

# Graph-of-Mark: Promote Spatial Reasoning in Multimodal Language Models with Graph-Based Visual Prompting

Giacomo Frisoni\*, Lorenzo Molfetta\*, Mattia Buzzoni\*, Gianluca Moro\*

Department of Computer Science and Engineering, University of Bologna  
 {giacomo.frisoni, lorenzo.molfetta, gianluca.moro}@unibo.it,  
 mattia.buzzoni@studio.unibo.it

## Abstract

Recent advances in training-free visual prompting, such as Set-of-Mark, have emerged as a promising direction for enhancing the grounding capabilities of multimodal language models (MLMs). These techniques operate by partitioning the input image into object regions and annotating them with marks—predominantly boxes with numeric identifiers—before feeding the augmented image to the MLM. However, these approaches treat marked objects as isolated entities, failing to capture the *relationships* between them. On these premises, we propose Graph-of-Mark (GoM), the first pixel-level visual prompting technique that overlays scene graphs onto the input image for spatial reasoning tasks. We evaluate GoM across 3 open-source MLMs and 4 different datasets, conducting extensive ablations on drawn components and investigating the impact of auxiliary graph descriptions in the text prompt. Our results demonstrate that GoM consistently improves the zero-shot capability of MLMs in interpreting object positions and relative directions, improving base accuracy in visual question answering and localization up to 11 percentage points.

GoM — <https://github.com/disi-unibo-nlp/graph-of-marks>

## 1 Introduction

It takes more than detecting objects to make sense of a visual scene. Often, it is the *arrangement* of objects in space that reveals the underlying semantics of complex 2D and 3D environments. Despite the proliferation and remarkable progress of multimodal language models (MLMs) (Yin et al. 2024), spatial reasoning remains a challenge in machine perception (Kamath et al. 2023). Recent empirical studies show that even state-of-the-art models tend to overlook positional aspects and view images as mere “bags of objects” (Herzig et al. 2023; Doveh et al. 2023). This limitation is deeply rooted in their architecture and training objectives, which prioritize global image-level understanding over explicit spatial supervision with localized representations and compositionality (Zhang et al. 2025). As a result, MLMs exhibit persistent difficulty in distinguishing even basic spatial concepts (Fu et al. 2024; Majumdar et al. 2024; Zhang et al. 2024b; Shiri et al. 2024) such as *left* and *right*, *above* and

\*These authors contributed equally.

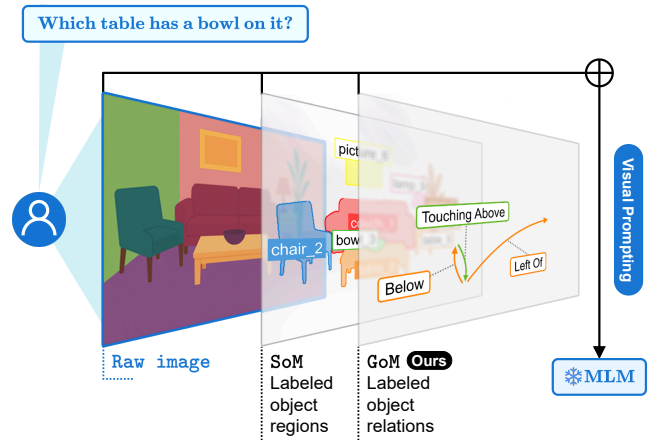


Figure 1: **Illustration of GoM.** A multimodal language model is prompted by anchoring the input image in scene graphs expressing spatial object relations that are relevant to solving the task query provided by the user.

*below*, *near*, and *behind*—the latter requiring depth cues not captured by standard RGB pixel data. Yet, spatial reasoning is far from an academic abstraction—it is fundamental to real-world applications, spanning fields from biomedical research and clinical practice (Domeniconi et al. 2014b; di Lena et al. 2015) to GUI agents, augmented reality, robotic manipulation, autonomous navigation and neuro-symbolic systems (Delvecchio, Molfetta, and Moro 2025).

