

## Draft and Refine with Visual Experts

Sungheon Jeong  
UC Irvine

sungheoj@uci.edu

Ryozo Masukawa  
UC Irvine

rmasukaw@uci.edu

Jihong Park  
MOLOCO

qkrwlghddlek@gmail.com

Sanggeon Yun  
UC Irvine

sanggeoy@uci.edu

Wenjun Huang  
UC Irvine

wenj unh3@uci.edu

Hanning Chen  
UC Irvine

hanningc@uci.edu

Mahdi Imani  
Northeastern University

m.imani@northeastern.edu

Mohsen Imani  
UC Irvine

m.imani@uci.edu

### Abstract

While recent Large Vision–Language Models (LVLMs) exhibit impressive multimodal reasoning abilities, they often produce ungrounded, hallucinated responses by over-relying on linguistic priors rather than visual evidence. This critical limitation arises from the lack of a quantitative measure of how much these models actually rely on visual inputs during reasoning. We propose **Draft and Refine (DnR)**, an agent framework driven by a novel question-conditioned utilization metric. This metric quantifies the model’s actual reliance on visual evidence by first constructing a query-conditioned relevance map to localize question-specific evidence, and then assessing dependence through relevance-based probabilistic masking. Guided by this metric, the DnR agent refines its initial ‘draft’ through targeted feedback from external visual experts. Each expert’s output (e.g., boxes, masks) is rendered as visual cues on the image, and the LVLM is re-queried to select the response that yields the greatest improvement in utilization. This process strengthens visual grounding of predictions without re-training or architectural changes. Experiments across a broad range of VQA and captioning benchmarks demonstrate consistent accuracy gains and reduced hallucination. These results show that quantifying visual utilization provides a principled path for designing more interpretable and evidence-driven multimodal agent systems that effectively leverage visual experts.

### 1. Introduction

Large vision language models (LVLMs) have begun to interface with diverse visual tools and experts, enabling complex reasoning that combines perception and language [13, 27, 56, 64, 70–72]. However, existing methods mainly rely on language-driven control, prompting models to call experts based on chain-of-thought or textual confidence [53,

56, 70, 71]. Such mechanisms inherit the biases and unreliability of the language model itself and seldom account for how effectively visual information is actually utilized [4, 10]. Learning-based coordination frameworks, on the other hand, require costly and inflexible joint optimization across multiple experts and tasks [7, 26, 31, 32]. This raises a fundamental question: *can a vision–language model autonomously determine when and which visual expert to invoke—guided by its own perceptual needs rather than linguistic biases?*

To address this challenge, we reconsider expert coordination from the perspective of visual utilization. Our initial intuition is simple: LVLMs should be guided to more effectively utilize visual information when forming its predictions [28, 75]. However, not all visual cues are equally useful, because different tasks and questions require attention to different parts of the image [22, 40]. Therefore, the key is not to maximize visual dependence globally but to evaluate how well the model utilizes the regions relevant to the given situation [5, 50]. Once the model identifies which visual information is most critical, the selected expert can provide additional evidence to refine the model’s reasoning and improve its decision. Building on this idea, we develop a Draft and Refine (DnR) framework that measures visual utilization and uses it as a quantitative criterion for selecting and incorporating visual experts. This formulation enables adaptive expert choice grounded in the model’s actual perceptual behavior rather than relying on linguistic priors or heavy joint supervision.

Existing LVLMs can describe what they see but cannot determine which visual elements are truly important or why they matter for a given task [55, 66]. Because there is no explicit label or metric to assess such dependence, we first define a measurable criterion that quantifies how much the model relies on visual information through a relevance-based perturbation process. This enables us to evaluate visual utilization without additional supervision.

After computing this utilization, we incorporate external visual experts to complement perceptual gaps by providing additional evidence. To align their outputs with the model’s perception, we introduce a visual rendering mechanism that highlights essential regions while suppressing irrelevant content. This controllable rendering mechanism enables flexible integration of diverse visual experts into general-purpose multimodal reasoning systems without re-training.

Through this formulation, our approach offers a practical framework to quantify and leverage visual information within multimodal reasoning. It allows models to evaluate how effectively they use visual evidence and refine reasoning based on perceptual relevance rather than linguistic confidence. The rendering mechanism serves as a flexible interface for integrating diverse visual experts without re-training, enabling agent-style coordination that generalizes across domains. Although rendering styles and parameters require dataset- and model-specific tuning, the framework consistently improves performance across VQA and captioning benchmarks, showing a strong correlation with task accuracy and a substantial reduction in hallucination. Collectively, these findings suggest that measurable visual grounding offers a principled basis for developing interpretable and evidence-driven AI agent systems that leverage visual tools.

## 2. Related Work

**Multimodal Large Language Models.** Recent advances in multimodal large language models (MLLMs) have significantly expanded the capacity of language models to reason over visual inputs [3, 8, 29, 35, 38, 45, 47, 51, 73]. By coupling pretrained vision encoders with powerful language backbones, these systems achieve strong zero-shot performance on captioning, visual question answering, and general visual reasoning tasks [35, 61]. However, their reasoning process remains predominantly language-driven, relying heavily on internal linguistic priors rather than on visual grounding [4, 10, 61]. When the visual representation is coarse or misaligned, such models tend to hallucinate plausible yet unsupported content [4, 10, 61]. This imbalance between linguistic reasoning and perceptual understanding has been repeatedly observed across recent evaluations, which report limited grounding fidelity and reduced robustness on fine-grained or localized visual reasoning benchmarks [4, 35, 61]. These findings underscore the need for frameworks that can explicitly assess, verify, and refine visual evidence throughout the reasoning process [4, 10, 35].

**Tool-Augmented and Agentic LLMs.** Recent research has transformed large language models into agentic systems that autonomously plan, execute, and integrate external tools or APIs for reasoning and perception [36, 53, 56, 70, 71]. These approaches span programmatic reasoning agents

that compose visual functions via code execution [13, 64, 65], tool-calling coordinators that orchestrate multiple pre-trained experts in a modular pipeline [27, 34, 56, 70, 72], and multimodal planners that integrate vision encoders for grounded decision-making [3, 8, 29, 38, 45, 47, 51, 62]. Despite their success, these systems generally rely on pre-defined or heuristic tool invocation, often selecting experts based on the LLM’s textual outputs or embedding similarities rather than measurable visual evidence [7, 26, 27, 30–32, 34, 53, 56, 70, 72], and thus lack quantitative criteria for determining when visual assistance is necessary or which expert should be called. In contrast, our framework formulates expert invocation as a measurable decision process, guided by the model’s actual utilization of visual evidence, thereby enabling an adaptive and systematic bridge between perception and reasoning.

**Hallucination Mitigation and Visual Grounding.** Addressing hallucination and weak grounding has become a central challenge in both language and vision–language modeling [4, 10, 35]. Existing approaches attempt to improve factual or visual alignment through faithfulness-oriented supervision [16, 33], retrieval-augmented reasoning [6, 68], or uncertainty-driven selection and abstention strategies [15, 20, 21, 59]. While these methods reduce hallucination, they generally operate as post-hoc corrections or apply global feature amplification without explicitly conditioning on the query, often leading to over-attention on irrelevant regions and diluted reasoning fidelity [1, 74]. In contrast, our framework introduces a query-conditioned relevance map that directly links visual importance to question semantics, enabling selective and evidence-grounded refinement rather than uniform enhancement.

## 3. Draft and Refine with Visual Experts

The Draft-and-Refine (*DnR*) framework improves a LVLM by refining its draft response using specialized visual experts. Given an image  $x$  and a question  $q$ , the LVLM first drafts an initial response based on its existing visual understanding. *DnR* then analyzes the  $x$  to identify regions most relevant to  $q$ , evaluates how the model’s reasoning depends on them, and selectively engages the expert that provides complementary visual evidence. By integrating this evidence, *DnR* turns the LVLM’s passive description into an active, evidence-anchored reasoning process that steers its answers toward visually grounded decisions rather than incidental linguistic context.

