
CREG: Compass Relational Evidence for Interpreting Spatial Reasoning in Vision-Language Models

Kaizhen Tan

Heinz College of Information Systems and Public Policy
Carnegie Mellon University
Pittsburgh, PA, USA
kaizhent@cmu.edu

Abstract

Vision-language models (VLMs) achieve strong performance on spatial reasoning benchmarks, yet how they encode directional relations remains opaque. Standard attribution methods such as GradCAM and attention rollout produce flat saliency maps that highlight *where* the model attends, but not *what direction* it infers between objects. We introduce CREG (Compass Relational Evidence Graph), a training-free interpretability framework that projects multi-layer contrastive Grad \times Act attributions into a reference-centered polar coordinate system, yielding a directional evidence distribution $\hat{P}(\theta)$ over compass sectors. We propose three complementary evaluation metrics—Direction Alignment Error (DAE), Edge Accuracy (EA), and Causal Occlusion Score (COS)—to quantify the faithfulness of directional explanations. On Qwen2-VL-7B across two benchmarks (VSR and COCO-Pairs), CREG substantially outperforms all standard attribution baselines: on COCO-Pairs, prediction-targeted CREG achieves DAE 55.5° and EA 0.553, reducing angular error by 16.1° over attention rollout (71.6°/0.433). On the smaller Qwen2-VL-2B, CREG’s advantage diminishes, suggesting that its contrastive multi-layer approach benefits most from the more structured spatial representations that emerge at larger model scales. Prediction-targeted attribution outperforms ground-truth-targeted, suggesting that CREG better aligns with the model’s internal decision signal rather than providing post-hoc rationalization. Causal occlusion experiments on 540 samples across both datasets support faithfulness (COS $\geq +0.42$). We further show that contrastive gradient targeting contributes a 20.5° DAE improvement.

1 Introduction

Vision-language models (VLMs) such as Qwen2-VL [Wang et al., 2024], LLaVA [Liu et al., 2024], and InternVL [Chen et al., 2024] have achieved remarkable progress on spatial reasoning tasks, correctly identifying that “the cup is to the left of the bottle” or “the cat is above the keyboard.” Yet the internal mechanisms by which these models encode and reason about directional spatial relations remain largely opaque.

Existing interpretability tools—GradCAM [Selvaraju et al., 2017], attention rollout [Abnar and Zuidema, 2020], and gradient-based saliency [Sundararajan et al., 2017]—reveal *where* a model attends in an image, but not *what directional relation* it infers between two objects. A saliency heatmap may highlight both the cup and the bottle, but it cannot tell us whether the model is encoding “left-of,” “right-of,” or something else entirely. This gap between “attention location” and “relational evidence” is precisely where spatial interpretability breaks down.

We introduce CREG (Compass **R**elational **E**vidence **G**raph), a training-free framework that converts standard token-level attributions into *directional evidence*. The key idea is a polar coordinate

projection: given a reference object A and a target object B in an image, CREG projects multi-layer contrastive Grad \times Act attributions into a reference-centered polar coordinate system, producing a normalized directional distribution $\hat{P}(\theta)$ over K compass sectors. This distribution directly answers the question: “In which direction does the model’s spatial evidence concentrate?”

CREG builds on three design choices, each validated by ablation:

1. **Contrastive gradient target.** Rather than computing gradients with respect to a single class logit, CREG uses the difference $z_{\text{gt}} - z_{\text{neg}}$ (where z_{neg} is the predicted or second-highest logit; see §3.2), which isolates directional evidence from class-general features. This contributes a 20.5° reduction in angular error.
2. **Multi-layer aggregation.** Attribution signals from multiple transformer layers are combined with magnitude-based softmax weights (no learned parameters), capturing both low-level spatial features and high-level relational abstractions.
3. **Gaussian-weighted polar binning.** Vision-token attributions are aggregated into K angular sectors with distance-dependent Gaussian weighting, producing a smooth, interpretable compass distribution.

To rigorously evaluate directional explanations, we propose three metrics: **Direction Alignment Error (DAE)**, which measures the angular gap between the predicted and true compass directions; **Edge Accuracy (EA)**, which tests whether the peak sector falls within the correct quadrant; and **Causal Occlusion Score (COS)**, which validates faithfulness by comparing model confidence drops when occluding the indicated vs. opposite directions.

On Qwen2-VL-7B across two benchmarks—VSR [Liu et al., 2023] (240 samples) and COCO-Pairs (300 samples with native bounding boxes)—CREG substantially outperforms all baselines. On COCO-Pairs, prediction-targeted CREG achieves DAE 55.5° and EA 0.553, reducing angular error by 16.1° compared to attention rollout ($71.6^\circ / 0.433$). On the smaller Qwen2-VL-2B, the advantage diminishes: baselines perform comparably, suggesting that CREG’s contrastive multi-layer approach benefits from the more structured representations that emerge at larger scales. Causal occlusion experiments on 540 total samples across both datasets support faithfulness ($\text{COS} \geq +0.42$). Prediction-targeted attribution outperforms ground-truth-targeted, suggesting better alignment with the model’s internal decision signal.

Our contributions are:

- CREG, a training-free interpretability framework that converts VLM attributions into reference-centered directional evidence distributions.
- A directional evaluation protocol (DAE, EA, COS) for quantifying the faithfulness of spatial explanations.
- Empirical analysis across 2 models, 2 datasets, and 7 baselines showing that CREG achieves the best directional alignment on the larger model, with causal occlusion supporting faithfulness and prediction-targeted attribution suggesting genuine alignment with the model’s decision signal.