A natural response to these limitations has been to inject spatial awareness into MLMs by fine-tuning (Chen et al. 2024a; Zhu et al. 2025a). However, this route is computationally expensive and inflexible as the models must be re-trained to handle new tasks and domains, a challenge often observed in cross-domain generalization (Domeniconi et al. 2014a; Moro et al. 2018). An increasingly attractive alternative has emerged in the form of prompting, where task-specific guidance is provided at inference time through carefully crafted inputs. Since textual prompts inherently struggle to convey fine-grained or dense information from image modalities (Lin et al. 2024), *visual prompting* offers a more effective mechanism to uncover the spatial reasoning potential hidden within MLMs. Existing methods can be divided

into two classes: *embedding-level (soft)*, which encodes visual items into latent features, and *pixel-level (hard)*, which renders visible marks—such as masks, boxes, arrows, numbers, or alphabets—onto the image. Among the latter, Set-of-Mark (SoM) (Yang et al. 2023) has rapidly established itself as a de facto standard for improving visual grounding by overlaying numbered regions. Its effectiveness has inspired a growing line of follow-up work, including OmniParser (Lu et al. 2024). Despite these advances, marked objects are predominantly treated as independent entities—overlooking the relational structure that governs scenes and leaving a rich modeling layer still untapped.

In this paper, we introduce Graph-of-Mark (GoM), the first training-free, pixel-level visual prompting technique that embeds graph-connected marks for zero-shot spatial inference. Before feeding the model, GoM draws a *scene graph* (SG) over the raw input image—nodes correspond to detected object instances, while edges denote their spatial relations (Figure 1). Designed as a lightweight *plug-and-play* module compatible with any MLMs, GoM constructs SGs automatically and at scale using open-vocabulary object detectors, segmenters, and depth estimators. Moreover, GoM is broadly applicable, as the underlying SGs demand no ground-truth annotations and are both task- and domain-agnostic.

We validate the effectiveness of GoM through rigorous quantitative evaluation using 3 open-source MLMs and 4 heterogeneous, publicly available datasets spanning 2 vision tasks: visual question answering (VQA) and referring expression comprehension (REC). Our results demonstrate that GoM consistently enhances the zero-shot capabilities of all tested models, surpassing the performance of the existing most popular image prompting techniques. We further conduct ablation studies to dissect the contribution of edge labels beyond simple connectivity and the relative effectiveness of numeric versus alphabetic labels. Finally, we examine the impact of auxiliary SG descriptions included in the text prompt, probing whether MLMs can internalize the relational structure from visual input alone. To support reproducibility and foster further innovation, we release GoM-code, preprocessed dataset images, and evaluation scripts—under an open-source MIT license. We hope that GoM will provide fertile ground for the development of spatial-aware MLMs and hybrid graph-language solutions.

## 2 Related Work

### Multimodal Language Models for Spatial Reasoning.

The advancement of MLMs has enabled unified reasoning over text and vision inputs within scalable general-purpose architectures (Chen et al. 2024b). Still, frontier models fall short in tasks involving visual evidence grounding (Xiao et al. 2024) and spatial intelligence (Yang et al. 2025b; Ramakrishnan et al. 2025). Models such as CLIP (Radford et al. 2021), SPHINX (Lin et al. 2023), LLaVA (Liu et al. 2023), VisionLLM (Wang et al. 2023a; Wu et al. 2024a), and CogVLM2 (Wang et al. 2024b), integrate visual and textual encoded representations using instruction tuning or contrastive pretraining to align modalities. Spatial information in these systems is typically encoded as plain text coordinates (Peng et al. 2023) or serialized region descrip-

tions (Chen et al. 2023a), placing the burden of geometric interpretation on the language backbone. Spatial reasoning is treated more directly in models that encode structural information at the region or scene level. SpatialVLM (Chen et al. 2024a) uses the supervision of synthetic 3D question-answer pairs to model object configurations, whereas Struct2D (Zhu et al. 2025a) applies structured 2D multimodal input to guide 3D perception. Shikra (Chen et al. 2023b) embeds coordinate references as positional tokens within the language stream, preserving compatibility with frozen decoders. ASM (Wang et al. 2024a) leverages densely annotated region-level supervision to improve spatial grounding, while SPHINX-V (Lin et al. 2025) directly encodes visual tags into the input image, enabling localized understanding with frozen LLMs. GLaMM (Rasheed et al. 2024) generates natural language responses while simultaneously providing the associated object segmentation masks, and Omni-RGPT (Heo et al. 2025) introduces Token Mark to achieve consistent region representation for both images and videos by embedding tokens directly into visual feature space and text prompts. VIS-PROG (Gupta and Kembhavi 2023) builds executable neuro-symbolic visual programs to enable spatial reasoning across different tasks. However, these methods require expensive re-training and custom architecture, limiting their adoption in general-purpose domains. GoM can be applied to existing models without modification, enabling spatial reasoning with minimal overhead.

