented generation on multi-modality documents*. *Preprint*, arXiv:2410.10594.
- Wenhan Yu, Wang Chen, Guanqiang Qi, Weikang Li, Yang Li, Lei Sha, Deguo Xia, and Jizhou Huang. 2025c. *Bbox docvqa: A large scale bounding box grounded dataset for enhancing reasoning in document visual question answer*. *Preprint*, arXiv:2511.15090.
- Haotian Zhang, Haoxuan You, Philipp Dufter, Bowen Zhang, Chen Chen, Hong-You Chen, Tsu-Jui Fu, William Yang Wang, Shih-Fu Chang, Zhe Gan, and Yinfei Yang. 2024a. *Ferret-v2: An improved baseline for referring and grounding with large language models*. *Preprint*, arXiv:2404.07973.
- Qingru Zhang, Minshuo Chen, Alexander Bukharin, Nikos Karampatziakis, Pengcheng He, Yu Cheng, Weizhu Chen, and Tuo Zhao. 2023. *Adalora: Adaptive budget allocation for parameter-efficient fine-tuning*. *arXiv preprint arXiv:2303.10512*.
- Tao Zhang, Ziqi Zhang, Zongyang Ma, Yuxin Chen, Zhongang Qi, Chunfeng Yuan, Bing Li, Junfu Pu, Yuxuan Zhao, Zehua Xie, Jin Ma, Ying Shan, and Weiming Hu. 2024b. *mr²ag: Multi-modal retrieval-reflection-augmented generation for knowledge-based vqa*. *Preprint*, arXiv:2411.15041.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. *Bertscore: Evaluating text generation with bert*. *Preprint*, arXiv:1904.09675.
- Yichi Zhang, Ziqiao Ma, Xiaofeng Gao, Suhaila Shakiah, Qiaozi Gao, and Joyce Chai. 2024c. *Groundhog: Grounding large language models to holistic segmentation*. *Preprint*, arXiv:2402.16846.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P. Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. *Judging llm-as-a-judge with mt-bench and chatbot arena*. *Preprint*, arXiv:2306.05685.
- Junjie Zhou, Zheng Liu, Shitao Xiao, Bo Zhao, and Yongping Xiong. 2024a. *Vista: Visualized text embedding for universal multi-modal retrieval*. *Preprint*, arXiv:2406.04292.
- Tianshuo Zhou, Sen Mei, Xinze Li, Zhenghao Liu, Chenyan Xiong, Zhiyuan Liu, Yu Gu, and Ge Yu. 2024b. *Marvel: Unlocking the multi-modal capability of dense retrieval via visual module plugin*. *Preprint*, arXiv:2310.14037.
- Yongxin Zhu, Zhen Liu, Yukang Liang, Xin Li, Hao Liu, Changcun Bao, and Linli Xu. 2023. *Locate then generate: Bridging vision and language with bounding box for scene-text vqa*. *Preprint*, arXiv:2304.01603.

A Appendix

A.1 Implementation

For all evaluated MLLMs, we use a unified decoding setup with temperature set to 0.1 and maximum generation length limited to 600 tokens. This low-temperature setting reduces randomness and ensures stable, reproducible outputs during evaluation. In addition, we fine-tune Qwen3-VL-8B⁵ using the LLaMA-Factory⁶ framework on a single H800 GPU with setting LoRA (Hu et al., 2022) rank to 8. The sequence length is set to 4096, with a batch size of 4 and gradient accumulation steps of 4. We train for 3 epochs on the training set with a learning rate of 1e-4. The code and data will be available at <https://github.com/NLP2CT/GranuRAG>.

A.2 Sensitivity Analysis of Overlap Threshold

We conduct a sensitivity analysis on the overlap threshold used in our redundancy filtering step. Table 8 reports the results across six threshold values ranging from 70% to 100%. As shown, performance follows a clear and intuitive trend: lower thresholds (e.g., 70%) lead to over-removal of informative elements, which hurts the quality of the generated output, while higher thresholds (e.g., 90% and above) retain too much redundancy and also degrade performance. The best results are achieved at 80%, with a ROUGE-L of 32.01, BERT-F1 of 51.60, and LLM Score of 83.90. Importantly, performance remains consistently strong across the 75%–85% range, indicating that our method is not overly sensitive to this hyperparameter. We select 80% as our default threshold since it sits at the natural midpoint of this effective range and yields the best overall performance.