2 Related Work

Spatial reasoning in VLMs. Spatial relation understanding has emerged as a key challenge for VLMs. Benchmarks such as VSR [Liu et al., 2023], WhatsUp [Kamath et al., 2023], Spatial-Bench [Cai et al., 2024], and SpatialSense+ [Wen et al., 2024] evaluate whether models can correctly classify binary or multi-class directional relations (*left, right, above, below*). Recent work has shown that VLMs often fail at spatial reasoning despite strong general VQA performance [Tong et al., 2024, Kamath et al., 2023], and that targeted interventions such as confidence-based attention scaling [Yang et al., 2025] can substantially improve accuracy. Our work is complementary: rather than improving accuracy, we aim to *interpret* how VLMs represent spatial direction internally.

Attribution methods and faithfulness evaluation. Standard attribution methods project model computations onto input features to identify influential regions. GradCAM [Selvaraju et al., 2017] computes gradient-weighted activation maps; attention rollout [Abnar and Zuidema, 2020] aggregates

attention weights across layers; integrated gradients [Sundararajan et al., 2017] accumulate gradients along an interpolation path; transformer-specific methods [Chefer et al., 2021] propagate relevance through attention layers. These methods produce 2D saliency maps indicating *where* the model looks. A growing body of work has questioned whether such saliency explanations are truly faithful [Adebayo et al., 2018], proposing evaluation protocols based on input perturbation [Petsiuk et al., 2018, Hooker et al., 2019]. For relational tasks, knowing *where* is insufficient—we need to know *what relation* the model infers between located regions. CREG addresses this gap by projecting token-level attributions into a relation-centric coordinate system, with COS providing a direction-specific faithfulness test.

Spatial and relational interpretability. Recent work has begun to bridge attribution and relational reasoning. GrAInS [Li et al., 2025] combines contrastive gradient attribution with activation steering for spatial tasks. RelatiViT [Zhang et al., 2024] adds relative position encodings to improve spatial relation prediction. Our approach differs from both: CREG is purely analytical (no model modification), operates as a post-hoc interpretability tool, and produces directional evidence distributions rather than attention maps or modified predictions. The polar projection idea draws on classical spatial cognition research [Marr, 1982], where directional reference frames are fundamental to spatial representation.

3 Method: CREG

Given an image containing a reference object A and a target object B with known bounding boxes, CREG computes a directional evidence distribution $\hat{P}(\theta)$ that summarizes how the VLM’s internal representations encode the spatial relation between A and B . The method is entirely training-free and architecture-agnostic, requiring only access to hidden states and gradients.

3.1 Problem Setup

We consider the task of spatial relation classification. Given an image I , object names and bounding boxes for A (reference) and B (target), the VLM is prompted to classify the relation into one of four cardinal directions: *left*, *right*, *above*, *below*. The VLM produces logits z_c for each class $c \in \{1, 2, 3, 4\}$ at the last token position.

3.2 Attribution Targets

We define two attribution modes and a contrastive mechanism. Let z_c denote the logit for class c and let $\hat{c} = \arg \max_c z_c$ be the model’s prediction.

GT-targeted attribution. The gradient target is the ground-truth class logit: $\tau_{\text{gt}} = z_{\text{gt}}$. This asks: “what features support the correct answer?”

Prediction-targeted attribution. The gradient target is the predicted class logit: $\tau_{\text{pred}} = z_{\hat{c}}$. This asks: “what features drive the model’s actual decision?” When $\hat{c} = \text{gt}$, both modes are identical.

Contrastive formulation. In either mode, CREG uses a **contrastive target** that subtracts the strongest competing logit to isolate direction-specific evidence from class-general features:

$$\tau = z_{\text{tgt}} - z_{\text{neg}}, \quad z_{\text{neg}} = \begin{cases} z_{\hat{c}} & \text{if } \hat{c} \neq \text{tgt (gt-targeted)} \\ z_{2\text{nd}} & \text{otherwise} \end{cases} \quad (1)$$

where $z_{\text{tgt}} \in \{z_{\text{gt}}, z_{\hat{c}}\}$ depending on the chosen mode, and $z_{2\text{nd}}$ is the second-highest logit. The gradient $\nabla_h \tau$ isolates features that *distinguish* the targeted direction from the model’s best alternative, suppressing direction-invariant contributions.

3.3 Multi-Layer Grad×Act Aggregation

For a set of transformer layers $\mathcal{L} = \{l_1, \dots, l_L\}$ (default: last 4 layers), we compute per-token relevance scores at each layer:

$$r_j^{(l)} = \left| \sum_d g_{j,d}^{(l)} \cdot h_{j,d}^{(l)} \right|, \quad j \in \mathcal{V} \quad (2)$$