**Image Prompting and Augmentation.** Visual prompting has become a standard lightweight mechanism for reducing hallucinations in MLMs (Hu et al. 2024) and steering vision-language models via spatial annotations (Li et al. 2024a; Cai et al. 2024), enabling task adaptation and spatial awareness without altering model parameters. Early methods enhanced captioning and VQA alignment using region segmentation and bottom-up attention (Anderson et al. 2018), or soft perturbations to improve grounding. Prompting strategies have since incorporated more structured spatial signals. Pixel-level modifications driven by text (Yu and Wang 2024) and training-free visual prompt learning through test-time optimization (Wu et al. 2024b) guide models toward relevant regions under varying prompt formulations. Region-based architectures (Guo et al. 2024; Zhang et al. 2024c; You et al. 2024; Zhang et al. 2024a) incorporate bounding-box descriptors for more accurate spatial alignment, while segmentation-based methods inject structured priors such as SAM-derived masks (Yuan et al. 2024; Kirillov et al. 2023) or exploit region-level contrastive prompts (Wan et al. 2024) to sharpen spatial attention. More recent formulations re-frame prompting as symbolic supervision over discrete visual elements. CPT (Yao et al. 2024) encodes color-coded markers as fill-in-the-blank targets. Set-of-Marks (Yang et al. 2023) overlays alphanumeric tags on segmented regions to support co-reference and alignment across modalities, while Magma (Yang et al. 2025a) enhances robot manipulation and UI navigation by introducing Trace-of-Mark to extract the evolving position of overlaid marks. SpatialRGPT (Cheng et al. 2024) learns regional representations from 3D scene graphs and integrates depth information into the visual en-

coder to enhance spatial reasoning, while ROI-aware inputs support medical VQA (Chen et al. 2024c; Zhu et al. 2025b). OmniParser (Lu et al. 2024) treats UI images as structured visual documents, decomposing them into typed layout elements. Although these methods achieve good results at lower costs, they do not explicitly leverage logical or positional spatial relationships, instead relying on existing models to interpret the augmented visual scene.

**Structured Graph Scene Injection.** Spatial layouts in vision-language tasks have been leveraged to construct scene-level graphs that encode object relations and guide attention toward regions relevant for reasoning (Yang et al. 2024a). Region-level graph structures have been combined with GCNs for image captioning (Yao et al. 2018) and VQA (Yusuf et al. 2025), and with random-walk mechanisms (Wang et al. 2023b) to support relational queries. Question-conditioned attention over scene graphs (Li et al. 2019) and contrastive learning for unsupervised graph construction (Souza et al. 2023) further enhance graph-based VQA representations. Scene graphs have also been integrated into multimodal pipelines that address referring expression resolution (Wu et al. 2023), visual common-sense reasoning (Wang et al. 2022), and object-level spatial attention in transformers (Kant et al. 2020). Several approaches enrich these graphs with external knowledge. KRISP (Marino et al. 2021) combines symbolic and LLM-derived features, while methods like MAIL (Dong et al. 2024) generate concept graphs through multi-stage prompting. The explicit incorporation of external graph knowledge into neural models via relational attention are explored in (Zhang and Zhao 2021). Recent methods integrate graph construction with structured prompting. ConceptGraphs (Gu et al. 2024) generates 3D scene graphs with the support of an LLM and converts them into structured textual descriptions to support spatial planning, while Compositional CoT (Mitra et al. 2024) employs synthetic textual scene graphs to prompt LLMs for enhanced multi-step reasoning and image interpretation. Despite these advancements, existing methods either fuse scene graphs through latent mechanisms or convert them into text, hindering direct visual grounding of structured information.  $\text{GoM}$ , on the other hand, is the first to embed the graph directly into the image, enabling multimodal models to leverage relational cues at the pixel level without relying on implicit or textual intermediaries that obscure structure and reduce interpretability.

### 3 Method

In this section, we describe  $\text{GoM}$  prompting, designed to shift MLM perception from object collections to object networks. A compact algorithmic overview of  $\text{GoM}$ , expressed in pseudocode, is made available in the Supplementary Material.