### 3.1. Query-Conditioned Relevance Map

To assess how effectively the LVLM grounds its reasoning in visual evidence, we focus on identifying the regions of image  $x$  that truly matter for a given question  $q$ . Rather than assuming that all parts of the image contribute equally, we recognize that only certain regions serve as meaning-

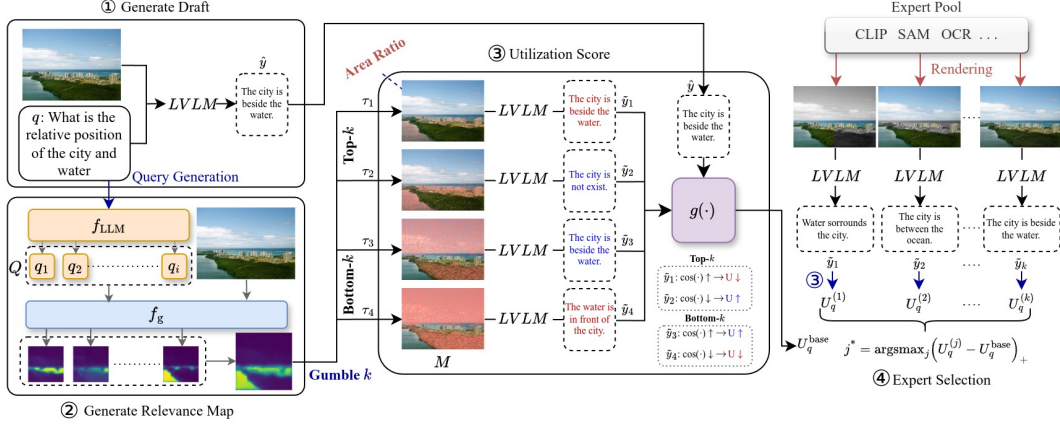


Figure 1. **Overview of the Draft-and-Refine (DnR) framework.** Given an image  $x$  and a question  $q$ , the LVLMM first generates an initial draft answer  $\hat{y}$  ①. The question is decomposed by  $f_{\text{LLM}}$  into a query set  $Q = \{q_i\}$ , and each query is grounded by  $f_g$  to produce spatial relevance maps, aggregated into  $r(x | q)$  ②. Gumbel- $k$  sampling masks Top- $k$  and Bottom- $k$  regions for perturbation, and a semantic encoder  $g(\cdot)$  measures similarity shifts between  $\hat{y}$  and perturbed predictions  $\hat{y}_r$  to compute the utilization score  $U_q^{\text{base}}$  ③. Expert models (e.g., CLIP, SAM, OCR) render structured visual evidence onto the image, producing refined outputs with updated utilization  $U_q^{(j)}$ . The expert with the largest gain  $U_q^{(j)} - U_q^{\text{base}}$  is selected for refinement ④.

ful evidence. Therefore, we introduce a *query-conditioned relevance map*  $r(x | q)$ , which localizes and highlights the areas of the  $x$  most informative for answering  $q$ .

**Extracting query terms.** We first transform the free-form question  $q$  into a set of explicit visual queries  $Q = \{q_1, q_2, \dots, q_m\}$ . Directly using  $q$  is often suboptimal, as natural questions may include abstract or relational expressions (e.g., “Is the person hungry?”) that are difficult for vision encoders to interpret. To bridge this gap, a large language model  $f_{\text{LLM}}$  reformulates  $q$  into visually grounded queries  $Q = f_{\text{LLM}}(q)$ , where each  $q_i$  denotes a concrete, visually identifiable concept such as an object or attribute. For example, for the question “What is the man wearing on his feet?”, the  $f_{\text{LLM}}$  generates queries like “shoes”, “feet”, and “clothing”. These queries serve as explicit textual anchors that guide vision–language alignment, as illustrated in Fig. 2.

**Computing query-conditioned regions.** Given the extracted queries  $Q$ , each  $q_i$  is used to localize its corresponding visual evidence within  $x$ . We employ a CLIP-based visual grounding model  $f_g$ , where each text query  $q_i$  guides the decoder to produce a spatial relevance map  $R(x | q_i) = f_g(x, q_i) \in [0, 1]^{H \times W}$ , representing pixel-wise relevance to  $q_i$ . This retrieval-like grounding process aligns textual concepts with their spatial counterparts.

$$r(x | q) = \frac{1}{m} \sum_{q_i \in Q} R(x | q_i), \quad Q = f_{\text{LLM}}(q). \quad (1)$$

Averaging across all query terms ensures consistent activation of semantically related regions while suppressing spurious responses as defined in Eq. (1). The resulting map

$r(x | q)$  provides the spatial foundation for evaluating how effectively the LVLMM leverages relevant evidence in the subsequent stages of DnR.

## 3.2. Question-Conditioned Utilization

After obtaining  $r(x | q)$ , the next objective is to assess how effectively the model utilizes the critical visual regions indicated by  $r(x | q)$ . The *question-conditioned utilization*  $U_q(x)$  quantifies how the model’s prediction varies when these regions, either highly relevant (Top- $k$ ) or irrelevant (Bottom- $k$ ), are perturbed. A higher  $U_q(x)$  indicates that the model responds sensitively to question-critical evidence while remaining stable against distractive regions, indicating more faithful and evidence-grounded reasoning.

### 3.2.1. Relevance-based probabilistic masking

To quantitatively evaluate how much the model’s prediction depends on the regions identified as relevant, we construct a probability distribution over the spatial regions of the image proportional to their relevance scores. Let  $\mathcal{U}(x)$  denote the set of candidate regions within  $x$ . The probability of each region  $u \in \mathcal{U}(x)$  is computed by normalizing its relevance value obtained from the relevance map.

$$P(u | x, q) = \frac{r(u | x, q)}{\sum_{u' \in \mathcal{U}(x)} r(u' | x, q)}. \quad (2)$$

This normalized distribution, Eq. (2), provides a probabilistic weighting over image regions according to their question-conditioned importance. Based on this distribution, two complementary masking strategies are employed, as illustrated in Fig. 3. **Top- $k$  masking** occludes highly relevant regions (high  $P(u | x, q)$ ) to measure degradation in

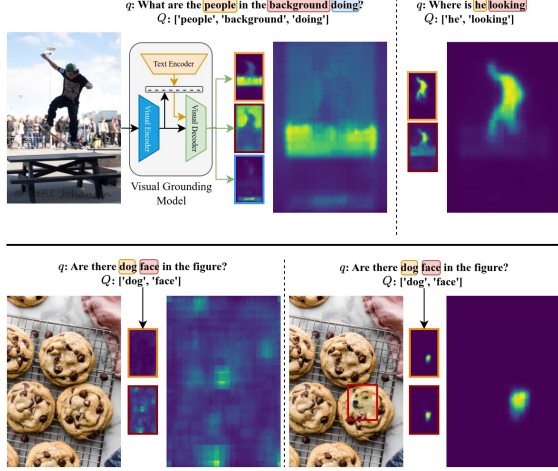


Figure 2. **Illustration of the query-conditioned relevance map.** For the same image (top row), different questions lead to distinct relevance regions aligned with the extracted query terms. Conversely, for the same question (bottom row), the relevance map varies with the image content, localizing evidence that matches the queried concept.

model prediction, while **Bottom- $k$  masking** occludes less relevant regions (low  $P(u | x, q)$ ) to assess prediction stability. Sampling from these two strategies yields a set of  $M$  stochastic binary masks, denoted as  $\mathcal{M}_q = \tau_1, \dots, \tau_M$ , each covering approximately a  $\rho$  fraction of the image area. To ensure stochastic region selection, we adopt Gumbel- $k$  sampling [23, 60], which enables diverse Top- $k$ /Bottom- $k$  mask generation from the relevance distribution. **Adaptive combination of Top- $k$  and Bottom- $k$  masking.** While Top- $k$  and Bottom- $k$  perturbations reveal complementary aspects of model behavior, their relative influence should depend on the sharpness and distribution of the relevance map  $r(x | q)$ . We therefore propose an *adaptive weighting factor*  $\alpha \in [0, 1]$  per sample by analyzing the information structure of  $r(x | q)$ . Specifically,  $\alpha$  is obtained through a function  $\alpha = \Phi_{\text{adapt}}(r(x | q))$  that integrates two cues (**entropy** and **contrast**) reflecting the focus and separability of the relevance distribution. We normalize  $r(x | q)$  into  $p_i = r_i / \sum_j r_j$  and compute the normalized entropy  $\mathcal{H}_{\text{norm}} = -\sum_i p_i \log p_i / \log N$ , and define contrast as  $C = (\max_i r_i - \min_i r_i) / \max_i r_i$ . The coefficients  $\beta_{\text{ent}}$  and  $\beta_{\text{ctr}}$  control their relative influence.