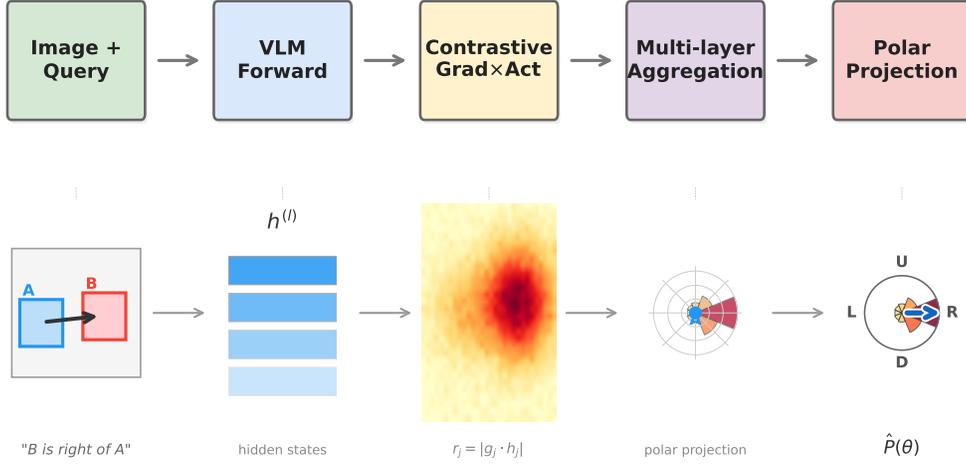


Figure 1: CREG pipeline overview. An image with reference/target objects is fed through a VLM. Contrastive Grad \times Act signals from multiple layers are aggregated and projected into a reference-centered polar coordinate system, producing the compass distribution $\hat{P}(\theta)$.

where $g^{(l)} = \nabla_{h^{(l)}} \tau$ is the gradient at layer l , $h^{(l)}$ is the hidden state, and \mathcal{V} is the set of vision token indices. Each layer’s relevance vector is normalized to $[0, 1]$.

Layers are aggregated with softmax weights proportional to signal magnitude:

$$w_l = \frac{\exp(\max_j r_j^{(l)})}{\sum_{l'} \exp(\max_j r_j^{(l')})}, \quad r_j = \sum_l w_l \cdot r_j^{(l)} \quad (3)$$

3.4 Polar Projection and Compass Binning

The combined relevance scores $\{r_j\}_{j \in \mathcal{V}}$ are mapped to a 2D token grid of size $H_g \times W_g$ (using the VLM’s image grid dimensions when available, or inferred via factorization of $|\mathcal{V}|$). Each grid cell (u, v) with $u \in \{1, \dots, H_g\}$, $v \in \{1, \dots, W_g\}$ is assigned pixel coordinates (x_{uv}, y_{uv}) by linearly mapping grid positions to image dimensions, and inherits relevance r_{uv} from the corresponding vision token.

We define a polar coordinate system centered on the reference object A ’s bounding box center (x_A, y_A) :

$$\theta_{uv} = \text{atan2}(-(y_{uv} - y_A), x_{uv} - x_A), \quad \rho_{uv} = \sqrt{(x_{uv} - x_A)^2 + (y_{uv} - y_A)^2} \quad (4)$$

where the convention $\theta = 0^\circ$ corresponds to “right” and $\theta = 90^\circ$ to “up.”

The compass distribution is computed by aggregating relevance into K uniform angular sectors (default $K = 8$, sector width 45°) with Gaussian distance weighting:

$$P(\theta_k) = \frac{\sum_{(u,v) \in S_k} r_{uv} \cdot \exp(-\rho_{uv}^2/2\sigma^2)}{\sum_{k'} \sum_{(u,v) \in S_{k'}} r_{uv} \cdot \exp(-\rho_{uv}^2/2\sigma^2)} \quad (5)$$

where S_k is the set of grid cells whose angle θ_{uv} falls in sector k , $\sigma = \sigma_r \cdot r_{\max} \cdot d_{AB}$ controls the spatial spread, and $r_{\max} = \rho_r \cdot d_{AB}$ defines the maximum radius of influence (d_{AB} is the Euclidean distance between A and B centers, $\sigma_r = 0.6$, $\rho_r = 2.0$).

The output $\hat{P}(\theta) = [P(\theta_1), \dots, P(\theta_K)]$ is a normalized distribution over compass directions, directly interpretable as “the model’s directional evidence points $X\%$ to the right, $Y\%$ upward, etc.”

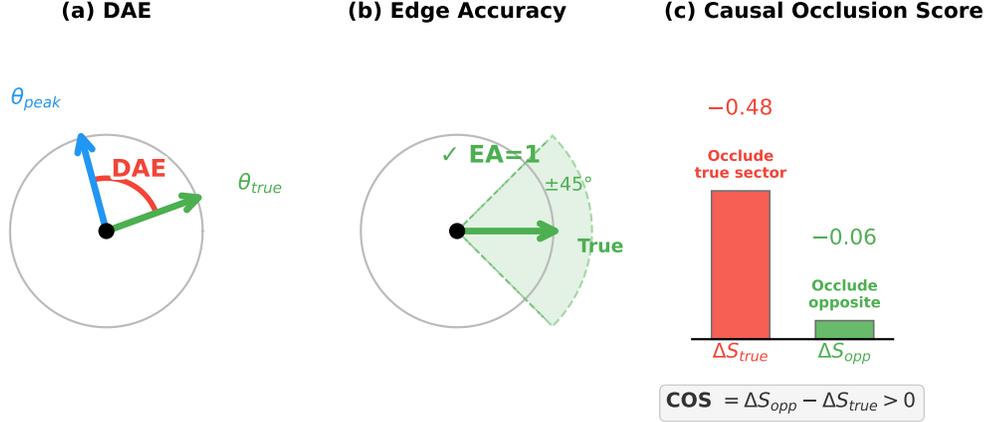


Figure 2: Intuition behind our evaluation metrics. (a) DAE: angular distance between the compass peak and the true $A \rightarrow B$ direction. (b) EA: whether the peak falls within $\pm 45^\circ$ of the truth. (c) COS: occluding the true-direction sector should cause a larger confidence drop than occluding the opposite sector.