#### Problem Definition

We consider MLMs capable of reasoning over both vision and language input modalities. Given an RGB image observation  $I \in \mathbb{R}^{H \times W \times 3}$  and an associated text prompt  $T$  comprising a task prefix and instance-specific query, these models generate a distribution  $P_{\text{MLM}}(\cdot | I, T)$  over textual comple-

tions that serve as responses. Although the exact parameterized MLM submodules (vision encoder, language encoder, and decoder) differ between models, the overarching computational flow remains consistent. Recent studies have demonstrated that visual tokens fed into MLMs largely influence the attention map values, thus indirectly controlling the model output (Wu et al. 2024b). Building on this premise, we propose a training-free method that strategically alters visual tokens to induce a global perspective encompassing not only individual objects—the primary units for visual reasoning—but also their interactions. To this end, we augment the input image  $I$  to produce a scene graph-annotated image  $I^{\text{SG}}$  that explicitly encodes spatial relationships between objects.

$$P_{\text{MLM}}(\cdot | \underbrace{\text{GoM}(I), T}_{I^{\text{SG}}})$$

#### Graph-Based Image Augmentation

**Object Detection and Segmentation.** Identifying the objects present in  $I$ , their classes, and positions constitutes the initial stage of the SG building. We adopt a coarse-to-fine approach to partition  $I$  into a set of regions  $R = \{r_1, \dots, r_{|R|}\}$ , where  $r_i$  is the spatial extent of the object  $i$ .

**Object detection** We first determine the 2D bounding box coordinates and the class label of each object. Unlike previous works that rely on single detection models, we harness an *ensemble* of independent detectors with complementary strengths to maximize object coverage in cluttered scenes. We keep only predictions whose confidence scores meet or exceed the threshold  $\tau_{\text{OD-min-conf}}$ . The weighted boxes fusion (WBF) heuristic (Solovyev et al. 2021) is used to merge multiple bounding boxes of the *same* class into a single, confidence-weighted average box; boxes are considered overlapping if their intersection over union (IoU) exceeds a specified threshold  $\tau_{\text{overlap-IoU}}$ .

**Object segmentation** We subsequently refine object representations from rectangular bounding boxes to precise region masks with contour outlines. When the segmenter fails, we fall back to using the box regions.

**Relation Estimation.** Having obtained object regions, we compute pairwise spatial relations between them—the  $\text{GoM}$  heart. We create an ontology organizing 7 relation types into 3 groups: (i) directional (above, below, left\_of, right\_of), (ii) depth stacking (in\_front\_of, behind), and (iii) general proximity (near) for cases where directional classification is ambiguous. To extend expressivity, (i)-type relations can be augmented with modifiers (touching, very\_close, close).

**Directional** For each object pair, we take the bounding box centers  $(cx_i, cy_i)$  and  $(cx_j, cy_j)$ , and calculate the displacement vector  $(dx, dy) = (cx_j - cx_i, cy_j - cy_i)$ . Relations are determined by the dominant displacement direction: if  $|dy| \geq |dx|$  and  $|dy| > \tau_{\text{dir-margin}}$ , the relation is above (if  $dy < 0$ ) or below (if  $dy > 0$ ). If  $|dx| > \tau_{\text{dir-margin}}$ , the relation is left\_of or right\_of based on the sign of  $dx$ .

**Depth stacking** A monocular model recovers metric depth from  $I$ . For each object pair, we sample the depth value at the bounding box centers, normalized to  $[0, 1]$  where higher

values indicate objects nearer to the camera. Depth relationships are established for all object pairs, even non-overlapped ones, with substantial z-difference: when  $|\text{depth}_j - \text{depth}_i| > \tau_{z\text{-diff}}$ , if  $\text{depth}_j > \text{depth}_i$ , object  $j$  is `in_front_of` object  $i$ , otherwise `behind`.

**Proximity** If the center-based distance between two objects is below  $\tau_{\text{near}}$  and they have not been assigned to a directional case, the `near` relation acts as fallback.

**Modifiers** We detail directional relations between objects of *different* classes based on their closeness. Objects are classified as `touching` when they exhibit overlap ( $\text{IoU} > \tau_{\text{touch-IoU}}$ ) or minimal separation between box edges ( $\leq \tau_{\text{touch-gap}}$ ). For non-touching objects, we employ normalized distance  $d_{\text{norm}}$ —the Euclidean distance between centers scaled by  $I$  dimensions—to establish proximity gradations: `very_close` ( $d_{\text{norm}} < \tau_{\text{v-close}}$ ), `close` ( $d_{\text{norm}} < \tau_{\text{close}}$ ).

**Filtering** The annotations computed so far are long-tailed and might incorporate information unrelated to the given task prompt  $T$ . We thus deploy a two-step filtering pipeline.