Eq. (3) defines this adaptive function as a weighted combination of the normalized entropy  $\mathcal{H}_{\text{norm}}$  and contrast  $C$ :

$$\alpha = \frac{\beta_{\text{ent}} \mathcal{H}_{\text{norm}} + \beta_{\text{ctr}} C}{\beta_{\text{ent}} + \beta_{\text{ctr}}}. \quad (3)$$

A higher  $\alpha$  corresponds to a sharper and more distinct relevance map, assigning greater weight to Top- $k$  masking, whereas a lower  $\alpha$  emphasizes Bottom- $k$  masking to assess stability under diffuse relevance. This adaptive mech-

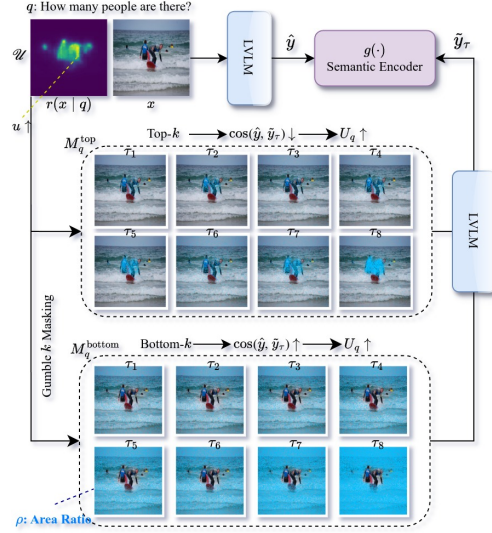


Figure 3. **Question-conditioned utilization computation.** Given a question  $q$  and image  $x$ , the relevance map  $r(x | q)$  guides *Gumbel Top- $k$ /Bottom- $k$*  masking over a ratio  $\rho$  of the image. Masked inputs  $\tau(x)$  are fed into the LVLM to obtain perturbed predictions  $\tilde{y}_\tau$ , compared with the original  $\hat{y}$  via a semantic encoder  $g(\cdot)$ , and aggregated with the adaptive factor  $\alpha$  to compute the final utilization score  $U_q(x)$ .

anism allows  $U_q(x)$  to dynamically decide which perturbation to emphasize based on the input, focusing on decisive evidence when the relevance map is confident and prioritizing stability when it is uncertain.

### 3.2.2. Semantic deviation measurement.

For each mask  $\tau \in \mathcal{M}_q$ , a masked image  $\tau(x)$  is generated by removing the selected regions. Let  $\hat{y} = f_{\text{VLM}}(x, q)$  denote the original prediction and  $\tilde{y}_\tau = f_{\text{VLM}}(\tau(x), q)$  the masked prediction. Both are projected into an embedding space by a semantic encoder  $g(\cdot)$ , yielding  $z_{\hat{y}} = g(\hat{y})$  and  $z_{\tilde{y}_\tau} = g(\tilde{y}_\tau)$ .

The semantic deviation for each mask is defined differently depending on its type (top- $k$  or bottom- $k$ ).

$$d_\tau(\hat{y}, \tilde{y}_\tau) = \begin{cases} 1 - \cos(z_{\hat{y}}, z_{\tilde{y}_\tau}), & \text{if } \tau \in \mathcal{M}_q^{\text{top}}, \\ \cos(z_{\hat{y}}, z_{\tilde{y}_\tau}), & \text{if } \tau \in \mathcal{M}_q^{\text{bottom}}. \end{cases} \quad (4)$$

As shown in Eq. (4), Top- $k$  masking measures the drop in semantic similarity after masking highly relevant regions, whereas Bottom- $k$  masking quantifies the stability of predictions when irrelevant regions are masked.

### 3.2.3. Utilization score.

Building upon the semantic deviations defined in Eq. (4), the final utilization score  $U_q(x)$  aggregates the semantic deviations from both Top- $k$  and Bottom- $k$  perturbations through the adaptive factor  $\alpha$ . It quantifies how the model's



Figure 4. Comparison of rendering strategies across different experts. Each column corresponds to an experts, and each row represents a rendering style.

prediction changes when query-relevant or irrelevant regions are masked, providing a balanced measure of evidence dependence and robustness.

$$U_q(x) = \alpha \cdot \mathbb{E}_{\tau \in \mathcal{M}_q^{\text{top}}} [d_\tau(\hat{y}, \tilde{y}_\tau)] + (1 - \alpha) \cdot \mathbb{E}_{\tau \in \mathcal{M}_q^{\text{bottom}}} [d_\tau(\hat{y}, \tilde{y}_\tau)]. \quad (5)$$

In Eq. (5), the expectation operator  $\mathbb{E}_{\tau \in \mathcal{M}_q^{\text{top}}}[\cdot]$  denotes the mean semantic deviation over all masks in each subset. Owing to the complementary definitions in Eq. (4), a higher  $U_q(x)$  consistently indicates a more faithful and visually grounded reasoning process. Specifically, the Top- $k$  term captures how **strongly the model relies** on question-critical regions, where larger deviations imply stronger evidence dependence, whereas the Bottom- $k$  term measures how **stable the prediction remains** when irrelevant regions are perturbed. Their contributions are adaptively balanced by the factor  $\alpha$ , yielding a unified measure of both dependence and robustness.

### 3.3. Expert Selection and Incorporation

*DnR* leverages  $U_q$  as a quantitative measure to guide the selection of visual experts. Given multiple candidates  $\{h_1, h_2, \dots, h_K\}$ , each expert generates structured outputs such as bounding boxes, masks, or depth maps, which the LVLM cannot directly process. To make these outputs compatible, they are rendered onto the original image to produce  $\hat{x}^{(j)} = R(x, h_j)$ , forming rendered visual evidence. This rendering converts the expert’s structural predictions into visual cues (e.g., gray, blur, or highlight regions) that the model can process through its existing vision encoder without any architectural modification (Fig. 4).

$$\tilde{y}^{(j)} = f_{\text{VLM}}(\hat{x}^{(j)}, q) = f_{\text{VLM}}(R(x, h_j), q) \quad (6)$$

Each  $\tilde{y}^{(j)}$  represents the model’s refined response obtained by re-querying  $f_{\text{VLM}}$  with the rendered input  $\hat{x}^{(j)} = R(x, h_j)$ , where the rendering  $R(\cdot)$  visually encodes the expert’s structural output onto the original image to serve as explicit visual evidence.

Subsequently, the framework computes the query-conditioned utilization  $U_q^{(j)}$  for each refined response  $\tilde{y}^{(j)}$  and compares it to the baseline utilization  $U_q^{\text{base}}$  obtained from the initial draft  $\hat{y}$  (Fig. 1). For each expert-rendered input  $\hat{x}^{(j)}$  and its prediction  $\tilde{y}^{(j)}$ , the utilization score  $U_q^{(j)}$  is re-computed by applying the original masks  $\tau \in \mathcal{M}_q$  to  $\hat{x}^{(j)}$  and using  $\tilde{y}^{(j)}$  as the new baseline for Eq. (4). The expert that yields the largest improvement is then selected:

$$j^* = \arg \max_j (U_q^{(j)} - U_q^{\text{base}})_+, \quad (7)$$

If no expert increases  $U_q(x)$  beyond  $U_q^{\text{base}}$ , the system concludes that the draft already captures sufficient visual grounding and skips further refinement, where  $(\cdot)_+$  ensures that only positive improvements over the baseline are considered.

### 3.4. Learned Expert Selection

Running all  $K$  experts for Eq. (7) is computationally expensive. To make *DnR* practical, a lightweight **selector** network  $S_\theta(j | s)$  is trained to predict the optimal expert  $j^*$  as defined in Eq. (7) directly from the state  $s$ .

The state  $s$  represents the query-conditioned context observed before expert invocation:

$$s = (x, Q, \hat{y}, r(x | q)), \quad (8)$$

where  $x$  is the image,  $Q$  the query set,  $\hat{y}$  the draft prediction, and  $r(x | q)$  the relevance map. The selector  $S_\theta$  is trained with the loss

$$\mathcal{L} = -\mathbb{E}_{s \sim \mathcal{D}} [\log S_\theta(j^* | s)]. \quad (9)$$

This setup enables scalable expert coordination as the number of experts increases, where direct evaluation becomes linearly expensive since each candidate requires separate rendering and utilization computation. Leveraging the utilization difference  $(U_q^{(j)} - U_q^{\text{base}})$  defined in *DnR*, the selector learns an approximate rule that converts the refinement process from exhaustive search to direct decision.