A.3 Synthesize Fine-tuned Data

For the fine-tuning task, we construct our test set by randomly selecting 20% of partial landmarks images (containing incomplete visual elements), with the remainder forming the training set. Critically, we ensure that approximately half of the scenic spots are entirely absent from the training set to evaluate the model’s robustness on unseen

⁵<https://huggingface.co/Qwen/Qwen3-VL-8B-Instruct>

⁶<https://github.com/hiyouga/LLaMA-Factory>

Threshold	ROUGE-L	BERT-F1	LLM Score
70%	23.69	39.35	68.00
75%	29.83	47.72	75.35
80%	32.01	51.60	83.90
85%	31.67	50.02	79.30
90%	29.84	48.69	76.58
100%	26.45	43.67	65.20

Table 8: Sensitivity analysis of the overlap threshold. **Bold** means the best performance.

locations. The related analysis is shown in section 6.1. We then employ Qwen3-VL-235B-A22B-Thinking (Team, 2025) to synthesize reasoning data for both sets by populating the prompt with three components: the landmark image, annotated visual elements with descriptions, and the target commentary text. We prompt the LLM to generate reasoning processes from input to target output using two approaches: its native chain-of-thought style and our proposed pipeline framework.

To ensure data quality, each synthesized reasoning process underwent automated scoring by LLM, with samples failing to fully meet our criteria being regenerated. When a sample failed validation three consecutive times, human annotators intervened to manually revise the reasoning process. Through this rigorous pipeline combining automated synthesis with quality control, we construct a high-quality dataset for fine-tuning LLMs to learn both standard CoT reasoning and our pipeline-based thinking process.

A.4 Attribution Evaluation

Beyond factual correctness, we also evaluate whether LLMs can trace their claims back to specific visual evidence. A description may sound plausible yet be entirely ungrounded in what the model actually perceives. We measure this with three metrics: Attribution Precision (AP), the fraction of cited elements that exist in the predicted visible set $\hat{E}(I)$; Attribution Recall (AR), the fraction of elements in $\hat{E}(I)$ that are actually cited in the output; and Unsupported Claim Rate (UCR), the fraction of output sentences not grounded by any element.

As shown in Table 9, all models improve consistently from Setting (A) to (C), confirming that structured evidence grounding helps anchor generated text to perceived visual content. Fine-tuning also brings clear gains: Qwen3-VL-8B[†] improves recall from 0.45 to 0.58 under Setting (C) compared to its base version. Among proprietary models, GPT-4o

Model	AP [↑]	AR [↑]	UCR [↓]
<i>Qwen3-VL-8B</i>			
Setting (A)	0.6997	0.2244	0.9417
Setting (B)	0.8153	0.2423	0.7814
Setting (C)	<u>0.9242</u>	<u>0.4467</u>	<u>0.6753</u>
<i>Qwen3-VL-8B[†]</i>			
Setting (A)	0.6500	0.3921	0.8392
Setting (B)	0.8420	0.4309	0.7513
Setting (C)	<u>0.9561</u>	<u>0.5847</u>	<u>0.6235</u>
<i>Qwen-VL-Max</i>			
Setting (A)	0.7675	0.3573	0.7779
Setting (B)	0.9474	0.4754	0.6279
Setting (C)	<u>0.9730</u>	<u>0.6194</u>	<u>0.5293</u>
<i>GPT-4.1-Mini</i>			
Setting (A)	0.6791	0.3339	0.8104
Setting (B)	0.9614	0.5844	0.7888
Setting (C)	<u>0.9865</u>	<u>0.6144</u>	<u>0.6482</u>
<i>GPT-4o</i>			
Setting (A)	0.6990	0.3562	0.8490
Setting (B)	0.9505	0.6097	0.7323
Setting (C)	0.9865	0.6620	0.6184
<i>Claude-3.5-Sonnet</i>			
Setting (A)	0.7616	0.2219	0.9414
Setting (B)	0.9479	0.2494	0.8125
Setting (C)	<u>0.9583</u>	<u>0.4297</u>	<u>0.6931</u>

Table 9: Attribution evaluation results. AP = Attribution Precision, AR = Attribution Recall, UCR = Unsupported Claim Rate. [†] denotes the fine-tuned variant. **Bold** marks the best score per metric. Underline means the best results for each LLM.

achieves the best overall profile, while Claude-3.5-Sonnet shows high precision but low recall, suggesting it cites carefully but sparingly. Notably, even the best configuration leaves over half of output sentences unsupported, indicating that faithful visual attribution remains an open challenge.