4 Evaluation Protocol

We propose three metrics to evaluate directional attribution quality. Each metric captures a different aspect of faithfulness and is validated against expected behavior under a random baseline.

Direction Alignment Error (DAE). DAE measures the circular angular distance between the predicted compass peak direction and the geometric ground truth:

$$\text{DAE} = \left| \left((\hat{\theta}_{\text{peak}} - \theta_{\text{true}} + 180^\circ) \bmod 360^\circ \right) - 180^\circ \right| \quad (6)$$

where $\hat{\theta}_{\text{peak}}$ is the center angle of the sector with maximum $P(\theta_k)$, and $\theta_{\text{true}} = \text{atan2}(-(y_B - y_A), x_B - x_A)$ is the geometric direction from A to B . $\text{DAE} \in [0^\circ, 180^\circ]$; lower is better. Under uniform random assignment over $K = 8$ sectors, the expected DAE is $\approx 90^\circ$.

Edge Accuracy (EA). EA is the fraction of samples where the peak direction falls within $\pm 45^\circ$ of the true direction (i.e., the correct 90° -wide quadrant):

$$\text{EA} = \frac{1}{N} \sum_{i=1}^N \mathbf{1}[\text{DAE}_i \leq 45^\circ] \quad (7)$$

Under random assignment with $K = 8$, the expected EA is ≈ 0.25 . EA provides a binary, easily interpretable measure of whether the attribution “got the direction right.”

Causal Occlusion Score (COS). COS tests whether the attributed direction is *causally* relevant to the model’s prediction. We occlude the image region in the indicated sector (filling with grey) and re-run inference. Let ΔS_k denote the change in log-probability of the ground-truth class after occluding sector k :

$$\Delta S_k = \log p_k(\text{gt}) - \log p_0(\text{gt}), \quad \text{COS} = \Delta S_{\text{opp}} - \Delta S_{\text{true}} \quad (8)$$

where $\log p_0(\text{gt})$ is the unoccluded baseline log-probability, and $\log p_k(\text{gt})$ is the log-probability after occluding sector k . If the attribution is faithful, occluding the true direction should cause a larger confidence drop than occluding the opposite direction, i.e., $\Delta S_{\text{true}} < \Delta S_{\text{opp}} \leq 0$, yielding $\text{COS} > 0$. For example, if occluding the true sector causes a -0.47 drop and the opposite sector causes only a -0.06 drop, then $\text{COS} = -0.06 - (-0.47) = +0.41$.

Metric validation. We verify that all metrics behave as expected under controlled conditions. A random-sector baseline achieves $\text{DAE} \approx 90^\circ$, $\text{EA} \approx 0.22$, consistent with theoretical expectations

(for $K = 8$ sectors with uniformly random peaks, the expected DAE is $\frac{1}{K} \sum_{i=0}^{K-1} \min(i \cdot 45^\circ, 360^\circ - i \cdot 45^\circ) = 90^\circ$, and the expected EA is $\frac{2}{K} = 0.25$; the slight deviation from 0.25 reflects the finite sample size and non-uniform true-direction distribution). On synthetic images with known object placements, CREG achieves DAE = 52.7° , EA = 0.75, and VLM accuracy = 1.0, confirming that the pipeline functions correctly when the spatial signal is unambiguous.

5 Experiments

5.1 Setup

Models. We evaluate on two VLMs from the Qwen2-VL family [Wang et al., 2024]: **Qwen2-VL-7B-Instruct** (8.29B params) and **Qwen2-VL-2B-Instruct** (2.21B params). Both use bfloat16 precision on a single NVIDIA A800-80GB GPU with no fine-tuning. Comparing two model scales allows us to test whether CREG’s advantages depend on model capacity.

Datasets. We evaluate on two benchmarks:

- **VSR** [Liu et al., 2023]: 240 COCO-based spatial reasoning samples with bounding boxes obtained by matching subject/object names to COCO instance annotations (21.2% coverage, 100% box validity; see Appendix A.2).
- **COCO-Pairs**: 300 samples constructed directly from COCO val2017 object annotations, with native bounding boxes (no heuristic matching). We select pairs of different-category objects with clear directional dominance ($> 1.5 \times$ displacement ratio), balanced across four classes (75 per class).

COCO-Pairs serves as a complementary benchmark with guaranteed box quality and balanced class distribution.

Baselines. All baselines share the same polar projection and compass binning post-processing (Eq. 5 with identical K , σ , and r_{\max}). The *only* difference between methods is the per-token relevance signal r_j fed into the projection. We compare:

- **Random**: uniform random relevance.
- **GradCAM** [Selvaraju et al., 2017]: gradient-weighted activation map from the last vision encoder layer, upsampled to the token grid. Target: z_{gt} .
- **Gradient norm**: ℓ_2 norm of $\nabla_{h^{(-2)}} z_{\text{gt}}$ per vision token.
- **Integrated Gradients** [Sundararajan et al., 2017]: accumulated gradients along a linear interpolation path (50 steps) from a zero baseline to the actual hidden state at layer -2 . Target: z_{gt} .
- **Attention rollout** [Abnar and Zuidema, 2020]: multiplicative aggregation of attention matrices across all layers, extracting the last-token-to-vision-token row.
- **Single-layer Grad \times Act**: Grad \times Act at layer -2 only (no multi-layer aggregation, no contrastive target).
- **Geometry oracle**: relevance set to 1 at the true B location and 0 elsewhere (theoretical ceiling).