## 4. Experiments

We evaluate *DnR* to validate the proposed  $U_q$  and its role in multimodal reasoning. The framework performs a single-step selection by choosing the candidate that maximizes  $U_q$ , and experiments examine how this choice influences accuracy, reduces hallucination, and promotes evidence-grounded reasoning.

VLM Backbone	VQA					Image Captioning				Knowledge VQA				Comprehensive Benchmarks		
	VQAv2 [12]	GQA [18]	VizWiz [14]	TextVQA [57]	OCR-VQA [44]	COCO [37]	NoCaps [2]	Flickr [48]	VCR [76]	VSR [67]	OK-VQA [43]	A-OKVQA [54]	ScienceQA [41]	MME [11]	MMBench [39]	SEED-Bench [25]
IDEFICS [24] (Draft / DnR)	37.8/47.85	24.1/25.5	43.33/36.67	30.15/30.32	47.24/47.74	135.66/137.7	106.53/108.5	74.2/80.4	15.58/21.11	52.76/54.27	43.5/44.4	69/69.5	43.95/44.02	1392.11/1431.58	50.01/50.26	32.16/32.75
Revision Rate	29.8	1.5	19.8	2	3.2	-	-	-	18.3	2.5	2.9	0.2	0.2	0.6	0.4	0.3
Correction / Degradation	46.2/14.3	51.3/1.1	6.3/24.1	2.3/0.3	12.5/2.1	-	-	-	56.3/7.1	52.9/32.3	33.3/1.4	8.6/1.9	94.3/0.4	80.3/16.6	43.8/12.2	33.3/19.6
Pearson/Spearman	0.143/0.064	0.449/0.364	0.248/0.259	0.026/0.073	0.129/0.222	-	-	-	0.38/0.421	0.158/0.173	0.351/0.210	0.1/0.166	0.351/0.277	0.12/0.148	0.12/0.06	0.354/0.375
InstructBLIP [8]	76.4/77.75	38.24/39.77	37.83/38.67	52.43/54.1	80.4/81.91	114.31/118.7	107.8/109.1	72.7/79.8	13.07/12.56	61.31/61.81	49.2/50.2	79.11/81.09	50.5/52.5	1294.74/1295.31	51.84/52.89	51.46/53.8
Revision Rate	7.2	13	24.2	2.1	1.2	26.2	31.1	36	4.1	1.2	4.5	1.8	1.2	0.1	1.6	3.2
Correction / Degradation	9.1/5.4	35.3/1.2	16.1/2	52.3/25.25	24.5/3.4	-	-	-	13.2/21.3	59.9/33.1	39.5/18.4	42.9/28.6	77.6/4.3	92.1/4.3	55.3/4.4	36.4/9.1
Pearson/Spearman	0.243/-0.024	0.290/0.286	0.290/0.286	0.066/0.076	0.426/0.415	-	-	-	0.2/0.204	0.444/0.361	0.168/0.09	0.136/0.104	0.608/0.421	0.152/0.128	0.01/0.02	0.062/0.145
MiniGPTv2 [77]	32.6/34.1	25.3/27.6	59.17/60.67	36.67/36.68	56.78/58.79	-	-	-	13.07/15.58	40.2/43.72	18.1/19.8	38.58/41.01	28.77/29.41	878.95/910.53	37.63/39.21	29.82/31.58
Revision Rate	17.5	3.1	3.8	15.5	1.8	-	-	-	7.2	0.8	4.5	6.8	0.2	0.2	0.4	7.6
Correction / Degradation	12.9/2.9	8.3/0.5	41.2/6.7	1.6/0.1	11.1/0.1	-	-	-	40.9/3.4	65.1/20.8	11.1/0.1	14.8/7.2	94.3/0.5	50.32/12.1	66.7/1.2	30.8/15.7
Pearson/Spearman	0.026/0.051	0.194/0.378	0.394/0.234	0.092/0.129	0.065/0.124	-	-	-	0.292/0.224	0.155/0.165	0.206/0.168	0.108/0.122	0.719/0.807	0.347/0.355	0.422/0.360	0.155/0.131
LLaVA 1.6 [38]	80.9/82.81	61.5/64.2	76.83/76.99	64.49/64.59	73.37/74.87	126.54/138.6	70.18/75.5	77.9/79.6	18.59/18.69	64.82/65.83	55.1/56.1	73.5/73.5	72.9/73.8	1694.74/1721.05	76.32/77.89	66.08/66.67
Revision Rate	5.1	4.5	1.8	3.9	1.8	22.8	37.1	23	0.5	0.2	2	0.8	1.4	0.4	0.5	0.3
Correction / Degradation	66.7/9.2	33.3/3.1	14.3/2.1	25/0.7	66.7/16.7	-	-	-	25.3/11.4	95.3/0.3	37.5/24.9	33.3/7.8	49.8/0.4	98.3/0.6	23.3/7.3	95.5/0.4
Pearson/Spearman	0.509/0.639	0.288/0.328	0.156/0.154	0.155/0.231	0.626/0.849	-	-	-	0.772/0.778	0.272/0.272	0.297/0.161	0.110/0.128	0.454/0.459	0.270/0.307	0.109/0.116	0.112/0.001
PaliGemma [9]	73.2/75.2	58.3/59.9	79.17/80.17	65.83/66.5	68.34/70.35	137.99/144.7	83.81/85.99	92.9/101.9	13.57/15.53	66.83/68.84	57.2/58.5	85.2/85.5	89.1/89.7	1434.21/1444.74	70.09/71.05	57.89/59.65
Revision Rate	9.2	1.5	2.8	1.2	3.5	22.6	38.2	53.2	2.2	0.9	3.2	1	0.5	0.2	0.5	4.1
Correction / Degradation	56.4/27.3	16.7/0.8	82.1/2.3	42.3/21.3	28.6/0.5	-	-	-	43.7/5.4	75.7/11.1	60.15/23.1	12.3/0.3	92.1/1.3	51.1/12.5	23.5/2.1	64.3/31.7
Pearson/Spearman	0.292/0.291	0.276/0.399	0.307/0.251	0.136/0.03	0.321/0.329	-	-	-	0.253/0.265	0.341/0.275	0.165/0.128	0.441/0.454	0.803/0.853	0.08/0.07	0.09/0.1	0.129/0.191
CogVLM [17]	82.05/82.85	56.13/57.74	48.5/50.33	68.51/69.68	82.91/82.91	86.57/94.8	80.3/81.7	67.7/70.5	12.56/13.29	62.81/65.33	58.6/58.6	84.5/84.5	61.5/62.22	1384.21/1423.68	76.84/77.89	58.48/59.06
Revision Rate	5.1	2.5	4.8	2.30	0.20	13.4	11.1	30.6	2.5	0.5	0.8	0.7	1	0.6	0.3	0.3
Correction / Degradation	66.7/2.3	32.2/1.3	5.7/0.4	42.9/28.6	1.3/0.1	-	-	-	56.3/19.4	99.7/0.1	33.1/4.6	3.4/2.1	75.43/21.33	66.6/31.1	70.1/2.9	98.1/0.3
Pearson/Spearman	0.684/0.735	0.223/0.357	0.220/0.136	0.318/0.111	0.263/0.28	-	-	-	0.110/0.117	0.458/0.456	0.500/0.623	0.163/0.14	0.479/0.412	0.420/0.470	0.01/0.03	0.153/0.227
Qwen2.5-VL [49]	83.95/85.45	57.03/58.31	73.01/73.83	83.92/84.25	72.36/74.86	61.96/64.7	71.8/74.2	57.3/55.9	30.65/31.66	79.9/80.4	58.2/58.3	74.4/76.4	86.5/87.1	2268.42/2276.32	86.05/86.84	80.12/81.29
Revision Rate	10.6	8.8	1.2	0.20	2.5	27.6	37.5	42.1	5.5	0.2	3.8	1.5	0.4	0.3	1.2	0.5
Correction / Degradation	40/4.3	16.7/8.3	41.3/1.0	33.3/5.7	89.3/3.4	-	-	-	45.5/18.2	93.9/0.3	13.3/6.7	33.2/16.8	82.4/5.3	65.6/32.1	55.5/42.3	29.4/17.6
Pearson/Spearman	0.389/0.509	0.136/0.222	0.19/0.224	0.230/0.112	0.499/0.425	-	-	-	0.637/0.631	0.25/0.25	0.02/0.023	0.43/0.467	0.644/0.703	0.5/0.5	0.188/0.083	0.170/0.07

Table 1. Comprehensive evaluation of diverse LVLMs on 16 multimodal benchmarks. Each cell reports *Draft* / *DnR* performance. Additional rows show revision rates (percentage of refined), correction vs. degradation (*False*  $\rightarrow$  *True* / *True*  $\rightarrow$  *False*) transitions, and correlation coefficients (Pearson / Spearman) measuring alignment between confidence or utilization scores and accuracy changes. Cells in green denote improvements after refinement, and those in blue indicate strong correlations ( $r > 0.2$ ).