A.5 Case Study

Table 11 presents a color-coded case study that directly compares element-level grounding and description quality across methods. Each color family corresponds to one architectural element, with darker shades indicating more accurate descriptions. Specifically, **blue** highlights track the Christian emblem, **green** highlights track windows, and **pink** denotes the pediment. This visualization reveals clear differences: the baseline misses the

pediment entirely, CoT hallucinates non-existent twin towers, while GranuRAG stays faithful to the visible evidence and aligns closely with the ground truth.

From an element selection perspective, the baseline fails to identify the prominent pediment, resulting in incomplete structural coverage. CoT, though more detailed, introduces twin towers that do not exist in the image, a typical hallucination risk in CoT reasoning. GranuRAG correctly identifies all key elements including the pediment and Christian emblem while avoiding both omissions and hallucinations. In terms of generation quality, the shade progression reflects meaningful differences: for the Christian emblem (blue), GranuRAG provides fine-grained symbolic details (e.g., three nails, IHS inscription) that match the ground truth, while baseline and CoT remain superficial. For windows (green), GranuRAG delivers complete structural descriptions versus partial coverage by other methods. Most notably, the pediment (pink) appears only in GranuRAG’s output, demonstrating that our method improves both element recall and description accuracy.

B Dataset Construction Details

B.1 Data Schema

We organize knowledge at two granularities to support hierarchical retrieval:

Landmarks -level representation:

$$x_{\text{attr}} = (\text{meta}, E, ED, I^{\text{rep}}), \quad (6)$$

where $\text{meta} = \{\text{Landmarks}, \text{Summary}, \text{Style}\}$ contains high-level context; $E = \{e_1, \dots, e_k\}$ is the landmarks-specific element inventory; $ED : E \rightarrow \text{Paragraphs}$ maps each element to its detailed description; and I^{rep} is a representative image for visualization.

Image-level representation:

$$z = (I, a(I), E^{\text{gt}}(I)), \quad (7)$$

where I is an image, $a(I)$ denotes its source landmarks, and $E^{\text{gt}}(I) \subseteq E$ specifies ground-truth visible elements.

The structured JSON schema follows:

```
{
  "Landmarks": "[NAME]",
  "Summary": "[description]",
  "Style": "[style_notes]",
  "Elements": ["[elem_1]", ...,
```

```
  "[elem_k]"],
  "ElementDescriptions": {
    "[elem_i]": "[paragraph]",
    ...
  },
  "Image": "[image_path]"
}
```

B.2 Construction Pipeline

Stage 1: Data Collection We collect textual materials from encyclopedic sources including Wikipedia and Baidu Baike, as well as from official heritage and tourism portals such as the Macao World Heritage website⁷ and the Macao Government Tourism Office website.⁸ Images are sourced from Google Images, Rednote, and Trip.com using landmark names as search queries to ensure viewpoint diversity. All data are collected for non-commercial research purposes only. In addition, we perform data sanitization to remove privacy-sensitive and potentially harmful content, employing a hybrid approach that combines LLM-based filtering with manual spot-checking. All images are collected exclusively for non-commercial academic research under the principle of fair use.

We perform comprehensive data sanitization through a rigorous manual screening process. Trained annotators inspect each image individually following a strict exclusion protocol: any image containing the following elements is immediately discarded and not included in the dataset:

- Visible watermarks, platform logos, or copyright overlays
- Timestamps, date stamps, or camera metadata overlays
- Geolocation tags, map overlays, or address information
- License plates, vehicle identifiers, or transportation tags
- Identity documents, tickets, receipts, or personal contact info
- Recognizable human faces and personal information

This policy prioritizes privacy protection over data retention. After manual screening, we apply quality filtering to retain only images with resolution $\geq 512\text{px}$ and no visible compression artifacts. The final dataset contains 1,422 images that have passed

⁷<https://www.wh.mo/cn/>

⁸<https://www.macaotourism.gov.mo/zh-hant/sightseeing>

Model	P	R	F1
Qwen3-VL-8B	0.723	0.493	0.470
Qwen-VL-Max	0.850	0.822	0.807
GPT-4.1-Mini	0.801	0.715	0.731
GPT-4o	0.699	0.662	0.649
Claude-3.5-Sonnet	0.754	0.741	0.758

Table 10: Element detection results with normalization applied. **Bold** means the best result.

both privacy and quality checks, ensuring suitability for academic research purposes.