5.2 Main Results

Table 1 presents the main results on Qwen2-VL-7B. CREG (pred-targeted) achieves the best directional alignment on both benchmarks. On COCO-Pairs, the improvement over attention rollout is particularly strong: DAE 55.5° vs. 71.6° (-16.1°), EA 0.553 vs. 0.433 ($+0.120$). The consistent advantage across two independent datasets confirms that CREG’s gains are not dataset-specific.

Table 2 compares model scales. On Qwen2-VL-7B, CREG clearly outperforms all baselines (-6.9° DAE on VSR, -14.5° on COCO-Pairs vs. best baseline). On Qwen2-VL-2B, however, CREG does *not* outperform the best baseline: on VSR the best baseline achieves lower DAE (69.5° vs. 71.9°), and on COCO-Pairs both methods tie at $63.8^\circ/62.3^\circ$ DAE. This scale-dependent pattern suggests that

Table 1: Directional attribution on **Qwen2-VL-7B** across two benchmarks. 95% bootstrap CIs for CREG. Pred-targeted CREG consistently achieves the best DAE and EA.

Method	VSR (n=240)		COCO-Pairs (n=300)	
	DAE↓	EA↑	DAE↓	EA↑
Geo. oracle	0.0	1.000	0.0	1.000
CREG (pred)	64.8 [58, 72]	.458 [.39, .53]	55.5 [50, 61]	.553 [.50, .61]
CREG (gt)	68.3 [61, 76]	.438 [.38, .50]	57.1 [52, 62]	.523 [.47, .58]
Attn. rollout	75.2	.392	71.6	.433
Grad norm	86.3	.300	76.4	.377
Integ. Grad	86.9	.292	78.2	.373
GradCAM	96.8	.196	84.9	.310
Random	91.8	.221	91.2	.217

Table 2: Cross-model comparison: **2B vs. 7B** on both benchmarks (gt-targeted CREG and strongest baseline).

Model	Method	VSR		COCO-Pairs	
		DAE↓	EA↑	DAE↓	EA↑
7B	CREG	68.3	.438	57.1	.523
	Best baseline	75.2	.392	71.6	.433
2B	CREG	71.9	.429	63.8	.503
	Best baseline	69.5	.442	62.3	.503

CREG’s contrastive multi-layer approach benefits from the more structured spatial representations that emerge in larger models. In smaller models, spatial features may be less differentiated across layers, reducing the advantage of multi-layer contrastive aggregation.

5.3 GT-Targeted vs. Prediction-Targeted Attribution

Table 3 addresses a key methodological concern: does CREG explain what the model *actually computes*, or merely find evidence for the correct answer? On both datasets, pred-targeted CREG outperforms gt-targeted, especially on incorrect samples (VSR: 66.8° vs. 74.7°; COCO-Pairs: 61.7° vs. 64.5°). This suggests that CREG better captures the model’s internal decision signal when targeting the predicted class. On correct samples, both modes yield identical results as expected, since $\hat{c} = \text{gt}$.

5.4 Causal Faithfulness

Table 4 provides evidence for causal relevance of the attributed direction on both benchmarks. Occluding the CREG-indicated direction substantially reduces model confidence ($\Delta S_{\text{true}} \approx -0.47$), while occluding the opposite direction has minimal effect ($\Delta S_{\text{opp}} \approx -0.04$). The consistent positive COS across datasets (+0.417 on VSR, +0.446 on COCO-Pairs) suggests that CREG identifies a direction that is relevant to the model’s prediction, though we note that this uses a single grey-fill occlusion strategy and does not constitute a definitive causal proof.

5.5 Ablation Study

Table 5: removing contrastive targeting degrades DAE by 20.5°, confirming it is the most critical component. Multi-layer aggregation contributes 2.6°; K is robust.

Table 3: Correct-only vs. incorrect-only analysis on Qwen2-VL-7B. Pred-targeted attribution explains the model’s actual decision.

Subset	GT-targeted		Pred-targeted	
	DAE↓	EA↑	DAE↓	EA↑
<i>VSR (VLM ACC = 55.8%)</i>				
All (n=240)	68.3	.438	64.8	.458
Correct (n=134)	63.2	.515	63.2	.515
Incorrect (n=106)	74.7	.340	66.8	.387
<i>COCO-Pairs (VLM ACC = 41.3%)</i>				
All (n=300)	57.1	.523	55.5	.553
Correct (n=124)	46.6	.645	46.6	.645
Incorrect (n=176)	64.5	.438	61.7	.489

Table 4: Causal Occlusion Score on both benchmarks (Qwen2-VL-7B, full datasets).

	n	ΔS_{true}	ΔS_{opp}	COS↑
VSR	240	-0.476	-0.060	+0.417
COCO-Pairs	300	-0.464	-0.018	+0.446

6 Analysis and Limitations

6.1 VLMs Lack Flip Equivariance in Spatial Representations

A natural desideratum for spatial attribution is *flip equivariance*: horizontally flipping an image should mirror the compass distribution accordingly. We test this by flipping images and bounding boxes, re-running CREG, and computing the Pearson correlation between the original and correctly mirrored compass distributions.

Surprisingly, the correlation is near zero ($r = -0.065$, $n = 60$). This suggests that VLMs encode spatial direction through *position-dependent features* (e.g., absolute patch positions, asymmetric attention patterns) rather than geometric invariants. This observation is consistent with prior findings that VLMs are sensitive to object position [Tong et al., 2024] and suggests that improving spatial equivariance may require architectural changes to VLM vision encoders.