We use LLaMA-3-70B [58] as the language backbone  $f_{\text{LLM}}$  for query decomposition, CLIP-L/14 [50] and SentenceTransformer all-MiniLM-L6-v2 [52] as the semantic encoder  $g(\cdot)$ , and CLIPSeg [42] as the visual grounding model  $f_g$ . Four visual experts (GroundingDINO [40], SAM [22], DepthAnything [69], and mDETR [19]) are chosen from distinct backbone families to provide complementary cues. For the  $U_q$  computation, we sample  $M = 16$  stochastic masks per image using Gumbel- $k$  sampling, with the masking ratio  $\rho$  set to 0.25 for Top- $k$ , 0.75 for Bottom- $k$ , and  $\beta_{\text{ent}} = \beta_{\text{ctr}} = 0.5$  for adaptive hybrid masking.

## 4.1. Comprehensive Evaluation

We evaluate *DnR* across five categories: VQA, image captioning, visual reasoning, knowledge VQA, and comprehensive benchmarks, each examining a distinct aspect of multimodal reasoning with subsets. All experiments use the *GRAY* rendering strategy and *Hybrid* masking mode for consistency. The utilization metric  $U_q$  acts as a unified indicator connecting visual evidence to model behavior, and its validity is tested through Pearson and Spearman correlations between utilization scores and accuracy gains.

### 4.1.1. VQA

*DnR* improves visual reasoning across diverse LVLMs, but its effect varies with each model’s baseline grounding. Models with weaker visual reliance, such as IDEFICS [24] and MiniGPT-v2 [77], show higher revision rates ( $\approx 8.3$ – $11.3\%$ ) and larger gains ( $+1.4$ – $3.0\%$ ), indicating that *DnR* supplies missing visual feedback. Stronger systems like LLaVA 1.6 [38], CogVLM [17], and Qwen2.5-VL [49] revise less often ( $\approx 3.0$ – $4.6\%$ ) yet still improve ( $+1.1$ – $1.5\%$ ), suggesting a stabilizing effect that refines already coherent reasoning. This contrast shows that *DnR*

enhances consistency while moderating excessive or insufficient visual dependency during inference.

*DnR* scales with each task’s visual demand. Perception-focused benchmarks such as VQAv2 [12] and VizWiz [14] show larger gains (1.8–2.7%), while text-centric tasks like OCR-VQA [44] and TextVQA [57] show smaller but consistent improvements (1.4–1.6%). An exception appears in IDEFICS–VizWiz, the drop occurs because *DnR* pushes the model to answer cases it previously avoided by saying “unanswerable,” and its limited capability on these queries naturally leads to more mistakes. A high correction-to-degradation ratio (31.8–6.5%) and positive Pearson and Spearman correlations (0.275/0.273) indicate that *DnR* systematically shifts predictions toward more evidence-grounded reasoning rather than random variation.

### 4.1.2. Image Captioning

In image captioning, where no explicit  $q$  is provided, *DnR* constructs object-centric  $Q$  to explore the scene and guide visual reasoning. The strength of improvement depends on each model’s initial captioning bias. Models producing linguistically generic descriptions, such as IDEFICS and InstructBLIP, show moderate CIDEr gains ( $+2.0$ – $4.4$ ) at low medium revision rates (17–36%), suggesting that *DnR* primarily supplements missing scene-specific details. In contrast, PaliGemma and LLaVA 1.6 yield larger gains ( $+6.6$ – $12.1$ ), especially on COCO [37] and Flickr [48], where object-level grounding is emphasized. Overall, *DnR* promotes more discriminative and visually grounded phrasing even when baseline captions are already coherent.

COCO and Flickr show high CIDEr gains supported by high revision frequencies (22–53%), indicating that *DnR* shifts captions away from generic phrasing toward grounded, object-centric descriptions. NoCaps [2], with

weaker visual constraints, yields smaller but steady gains (+1.7–5.3). A minor drop appears for Qwen2.5-VL on Flickr (57.3→55.9) because its global pooling, amplified by *DnR*'s focus on high-confidence cues, suppresses fine-grained details. Overall, the positive correlation between revision rate and CIDEr gain shows that *DnR* steers captions toward visual fidelity rather than altering them arbitrarily.

#### 4.1.3. Visual Reasoning

Across reasoning benchmarks, *DnR* strengthens both semantic and spatial reasoning by aligning model decisions with visual evidence. On VCR [76], visually weaker models such as IDEFICS and MiniGPT-v2 revise more often (7–18%) and gain larger improvements (+2.0–5.5), showing that *DnR* supplies missing visual cues. More visually coherent models (LLaVA 1.6, CogVLM, Qwen2.5-VL) revise infrequently (<6%) yet still improve (+0.1–1.0), with high Pearson correlations (0.63–0.77) indicating stable alignment between utilization shifts and accuracy. InstructBLIP shows a slight drop (-0.5%) when grounding fails to influence its reasoning stage. Overall, *DnR* corrects visually under-grounded models and refines visually strong ones without destabilizing performance.

Meanwhile, in VSR [67], where spatial relationships dominate, revisions occur rarely (<1%) but consistently yield measurable gains (+0.5–2.5%), showing that *DnR* corrects spatial misinterpretations while maintaining overall stability. CogVLM and LLaVA 1.6 achieve near-perfect correction ratios (99.7% and 95.3%) and moderate correlations ( $\approx 0.27$ –0.46), demonstrating that *DnR* fine-tunes spatial grounding with minimal perturbation. Overall, *DnR* acts as a dual-function mechanism reinforcing under-grounded reasoning and regularizing over-grounded logic, thereby improving both semantic and spatial fidelity across models.

#### 4.1.4. Knowledge-based VQA

Across knowledge-based benchmarks, *DnR* reinforces perceptual grounding but cannot substitute for missing conceptual knowledge. This limitation is evident from the low revision rates ( $\leq 6.8\%$ ) and small accuracy gains (+0.5–2.0%), indicating that additional visual cues cannot recover facts the model does not know. Weak grounded models, such as IDEFICS and MiniGPT-v2, achieve higher correction ratios (11–33%) despite minimal revisions, showing that *DnR* helps retrieve existing knowledge previously inaccessible due to poor visual grounding.

In contrast, stronger models like PaliGemma, LLaVA 1.6, and Qwen2.5-VL maintain steady improvements (+0.6–1.3%) with high utilization–accuracy correlation (Pearson/Spearman  $\approx 0.7$ –0.85 in ScienceQA [41]), confirming that their factual reasoning is already visually aligned. A-OKVQA [54], however, shows near-zero correlation ( $\leq 0.2$ ) and minimal accuracy change, demonstrating that when conceptual knowledge is missing, *DnR* offers no

benefit. In summary, the low revision activity, stable gains, and strong correlations together confirm that *DnR* improves access to existing knowledge through perceptual reinforcement, not knowledge augmentation.

#### 4.1.5. Comprehensive Benchmark

Comprehensive benchmarks evaluate overall multimodal consistency rather than task-specific accuracy. Because these datasets contain stable and unambiguous inputs, *DnR* yields few revisions ( $\leq 2\%$ ) and small but consistent gains, indicating a shift from error correction to stability refinement. Moderate utilization–accuracy correlations (0.3–0.5) confirm that such adjustments are selective rather than random. Across models, IDEFICS and MiniGPT-v2 gain mainly in perception-oriented metrics (+31–39 in MME [11]) as *DnR* compensates weak grounding, whereas stronger systems such as LLaVA 1.6, CogVLM, and Qwen2.5-VL show minimal revision but sustained coherence ( $>95\%$  correction). Overall, once multimodal grounding stabilizes, *DnR* works primarily to preserve consistency rather than alter predictions, defining its operating range as stability refinement.