Stage 2: Knowledge Structuring Raw textual materials are consolidated into the normalized JSON schema through a four-step process. We first extract element phrases from authoritative descriptions, then apply normalization rules (Appendix B.3). Next, we generate element descriptions via LLM (Appendix B.4) and finally validate all outputs (Appendix B.5).

Stage 3: Image Annotation For each image I , we provide Qwen3-VL-Max (Team, 2025) with the element inventory E and descriptions ED as context. The model proposes candidate visible elements $\hat{E}(I) \subseteq E$, which human annotators then refine into final ground-truth labels $E^{\text{gt}}(I)$ following strict visibility guidelines detailed in Appendix B.6.

Stage 4: Indexing We embed all images using CLIP and L2-normalize the resulting vectors to 512 dimensions. A FAISS flat L2 index is built with aligned metadata that maps vector IDs to landmark names, image paths, and ground-truth element sets, enabling source attribution during retrieval.

B.3 Normalization Procedures

We apply three normalization strategies to ensure cross-landmark consistency. Language unification consolidates multilingual variants into canonical keys and enforces consistent orthography. Synonym collapsing merges semantically equivalent terms under lightweight taxonomy tags, such as grouping “bell tower” with “campanile,” “granite stone” with “granite,” and “carved motif” with “ornamental carving.” Cross-landmark alignment ensures that shared architectural features like “Baroque façade” receive identical identifiers across different landmarks. These steps are critical for two reasons: they prevent the element matching stage

(Section 4.2) from failing due to superficial lexical mismatches across knowledge sources, and they enable fair evaluation by ensuring the same element is not counted as two different entities, which would artificially deflate Precision and Recall. As shown in Table 10, even with normalization applied, element detection remains challenging, confirming that the task difficulty is genuine rather than an artifact of inconsistent naming.

B.4 Element Description Generation

For each element $e \in E$, we prompt Qwen3-VL-Max to generate a factual paragraph of one to three sentences describing its visual characteristics, cultural significance, and historical context. We enforce three constraints during generation: factuality requires all statements to be source-grounded with no unverifiable claims; scope constrains the content to the specific element rather than general landmark background; and tone maintains a neutral encyclopedic style that avoids promotional language. Generated paragraphs are stored under their corresponding keys in ED.

B.5 Validation and Quality Control

Quality control combines automatic checks with human review. Automatic validation verifies key consistency ($\text{keys}(\text{ED}) = E$), ensures all fields are non-empty, and confirms orthographic normalization. Human reviewers then perform three additional passes: factuality review spot-checks generated descriptions against source materials and regenerates entries containing unsupported claims; style filtering removes promotional language and unattributed superlatives; and terminological harmonization ensures consistent wording for shared elements across landmarks.

B.6 Ground-Truth Annotation Protocol

Human annotators refine LLM-proposed labels $\hat{E}(I)$ to produce $E^{\text{gt}}(I)$ following strict visibility rules. Visual identifiability requires that elements qualify only if they are confidently recognizable from image pixels alone, without relying on prior landmark knowledge; for instance, an external tower invisible in an interior courtyard photo cannot be labeled. For partial occlusion handling, partially occluded elements are included only when discriminative visual cues remain present, such as a half-visible ornamental scroll whose characteristic shape remains recognizable. Synonym resolution requires annotators to select canonical keys when

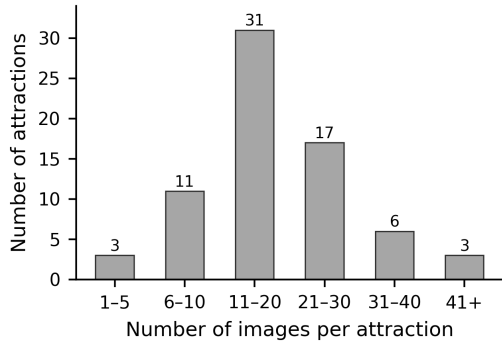


Figure 9: Distribution of image counts per landmark. The dataset shows a single dominant mode around 11–30 images (48 of 71 landmarks, 67.6%), with only a few landmarks providing more than 40 views.

multiple inventory terms refer to near-identical visual components. Uncertain cases undergo ambiguity escalation, where they are marked for expert review or excluded to avoid false positives.

The annotation workflow presents each image I alongside landmark metadata (meta, E , ED) and LLM-proposed candidates $\hat{E}(I)$. Annotators edit the proposals by adding missed but visible elements, removing hallucinated or non-verifiable elements, and correcting keys to canonical form. This protocol ensures ground-truth labels reflect genuine visual evidence, enabling fair evaluation while accounting for real-world photography ambiguities such as viewpoint limitations and occlusions.