6.2 Per-Class Analysis

The per-class VLM accuracy on VSR shows an asymmetry: *left* (73.8%) and *right* (67.4%) are substantially more accurate than *above* (58.7%) and *below* (30.0%). This horizontal-vertical gap may reflect training data biases or differences in how vision transformers encode vertical vs. horizontal spatial relationships. Notably, the *below* class is particularly weak, suggesting a systematic blind spot in vertical reasoning that merits further investigation. Figure 4 visualizes this asymmetry.

6.3 Does CREG Interpret Relational Reasoning or Merely Localize Objects?

A natural concern is whether CREG’s directional scores primarily reflect *target-object localization*—i.e., simply highlighting object B and letting the polar projection do the geometric work—rather than interpreting internal relational reasoning. Several observations address this concern, though we acknowledge it is not fully resolved.

First, all baselines share the same polar projection post-processing (Eq. 5). Since the only difference is the per-token relevance signal, CREG’s improvements over baselines cannot be attributed to the polar geometry alone—they must reflect a better relevance signal.

Second, the geometry oracle (which places all relevance at object B ’s location) achieves perfect $\text{DAE} = 0^\circ$, showing that pure localization *is* sufficient for directional scoring. However, the gap between CREG ($\text{DAE} 55.5^\circ$) and the oracle ($\text{DAE} 0^\circ$) indicates that CREG does *not* simply localize

Table 5: Ablation on VSR (60 samples, Qwen2-VL-7B). Default: $K = 8$, 4-layer, contrastive.

Configuration	DAE↓	EA↑
CREG (default)	61.8	.483
$K = 4$ sectors	60.9	.517
Single layer ($l = -2$)	64.4	.467
No contrastive	82.3	.333

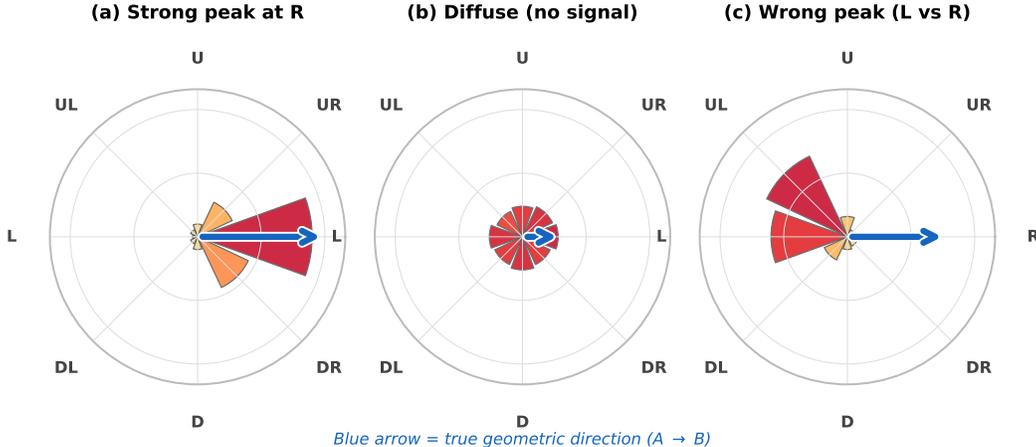


Figure 3: Compass distribution $\hat{P}(\theta)$ examples. Red arrow = true direction. (a) Strong peak at correct direction. (b) Diffuse: no clear directional signal. (c) Wrong peak direction.

B —it captures a diffuse, multi-region attribution pattern that partially but imperfectly aligns with the true direction.

Third, the prediction-targeted vs. gt-targeted comparison (§5.3) shows that CREG’s directional signal tracks the model’s actual prediction, not just the geometric ground truth. On incorrect samples, pred-targeted DAE (66.8°) is lower than gt-targeted (74.7°), suggesting that the relevance pattern reflects the model’s internal decision signal rather than a generic object-localization response.

Nevertheless, we acknowledge that disentangling relational interpretation from direction-aware localization remains an open challenge. Future work should include control experiments such as: uniform relevance inside the target box (pure localization baseline), randomized reference centers, and subject/object swap tests to more rigorously isolate relational evidence from geometric artifacts.

6.4 On the Choice of Attribution Target

A critical methodological consideration is whether attribution should target the ground-truth class or the model’s actual prediction. When VLM accuracy is below 100%, gt-targeted attribution risks extracting evidence for the correct answer that the model does not actually rely on. Our experiments (§5.3) show that prediction-targeted CREG achieves lower DAE than gt-targeted, especially on incorrect samples (DAE 66.8° vs. 74.7°). This suggests that prediction-targeted attribution better aligns with the model’s internal decision signal, and we recommend it as the default for interpretability applications. On correct samples, both modes produce identical results, since the predicted and ground-truth classes coincide.

6.5 Limitations

- **Single model family.** We evaluate on Qwen2-VL (7B and 2B). While CREG requires only hidden states and gradients, generalization to architecturally different VLMs (LLaVA, InternVL) remains to be verified.