## 4.2. Hallucination

We evaluate hallucination under the *highlight* rendering configuration, which emphasizes salient regions while preserving context. Four benchmarks are used: HaloQuest [63] and MMHal-Bench [75] for hallucination-oriented VQA, VizWiz for real-world *unanswerable* cases, and COCO Caption for free-form captioning. These benchmarks are chosen because they provide detailed hallucination annotations rather than binary labels. Each response was categorized as **Hallucination (H)**, **Misperception (M)**, **Grounded (G)**, or **Correct (C)** by ChatGPT 4o-mini [46].

*DnR* consistently reduced **hallucinations** across all benchmarks, with absolute decreases of 1–9 pp and percentage drops of 8–29%. InstructBLIP (6.78 pp, 26.7%) and Qwen2.5-VL (5.96 pp, 35.0%) showed the strongest declines, while LLaVA 1.6 (29.1%) achieved the largest proportional reduction. Average **Misperception** decreased by 1–3 pp. InstructBLIP (-3.97 pp, -13%) and LLaVA 1.6 (-1.53 pp, -10%) showed the strongest declines, while in CogVLM and PaliGemma, some samples shifted from hallucination to misperception.

Average **grounding** increased by 0.6–5.9 pp across datasets. CogVLM (+3.94 pp, +33.6%) and Qwen2.5-VL (+6.23 pp, +22.0%) achieved the largest gains. LLaVA 1.6 and Qwen2.5-VL showed the clearest shifts from hallucination to grounded responses on MMHal-Bench, while PaliGemma exhibited a similar transition on HaloQuest. Average **correctness** increased by 0.5–2 pp across benchmarks. PaliGemma (+3.46 pp, +7.4%) and MiniGPT-v2 (+0.78 pp, +4.56%) showed the most notable gains. These

Model	HaloQuest [63]			MMHal-Bench [75]				VizWiz [14]				COCO Caption [37]		
	↓H	M	C↑	↓H	M	G	C↑	↓H	M	G	C↑	↓H	M	G↑
IDEFICS [24]	43.34/ <b>40.87</b>	22.05/23.76	34.62/ <b>35.37</b>	13.54/ <b>11.50</b>	44.79/44.75	12.50/11.46	29.17/ <b>32.29</b>	12.06/17.09	32.16/30.65	19.10/22.11	36.68/30.15	28.14/ <b>22.11</b>	32.66/37.69	39.20/ <b>40.20</b>
InstructBLIP [8]	33.73/ <b>33.02</b>	20.48/21.54	45.79/ <b>45.44</b>	35.42/ <b>28.12</b>	38.54/45.83	9.38/8.33	16.67/ <b>17.71</b>	23.62/ <b>18.59</b>	23.62/26.13	20.60/22.11	32.16/ <b>33.17</b>	37.69/ <b>23.62</b>	29.15/34.17	33.16/ <b>42.22</b>
MiniGPT-v2 [77]	20.22/ <b>19.42</b>	29.75/29.52	50.03/ <b>51.06</b>	35.42/ <b>27.08</b>	48.96/54.17	6.25/8.33	9.38/ <b>10.42</b>	11.56/ <b>10.55</b>	26.13/24.85	10.05/12.06	52.26/ <b>52.54</b>	-	-	-
LLaVA 1.6 [38]	26.33/ <b>25.65</b>	13.96/14.57	59.71/ <b>59.79</b>	8.33/ <b>4.17</b>	26.04/26.04	21.88/23.96	43.75/ <b>45.83</b>	1.01/ <b>0.53</b>	13.07/17.06	20.10/16.58	65.83/65.83	24.62/ <b>20.60</b>	29.15/30.65	46.24/ <b>48.75</b>
PaliGemma [9]	20.67/ <b>16.41</b>	14.96/16.63	64.37/ <b>66.96</b>	12.50/ <b>8.33</b>	31.25/33.33	14.58/9.38	41.67/ <b>48.96</b>	2.01/ <b>1.11</b>	12.56/10.45	14.57/17.09	70.85/ <b>71.36</b>	30.65/ <b>28.14</b>	26.63/26.13	42.72/ <b>45.73</b>
CogVLM [17]	19.24/ <b>17.82</b>	15.60/16.72	65.16/ <b>65.46</b>	6.25/ <b>4.17</b>	52.08/48.96	9.38/17.71	32.29/29.17	20.60/ <b>16.08</b>	15.08/17.59	22.61/24.12	41.71/ <b>42.21</b>	41.21/ <b>36.18</b>	21.61/24.62	37.19/ <b>39.19</b>
Qwen2.5-VL [49]	3.48/ <b>3.07</b>	11.90/12.30	84.63/84.63	4.17/ <b>3.12</b>	20.83/14.58	31.25/35.42	43.75/ <b>46.88</b>	0.49/ <b>0.21</b>	10.55/12.86	25.14/22.10	63.82/ <b>64.83</b>	48.24/ <b>26.13</b>	24.62/29.15	27.14/ <b>44.72</b>

Table 2. **Hallucination-benchmark comparison.** Each cell reports Draft/*DnR* results across four hallucination-oriented datasets. Green cells indicate reduced hallucination (H↓), while blue cells highlight improved grounding or correctness (G or C↑). Values denote the percentage proportion of each category, where Hallucination+Misperception+Grounded+Correct=100(%).

Dataset	Metric	Exhaustive	Policy-Driven	$\Delta$ Perf.	$\Delta$ Cost (%)
VQAv2 [12]	Acc. (%)	75.2	74.9	-0.3	-72.64
COCO [37]	CIDEr	144.7	143.5	-1.2	-52.72
VSR [67]	Acc. (%)	68.8	68.7	-0.07	-69.97
ScienceQA [41]	Acc. (%)	89.7	89.5	-0.2	-78.16
MME [11]	Score	1444.7	1424.4	-20.35	-78.48
<b>Average</b>	-	364.6	360.2	-4.42	-70.39

Table 3. Performance and cost comparison between *Exhaustive* and *Policy-Driven* selection for PaliGemma. Both  $\Delta$  values denote relative differences, with cost reduction reported in percentage.

results indicate modest yet consistent improvements in response accuracy across all models.

Across all models, hallucinated responses were largely redirected toward **Misperception** and **Grounded** categories, indicating a shift from unfounded to visually supported reasoning. A minor exception occurred on VizWiz, where IDEFICS previously overused “unanswerable” responses; after applying *DnR*, its behavior became less conservative, leading to fewer such cases and more contextually grounded answers, albeit with slightly lower raw scores. LLaVA 1.6 and Qwen2.5-VL showed the strongest transitions toward grounded reasoning, while CogVLM and PaliGemma demonstrated moderate yet consistent shifts.

### 4.3. Policy-Driven Expert Selection

We train an expert selector  $S_\theta$  to replace exhaustive expert evaluation. The selector is a three-layer MLP defined over the state  $s$  in Eq. (8) and optimized with the loss in Eq. (9). Using the query set  $Q$  instead of the raw question  $q$  improves performance through stronger alignment with visual representations. Given  $s$ ,  $S_\phi(s) = \arg \max_j U_q^{(j)}$  predicts which expert or initial draft yields the highest  $U_q$ .

PaliGemma [9] was selected for efficiency analysis as it offered the fastest and most reliable inference across datasets. The performance difference remained minimal across four benchmarks (excluding MME with a distinct scoring scale), averaging around -0.4 pp, while computational cost decreased by approximately 70% on average. COCO exhibited the smallest cost reduction (-52.7%) since

it inherently involves more experts during generation, leaving less room for pruning. Occasional mispredictions occurred where the selector chose the *Refine* stage instead of *Draft*, but these cases typically satisfied  $\hat{y} = \tilde{y}^{(j)}$  and  $|U_q^{\text{base}} - U_q^{(j)}| < \epsilon$ , indicating that such swaps had negligible impact on the final outcome in Table 3.

These results demonstrate the potential of framing expert selection as a learnable task guided by  $U_q$ . Rather than remaining an interpretive metric,  $U_q$  becomes a concrete objective for deciding when and how an expert should intervene. This allows the framework to scale to larger and more diverse expert pools under a unified selection criterion.

## 5. Discussion

In our experiments, performance gains show a clear linear correlation with utilization, confirming that higher  $U_q$  reflects stronger visual grounding. However, both the absolute level of  $U_q$  and its change  $\Delta U_q$  vary across datasets, architectures, and inputs. Some models maintain high but stable  $U_q$  with little variation, indicating consistent but less adaptive grounding, whereas others show larger  $\Delta U_q$  shifts with stronger gains. These patterns reveal distinct reasoning behaviors and attention distributions among LVLMS. Adaptive normalization across domains may further improve the interpretability and stability of  $U_q$  without manual tuning.