B.7 Detailed coverage distribution

Beyond the aggregate statistics reported in the main text (mean 34%, median 29%), the image distribution statistic is shown in Figure 9, which reflects real-world photography patterns: most landmarks contain 11–30 images, while iconic locations exceed 40 images. The element coverage ratio $\text{Coverage}(I) = |E^{\text{gt}}(I)|/|E|$ shows considerable variation. The interquartile range spans 0.18 to 0.47, meaning half of all images show between 18% and 47% of available elements. The distribution exhibits notable tail behavior: 12% of images show fewer than 15% of elements due to extreme close-up framing, while 5% show more than 70% of elements in comprehensive panoramas. Stratifying by view type reveals that close-up views average 22% coverage, mid-range views average 36%, and panoramic views average 58%.

This granular breakdown confirms that multi-view aggregation is essential for our task. No single view type reliably captures complete element in-

ventories, which necessitates cross-view evidence fusion for comprehensive retrieval.

Input

Please carefully observe the visible details in the image corresponding to these elements, focusing your description on what is actually visible in the image. Below is the complete element description (ED) for the image.

ED:

Twin Towers: The top floor features symmetrical bell towers on both sides...

Windows: Smaller side windows with arched lintels and surrounding relief decorations...

Pediment: The central pediment is ...

Christian Emblem: Below the emblem are three nails commemorating...

Original Image**GT Text**

The facade of this church spans 24.6 meters in width. Its pediment displays quintessential Baroque style, adorned centrally with an exquisite Jesuit coat of arms. This christian emblem, rendered in symmetrical geometric patterns and relief carvings, serves as a common religious symbol in Jesuit churches. Directly below the coat of arms, three nails symbolize the historical events of Jesus' crucifixion. Above the coat of arms, the letters "IHS" are carved, derived from the initials of Jesus' Greek name "IHSOUS," commonly used as an abbreviation for Jesus Christ, signifying "Jesus is the Savior of mankind." The pilasters flanking the coat of arms are adorned with broken pediments. The central section of the church's second level features a large rectangular window with louvered panels, flanked by smaller arched windows. The lintels are decorated with scrollwork and complemented by carved details. This layout aligns with the Baroque architectural pursuit of symmetry and opulence.

Baseline**CoT****GranuRAG****Annotated Image****Annotated Image****Annotated Image****Generated Text**

The church's main facade presents a massive wall 24.6 meters wide, divided into three vertical and three horizontal sections. The central section features the Christian emblem, with three nails beneath it symbolizing the crucifixion of Jesus. The second level features three windows: the outer ones are smaller, adorned with arched lintels and surrounding relief carvings, while the central window is a rectangular shuttered opening. The overall design is simple yet solemn, showcasing the unique charm of religious art.

Generated Text

The exterior facade of this church displays a majestic Baroque style, spanning a width of 24.6 meters. The symmetrical bell towers at both ends feature glazed tile roofs. Three window openings in the left tower house large bronze bells, while the right tower also holds a bronze bell within its window openings. At the building's central pediment, the Jesuit emblem is carved, with the 'IHS' above symbolizing the name of Jesus. The second level features three windows: the outer ones are smaller, adorned with arched lintels and surrounding relief decorations, while the central window comprises a large shuttered opening. The overall design harmoniously blends religious artistry with functionality, conveying the building's solemnity and sacredness.

Generated Text

The church's main facade presents a massive wall 24.6 meters wide, divided into three vertical and three horizontal sections by two multi-tiered cornices and six fluted pilasters with multiple recesses. The pediment section of the building displays a quintessential Baroque style, adorned centrally with an exquisite Christian emblem. This emblem, rendered in symmetrical geometric patterns and relief carvings, serves as a common religious symbol found in Jesuit churches. Directly beneath the emblem, three nails symbolize the historical crucifixion of Jesus, one of Christianity's most significant emblems. Above the emblem, the letters "IHS" derive from the first three letters of Jesus' Greek name "IHSOUS," commonly used as an abbreviation for Jesus Christ, signifying "Jesus is the Savior of Humanity." The second level features a larger rectangular shuttered window at the center flanked by smaller arched windows. The lintels bear curved ornamentation complemented by sculpted details. This layout harmonizes with Baroque architecture's pursuit of symmetry and opulence while balancing natural light and aesthetic functionality.

Table 11: The case study example.
10491