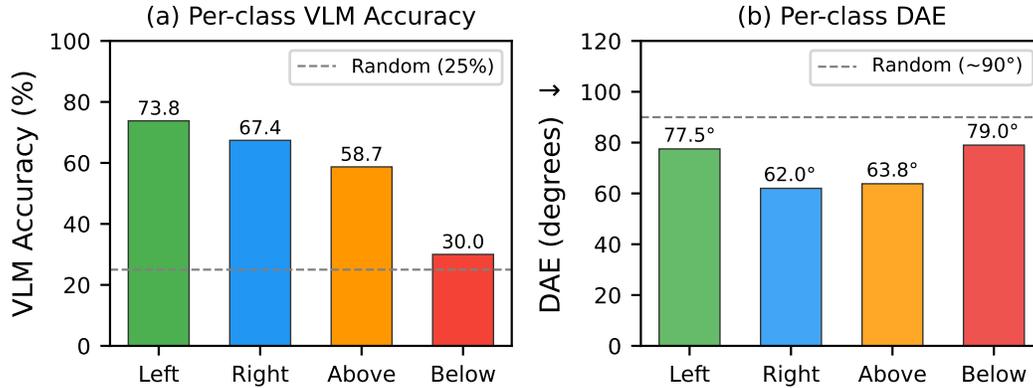


Figure 4: Per-class analysis on VSR. (a) VLM accuracy is much higher for horizontal relations than vertical ones, with *below* at only 30%. (b) DAE shows a similar pattern—vertical relations have higher angular error. Dashed lines indicate random baselines.

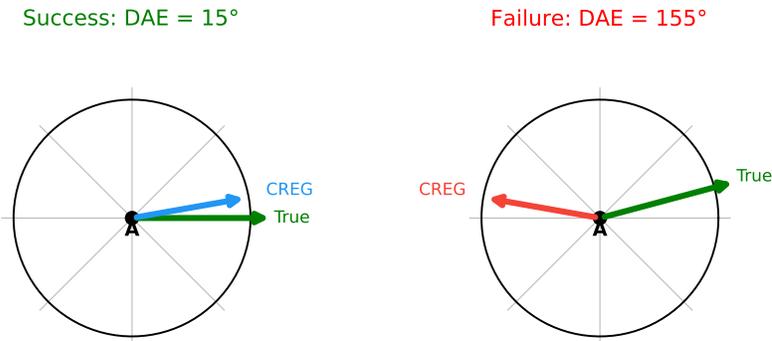


Figure 5: Success vs. failure illustration. Green arrow = true direction; colored arrow = CREG peak. Left: CREG correctly identifies the right direction (DAE = 15°). Right: CREG peak points opposite to the true direction (DAE = 155°).

- **Scale-dependent advantage.** On Qwen2-VL-2B, CREG does not outperform the best baseline, indicating that the method’s benefits may depend on the model having sufficiently structured spatial representations.
- **Two benchmarks from COCO.** Both VSR and COCO-Pairs are based on COCO images. Evaluation on datasets with different image distributions (e.g., synthetic scenes, indoor environments) would strengthen generalization claims.
- **Moderate absolute improvement.** CREG’s best DAE (55.5° on COCO-Pairs) is substantially better than baselines (71–91°) but remains far from the geometry oracle (0°), reflecting fundamental challenges in interpreting spatial reasoning.
- **Four cardinal relations only.** Depth, distance, and containment are not addressed.
- **COS limitations.** The occlusion uses a single grey-fill strategy with only one control condition (opposite-sector occlusion). Different intervention operators (blur, inpainting, patch drop) and additional controls (adjacent-sector occlusion) may yield different results. We do not claim definitive causal proof; COS provides supportive evidence that the attributed direction is relevant to the model’s prediction.
- **Statistical coverage.** We report 95% bootstrap confidence intervals for CREG in the main table but do not provide formal significance tests for all pairwise comparisons. The ablation study uses 60 samples; while directionally informative, wider-scale ablations would strengthen the component analysis.

7 Conclusion

We introduced CREG, a training-free framework for interpreting spatial reasoning in vision-language models. By projecting contrastive multi-layer Grad \times Act attributions into a reference-centered polar coordinate system, CREG transforms opaque token-level saliency into interpretable directional evidence distributions. On Qwen2-VL-7B across two benchmarks (VSR and COCO-Pairs), CREG substantially outperforms all standard attribution baselines on directional alignment metrics, with causal occlusion experiments supporting the faithfulness of the attributed directions on 540 samples. On the smaller Qwen2-VL-2B, the advantage diminishes, suggesting that CREG’s approach benefits from the more structured spatial representations that emerge at larger model scales. Prediction-targeted attribution achieves lower DAE than gt-targeted, suggesting better alignment with the model’s internal decision signal.

CREG opens several directions for future work: extending to depth and distance relations, evaluating on architecturally diverse VLMs (e.g., LLaVA, InternVL) to test cross-architecture generality, testing on non-COCO-derived benchmarks to strengthen generalization claims, and designing control experiments (e.g., randomized reference centers, subject/object swap tests) to more rigorously disentangle relational interpretation from direction-aware localization. More robust causal validation through alternative intervention operators (blur, inpainting) would further strengthen the faithfulness evidence.