## 6. Conclusion

We presented *Draft and Refine (DnR)*, a scalable agent-style framework that coordinates multiple visual experts to quantify and improve how LVLMS use visual evidence. *DnR* employs a relevance map to assess visual reliance and selects expert-guided refinements through a lightweight, modality-agnostic interface rather than heuristic control. Experiments across diverse benchmarks demonstrate consistent accuracy gains, reduced hallucination, and clearer attribution of visual reasoning. Overall, *DnR* provides a principled criterion for evaluating and leveraging visual experts, creating a scalable foundation for systematically integrating and expanding expert-driven multimodal reasoning.

## Acknowledgments

This work was supported in part by the DARPA Young Faculty Award, the National Science Foundation (NSF) under Grants #2127780, #2319198, #2321840, #2312517, and #2235472, the Semiconductor Research Corporation (SRC), the Office of Naval Research through the Young Investigator Program Award, and Grants #N00014-21-1-2225 and N00014-24-1-2547, Army Research Office Grant #W911NF2410360. Additionally, support was provided by the Air Force Office of Scientific Research under Award #FA9550-22-1-0253.

## References

- [1] Aishwarya Agrawal, Rishi Nand, and Dhruv Batra. Don't just assume; look and answer: Overcoming priors for visual question answering. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. 2
- [2] Harsh Agrawal, Karan Desai, Yufei Wang, Jiasen Lu, Vlad I Morariu, Stefan Lee, Mark Yatskar, Devi Parikh, and Dhruv Batra. nocaps: novel object captioning at scale. In *ICCV*, 2019. 6
- [3] Jean-Baptiste Alayrac et al. Flamingo: a visual-language model for generalist vision + language understanding, 2022. Preprint. 2
- [4] Zechen Bai, Pichao Wang, Tianjun Xiao, Tong He, Zongbo Han, Zheng Zhang, and Mike Zheng Shou. Hallucination of multimodal large language models: A survey. *arXiv preprint arXiv:2404.18930*, 2024. 1, 2
- [5] Hila Chefer, Shir Gur, and Lior Wolf. Transformer interpretability beyond attention visualization. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 1
- [6] Hao Chen, Jiayi Liu, and Rui Feng. Groundingqa: Enhancing visual question answering via evidence retrieval and verification. *arXiv preprint arXiv:2402.09412*, 2024. 2
- [7] Zhen Chen, Rui Han, Xin Zhao, et al. Vcoder: Versatile vision encoders for multimodal large language models. *arXiv preprint arXiv:2403.09820*, 2024. 1, 2
- [8] Wenliang Dai, Junnan Li, Dongxu Li, Anthony Meng Huat Tan, Shafiq Joty, and Steven Hoi. Instructblip: Towards general-purpose vision-language models with instruction tuning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023. 2, 6, 8
- [9] Sagar Desai, Carlos Riquelme, Basil Mustafa, and et al. Paligemma: A unified multimodal model for vision-language understanding. *arXiv preprint arXiv:2405.14289*, 2024. 6, 8
- [10] Alessandro Favero, Luca Zancato, Matthew Trager, Sidharth Choudhary, Pramuditha Perera, Alessandro Achille, Ashwin Swaminathan, and Stefano Soatto. Multi-modal hallucination control by visual information grounding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 1, 2
- [11] Chaoyou Fu, Wenhao Zhang, Yichong Zhang, Dongdong Chen, Shuhui Han, Jing Liao, Errui Ding, and Hui Zeng. Mme: A comprehensive evaluation benchmark for multimodal large language models. *arXiv preprint arXiv:2403.05532*, 2024. 6, 7, 8
- [12] Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In *CVPR*, 2017. 6, 8
- [13] Tanmay Gupta, Aniruddha Kamath, Cordelia Schmid, et al. Vipergpt: Visual inference via python execution for reasoning. *arXiv preprint arXiv:2303.08128*, 2023. 1, 2
- [14] Danna Gurari, Qing Li, Abigale Stangl, Anhong Guo, Chi Lin, Kristen Grauman, Jiebo Luo, and Jeffrey P Bigham. Vizviz grand challenge: Answering visual questions from blind people. In *CVPR*, 2018. 6, 8
- [15] Fabiola Gurrola-Kim, Sanja Fidler, and Yanwei Fu. Reliable vqa: Abstain rather than answer. In *European Conference on Computer Vision (ECCV)*, 2022. 2
- [16] Jun Han, Yifan Li, Ziyu Zhao, et al. Grounding by attention: A strong baseline for faithful multimodal reasoning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2023. 2
- [17] Shizhe Hong, Weihao Yu, Yuntao Chen, and et al. Cogvlm: Visual language models for vision-centric tasks. *arXiv preprint arXiv:2403.00504*, 2024. 6, 8
- [18] Drew A Hudson and Christopher D Manning. Gqa: A new dataset for real-world visual reasoning and compositional question answering. In *CVPR*, 2019. 6
- [19] Aishwarya Kamath, Mannat Singh, Gabriel Ilharco, Pratyay Banerjee, Tanmay Kumar, Sayna Ebrahimi, Yong Jae Lee, Marcus Rohrbach, Douwe Kiela, and Aniruddha Kembhavi. Mdetr: Modulated detection for end-to-end multi-modal understanding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. 6
- [20] Kashif Khan and Yanwei Fu. Consistency and uncertainty: Identifying unreliable responses for selective vqa. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 2
- [21] Minseok Kim, Hyojin Kim, and Chang D Yoo. Single-modal entropy-based active learning for visual question answering. In *British Machine Vision Conference (BMVC)*, 2021. 2
- [22] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollár, and Ross Girshick. Segment anything. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023. 1, 6
- [23] Wouter Kool, Herke van Hoof, and Max Welling. Stochastic beams and gumbel-top-k sampling for diverse generation. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019. 4
- [24] Hugo Laurençon, Anton Lozhkov, Lucile Saulnier, and et al. Idefics 2: Image-aware decoder-only transformers. *Hugging Face Technical Report*, 2024. 6, 8
- [25] Bowen Li, Yichong Zhang, Bowen Zhao, Tianyu Lin, Xinyue Wang, Dongdong Chen, and Hui Zeng. Seed-bench: Benchmarking multimodal large language models. *arXiv preprint arXiv:2403.17839*, 2024. 6

- [26] Chen Li, Wei Zhang, et al. Sphinx: The joint mixing of weights, tasks, and visual embeddings for multi-modal large language models. *arXiv preprint arXiv:2405.06560*, 2024. 1, 2
- [27] Dongxu Li, Bo Zhao, et al. Visionllm v2: An end-to-end generalist multimodal large language model. *arXiv preprint arXiv:2406.01479*, 2024. 1, 2
- [28] Haotian Li, Pengchuan Liu, Yue Li, Silvio Savarese, and Steven C. H. Hoi. Visual instruction tuning. *arXiv preprint arXiv:2304.08485*, 2023. 1
- [29] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. BLIP-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *Proceedings of the 40th International Conference on Machine Learning*, pages 19730–19742. PMLR, 2023. 2
- [30] Jiacheng Li, Hangjie Deng, et al. Hyperclip: Adapting vision-language models with hypernetworks. *arXiv preprint arXiv:2404.01379*, 2024. 2
- [31] Jinyu Li, Chen Zhang, Jiaqi Xu, and Ziwei Liu. Mova: Adapting mixture of vision experts to multimodal context. *arXiv preprint arXiv:2402.05000*, 2024. 1
- [32] Qiang Li, Jiaming Zhao, et al. Hyperllava: Dynamic visual and language expert tuning for multimodal large language models. *arXiv preprint arXiv:2404.08580*, 2024. 1, 2
- [33] Zujie Li, Zhecan Wang, Yue Wu, and Wen Gao. Faithfulvqa: A benchmark for faithful visual question answering. In *European Conference on Computer Vision (ECCV)*, 2022. 2
- [34] Zihan Li, Yifan Yuan, Jing Han, et al. Meta-prompting for automating zero-shot visual recognition with llms. *arXiv preprint arXiv:2401.02345*, 2024. 2
- [35] Zongxia Li, Xiyang Wu, Hongyang Du, Fuxiao Liu, Huy Nghiem, and Guangyao Shi. A survey of state of the art large vision language models: Alignment, benchmark, evaluations and challenges. *arXiv preprint*, 2025. 2
- [36] Hao Liang et al. Octotools: An agentic framework with extensible tools for complex reasoning. *arXiv preprint arXiv:2405.08588*, 2024. 2
- [37] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014. 6, 8
- [38] Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. Llava-1.6: Improved multimodal reasoning with visual instruction tuning. *arXiv preprint arXiv:2401.02418*, 2024. 2, 6, 8
- [39] Haotian Liu, Chengzhi Li, Ming Zhang, Wenqi Shao, Bo Li, Ziyu Wang, Qingyang Ye, et al. Mmbench: Is your multi-modal model an all-around player? *arXiv preprint arXiv:2307.06281*, 2024. 6
- [40] Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Chunyuan Li, Peng Gao, Yong-Lu Li, Hongyang Zhang, Tianrui Hui, Qiang Li, Hang Su, Jun Zhu, and Lei Zhang. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 1, 6
- [41] Pan Lu, Swaroop Mishra, Tony Xia, Lin Qiu, Kai-Wei Chang, Song-Chun Zhu, Peter Clark, and Yizhou Liang. Learn to explain: Multimodal reasoning via thought chains for science question answering. In *NeurIPS*, 2022. 6, 7, 8
- [42] Timo Lüddecke and Alexander S Ecker. Clipseg: Image segmentation using text and image prompts. In *CVPR*, 2022. 6
- [43] Kenneth Marino, Mohammad Rastegari, Ali Farhadi, and Roozbeh Mottaghi. Ok-vqa: A visual question answering benchmark requiring external knowledge. In *CVPR*, 2019. 6
- [44] Ankan Kumar Mishra, Karteek Alahari, and CV Jawahar. Ocr-vqa: Visual question answering by reading text in images. In *ICDAR*, 2019. 6
- [45] OpenAI. Gpt-4v(ision): Multimodal capabilities of gpt-4. *OpenAI Technical Report*, 2023. 2
- [46] OpenAI. Chatgpt 4o-mini. <https://chat.openai.com/>, 2025. Used as automatic evaluator for multimodal hallucination analysis. 7
- [47] Zhiliang Peng, Wenhui Wang, Dong Li, Yaru Hao, Shaohan Huang, Shuming Ma, and Furu Wei. Kosmos-2: Grounding multimodal large language models to the world. *arXiv preprint arXiv:2306.14824*, 2023. 2
- [48] Bryan A Plummer, Liwei Wang, Chris M Cervantes, Juan C Caicedo, Julia Hockenmaier, and Svetlana Lazebnik. Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. In *ICCV*, 2015. 6
- [49] Qwen Team. Qwen2.5-vl: A strong multimodal foundation model. Technical report, Alibaba Group, 2025. 6, 8
- [50] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, 2021. 1, 6
- [51] Michael Reid et al. Gemini 1.5: Unlocking multimodal understanding across domains. In *arXiv preprint arXiv:2403.05530*, 2024. 2
- [52] Nils Reimers and Iryna Gurevych. Sentence-bert: Sentence embeddings using siamese bert-networks. In *EMNLP*, 2019. 6
- [53] Timo Schick, Jane Dwivedi-Yu, Roberto Dessì, Roberta Raileanu, Maria Lomeli, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. Toolformer: Language models can teach themselves to use tools. *CoRR*, abs/2302.04761, 2023. 1, 2
- [54] Dustin Schwenk, Christopher Clark, Kenneth Marino, Roozbeh Mottaghi, and Aniruddha Kembhavi. A-okvqa: A benchmark for visual question answering using world knowledge. In *CVPR*, 2022. 6, 7
- [55] Ramprasaath R. Selvaraju, Stefan Lee, Yilin Shen, Hongxia Jin, Shalini Ghosh, Larry Heck, Dhruv Batra, and Devi Parikh. Taking a hint: Leveraging explanations to make vision and language models more grounded. In *International Conference on Computer Vision (ICCV) Workshop on Debugging Machine Learning Models*, 2019. 1
- [56] Yongliang Shen, Kaitao Song, Xu Tan, Dongsheng Li, Weiming Lu, and Yueting Zhuang. Hugginggpt: Solving ai

- tasks with chatgpt and its friends in hugging face. *CoRR*, abs/2303.17580, 2023. 1, 2
- [57] Amanpreet Singh, Vivek Natarajan, Yu Jiang, Xinlei Chen, Marcus Rohrbach, Dhruv Batra, and Devi Parikh. Towards vqa models that can read. In *CVPR*, 2019. 6
- [58] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.12360*, 2024. 6
- [59] Shreyas Tuli, Xiang Liang, Hyun Lee, et al. Vision-amplified semantic entropy for hallucination detection in multimodal large language models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 2
- [60] Nicolas Vieillard, Olivier Pietquin, and Matthieu Geist. Diverse sampling with gumbel-top-k. In *International Conference on Machine Learning (ICML)*, 2020. 4
- [61] Andrés Villa, Juan Carlos León Alcázar, Álvaro Soto, and Bernard Ghanem. Behind the magic, merlim: Multimodal evaluation benchmark for large image-language models. *arXiv preprint arXiv:2312.02219*, 2023. 2
- [62] Yuhang Wang et al. Argus: Vision-centric reasoning with grounded chain-of-thought. *arXiv preprint arXiv:2403.09719*, 2024. 2
- [63] Zhecan Wang, Garrett Bingham, Adams Wei Yu, Quoc V Le, Thang Luong, and Golnaz Ghiasi. Haloquest: A visual hallucination dataset for advancing multimodal reasoning. In *European Conference on Computer Vision*, pages 288–304. Springer, 2024. 7, 8
- [64] Tianhao Wu, Nikhil Nair, and Jiajun Xu. Visual programming: Compositional visual reasoning without training. *arXiv preprint arXiv:2302.14838*, 2023. 1, 2
- [65] Tianhao Wu et al. Visual chatgpt: Talking, drawing and editing with visual foundation models. *arXiv preprint arXiv:2303.04671*, 2023. 2
- [66] Linhui Xiao, Xiaoshan Yang, Xiangyuan Lan, Yaowei Wang, and Changsheng Xu. Towards visual grounding: A survey. *arXiv preprint arXiv:2412.20206*, 2024. 1
- [67] Zhiqiang Xie, Zhiyong Huang, Wenqi Shi, and Yuxin Liu. Visual spatial reasoning with graph-structured representation. In *AAAI*, 2023. 6, 7, 8
- [68] Chen Yang, Yifan Wu, Ming Xu, et al. Retrieval-augmented visual question answering for grounded reasoning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2023. 2
- [69] Xiaokang Yang, Yixiao Ge, and Ying Shan. Depth anything v2: Unleashing the power of large-scale pretraining. *arXiv preprint arXiv:2406.11886*, 2024. 6
- [70] Zhengyuan Yang, Linjie Li, Jianfeng Wang, Kevin Lin, Ehsan Azarnasab, Faisal Ahmed, Zicheng Liu, Ce Liu, Michael Zeng, and Lijuan Wang. Mm-react: Prompting chatgpt for multimodal reasoning and action. *CoRR*, abs/2303.11381, 2023. 1, 2
- [71] Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. ReAct: Synergizing reasoning and acting in language models. In *International Conference on Learning Representations (ICLR)*, 2023. 1, 2
- [72] Jing Ye et al. mplug-owl2: Revolutionizing multi-modal large language model with modality collaboration. *arXiv preprint arXiv:2402.13273*, 2024. 1, 2
- [73] Tao Yin et al. A survey of multimodal large language models: From visual embed to language reasoning. *arXiv preprint*, 2024. 2
- [74] Xi Yin, Yuchen Zhou, Jianfeng Gao, et al. Cq-vqa: Visual question answering on categorized questions. In *European Conference on Computer Vision (ECCV)*, 2020. 2
- [75] Xiang Yin, Lei Zhu, Dan Xu, Hengshuang Zhao, and Jing Shao. Mmhal-bench: A benchmark for multimodal hallucination evaluation. *arXiv preprint arXiv:2310.01817*, 2023. 1, 7, 8
- [76] Rowan Zellers, Yonatan Bisk, Ali Farhadi, and Yejin Choi. From recognition to cognition: Visual commonsense reasoning. In *CVPR*, 2019. 6, 7
- [77] Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohan S. Kankanhalli. Minigt-v2: Efficient vision-language alignment. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023. 6, 8