References

- Samira Abnar and Willem Zuidema. Quantifying attention flow in transformers. In *ACL*, 2020.
- Julius Adebayo, Justin Gilmer, Michael Muelly, Ian Goodfellow, Moritz Hardt, and Been Kim. Sanity checks for saliency maps. In *NeurIPS*, 2018.
- Yuxuan Cai et al. Spatialbench: Spatial understanding benchmark for vlms. In *arXiv preprint*, 2024.
- Hila Chefer, Shir Gur, and Lior Wolf. Transformer interpretability beyond attention visualization. In *CVPR*, 2021.
- Zhe Chen, Jiannan Wu, Wenhai Wang, Weijie Su, Guo Chen, Sen Xing, Muyan Zhong, Qinglong Zhang, Xizhou Zhu, Lewei Lu, et al. Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. In *CVPR*, 2024.
- Sara Hooker, Dumitru Erhan, Pieter-Jan Kindermans, and Been Kim. A benchmark for interpretability methods in deep neural networks. In *NeurIPS*, 2019.
- Amita Kamath, Jack Hessel, and Kai-Wei Chang. What’s “up” with vision-language models? investigating their struggle with spatial reasoning. In *EMNLP*, 2023.
- Xiangyu Li et al. Grains: Gradient attribution for interpretable spatial reasoning. *arXiv preprint*, 2025.
- Fangyu Liu, Guy Emerson, and Nigel Collier. Visual spatial reasoning. In *TACL*, 2023.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. In *NeurIPS*, 2024.
- David Marr. *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. MIT Press, 1982.
- Vitali Petsiuk, Abir Das, and Kate Saenko. RISE: Randomized input sampling for explanation of black-box models. In *BMVC*, 2018.
- Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *ICCV*, 2017.
- Mukund Sundararajan, Ankur Taly, and Qiqi Yan. Axiomatic attribution for deep networks. In *ICML*, 2017.

- Shengbang Tong, Zhuang Liu, Yuexiang Zhai, Yi Ma, Yann LeCun, and Saining Xie. Eyes wide shut? exploring the visual shortcomings of multimodal llms. In *CVPR*, 2024.
- Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, et al. Qwen2-vl: Enhancing vision-language model’s perception of the world at any resolution. In *arXiv preprint arXiv:2409.12191*, 2024.
- Xiuye Wen et al. Spatialsense+: Enhanced spatial relation understanding. In *ICLR*, 2024.
- Zhiqiang Yang et al. Adaptvis: Adapting vision transformers for visual spatial reasoning via inference-time confidence-based attention scaling. In *ICML*, 2025.
- Wei Zhang et al. Relativit: Relative position enhanced vision transformer for spatial relation prediction. In *ICLR*, 2024.

A Appendix

A.1 Prompt Sweep

We evaluated three prompt templates on a 30-sample development subset of VSR (Table 6). Prompt A achieves the highest VLM accuracy and is used for all main experiments.

Table 6: Prompt sweep results on VSR development subset (30 samples).

Prompt Template	VLM ACC	DAE
A: “In this image, where is the {tgt} relative to the {ref}? Choose one option: 1) to the left, 2) to the right, 3) above, 4) below. Answer with just the number.”	0.567	57.9°
B: “Look at the image. Is the {tgt} to the left (1), right (2), above (3), or below (4) the {ref}? Answer with one number.”	0.367	64.3°
C: “I can see two objects: {ref} and {tgt}. What is the spatial position of {tgt} relative to {ref}? Options: 1=left, 2=right, 3=above, 4=below. Answer:”	0.533	62.9°

A.2 VSR Bounding Box Matching Audit

VSR annotations do not include bounding boxes. We obtain boxes by matching subject/object names from VSR captions to COCO instance segmentation categories. Table 7 reports the matching quality.

Table 7: VSR bounding box matching audit. Coverage indicates the fraction of positive VSR samples where both subject and object names match a COCO category. Box validity indicates the fraction of matched samples with non-degenerate boxes.

Split	Positive	Matched	Coverage / Validity
Test	1181	240	20.3% / 100%
Dev	564	121	21.5% / 100%
Train	3876	828	21.4% / 100%
Total	5621	1189	21.2% / 100%

The 21% coverage reflects our conservative matching strategy: we require exact or substring matches between VSR entity names and COCO category labels. This produces a high-confidence subset where bounding boxes are reliable, at the cost of reduced sample size.

A.3 Implementation Details

CREG is implemented in Python using PyTorch. Key hyperparameters:

- Compass sectors: $K = 8$ (sector width 45°)
- Hidden-state layers: $\{-2, -3, -4, -5\}$ (last 4 layers)
- Gaussian $\sigma = 0.6 \times r_{\max}$, where $r_{\max} = 2.0 \times d_{AB}$
- Gradient target: contrastive (Eq. 1)

All baselines use the same polar projection framework (Eq. 5) with the same K , σ , and r_{\max} settings. The only difference is the per-token relevance signal fed into the projection. Baseline-specific details:

- **GradCAM:** We extract the gradient-weighted activation map from the last layer of the vision encoder (ViT). The spatial map is upsampled to the token grid resolution via bilinear interpolation. The gradient target is z_{gt} .
- **Integrated Gradients:** We use 50 interpolation steps along a linear path from a zero baseline to the actual hidden state at layer -2 . The accumulated gradient magnitudes provide per-token relevance. The gradient target is z_{gt} .

- **Attention rollout:** Attention matrices from all layers are multiplied cumulatively (with identity residual added at each layer). We extract the row corresponding to the last text token and restrict to vision-token columns.
- **Gradient norm:** The ℓ_2 norm of $\nabla_{h^{(-2)}} z_{\text{gt}}$ is computed per vision token at layer -2 .

Experiments run on a single NVIDIA A800-80GB GPU. The full 240-sample VSR evaluation (including all neural baselines) completes in approximately 4 minutes. No hyperparameter tuning was performed on the test sets; all hyperparameters (K , σ_r , ρ_r , layer set) were fixed before evaluation based on preliminary experiments on a separate 30-sample development set.