**Objects** We retain only objects relevant to the query—those explicitly or implicitly mentioned through their labels, aliases, or synonyms. For efficiency, the mention detection pipeline is incremental, progressing from lexical matching to semantic embedding comparison with cosine similarity  $> \tau_{\text{query-obj}}$ . If multiple objects are relevant to the query, we retain only them and their interrelations. If a single object is relevant, we keep that object, the relations for which it is the head, and the tail objects with which it interacts. If no object is relevant, we retain all objects and relations to ensure comprehensive coverage when query specificity is insufficient.

**Relations** For each object, we retain only the top- $k$  relations having it as head. Per-object ranking (the lower the better) employs a two-tier sorting, with query relevancy as the primary key (0 for relevant, 1 otherwise) and spatial distance as the secondary key. In this way, for instance, depth relations between distant objects are naturally excluded in favor of nearby interactions. A relation is judged relevant if its label matches at least one query relation term through the same matching pipeline introduced for the object case. Query relation terms include mentioned paraphrases and synonyms of target relations. To avoid redundancy, we iterate through the relation list (same order as presented) and retain only the first occurrence linking any two objects, thus preventing the modeling of both direct and inverse relations.

**Scene Graph Rendering.** Having defined meaningful objects (nodes) and spatial relations (edges), we possess all the ingredients to construct SGs overlaid on  $I$ .

**Node mark type** Object regions are rendered as mask marks with class-specific coloring (solid border, semi-transparent fill), ensuring objects of the same class share identical colors. To enable MLM reference to specific objects in textual outputs, we render unique ID marks adjacent to each object with interpretable and speakable content. We experiment with both numeric IDs and their textual extensions (formatted as `<class>_<id>`). ID marks appear within rectangular boxes sharing the border color of their corresponding object, with box fill and font color dynamically optimized for background contrast.

**Edge mark type** Relations are visualized as directed head-to-tail arrows, colored to match the head object. In addition to connectivity, we investigate the impact of relation labels rendered as textual marks within boxes having the same border color as the arrow, white fill, and black font.

**Mark allocation** Effective mark (mask, ID/label, arrow) placement requires collision-free positioning to prevent MLM confusion. The anti-overlapping strategy from SoM works exclusively to object IDs constrained within object regions, limiting its applicability to marks with minimal spatial footprint. Since our approach incorporates textual object IDs and edge marks extending beyond single object regions, we propose a novel allocation algorithm. First, we draw edge labels at object midpoints, as these marks are more intrusive visually. Object ID marks are subsequently positioned at region centroids, remaining within boundaries when spatial constraints permit; otherwise, they are relocated externally while avoiding conflicts with existing ID and label marks. Arrows are created with slightly shortened endpoints to accommodate textual labels. When multiple arrows originate from the same object in similar directions, we increase their curvature radius progressively to prevent overlap. Following initial placement, we use a resolver to iteratively displace conflicting marks along coordinate axes with small incremental steps until all collisions are eliminated. To preserve visual coherence between displaced edge labels and their corresponding arrows, we render dashed lines connecting each label to its original midpoint position.

## Prompting

We inspect two prompting modes, depending on the modalities involved in communicating spatial information.

**Visual SG** We populate  $T$  with plain task instructions. The MLM must interpret SG meaning solely through  $I^{\text{SG}}$ .

**Visual and Textual SG** We populate  $T$  with both task instructions and a verbalized SG representation ( $T^{\text{SG}}$ ), complementing  $I^{\text{SG}}$ . We use triplet arrow notation<sup>1</sup> with interpretative guidance to facilitate MLM understanding. The prompt template is reported in the Supplementary Material. Following SoM, we emphasize that GoM supports tasks requiring visual outputs by enabling reverse mapping from spoken mark IDs back to their corresponding object regions.

## 4 Experimental Setup

### Datasets and Metrics

We evaluate GoM’s performance across a suite of VQA and REC tasks reflecting complementary dimensions of spatial visual understanding. For VQA, we include GQA (Hudson and Manning 2019), VQAv1 (Antol et al. 2015), and VQAv2 (Goyal et al. 2017). All three feature open-ended questions and are evaluated using accuracy-based metrics following the official evaluation protocols defined by their respective creators. GQA demands multi-step spatial reasoning and fine-grained object relationships. VQAv1 covers general queries that combine vision, language, and common-sense reasoning. VQAv2 mitigates language bias by pairing

<sup>1</sup>`<head>-(<relation>)-><tail>`.

Question: “Is the potted plant below the oven?”

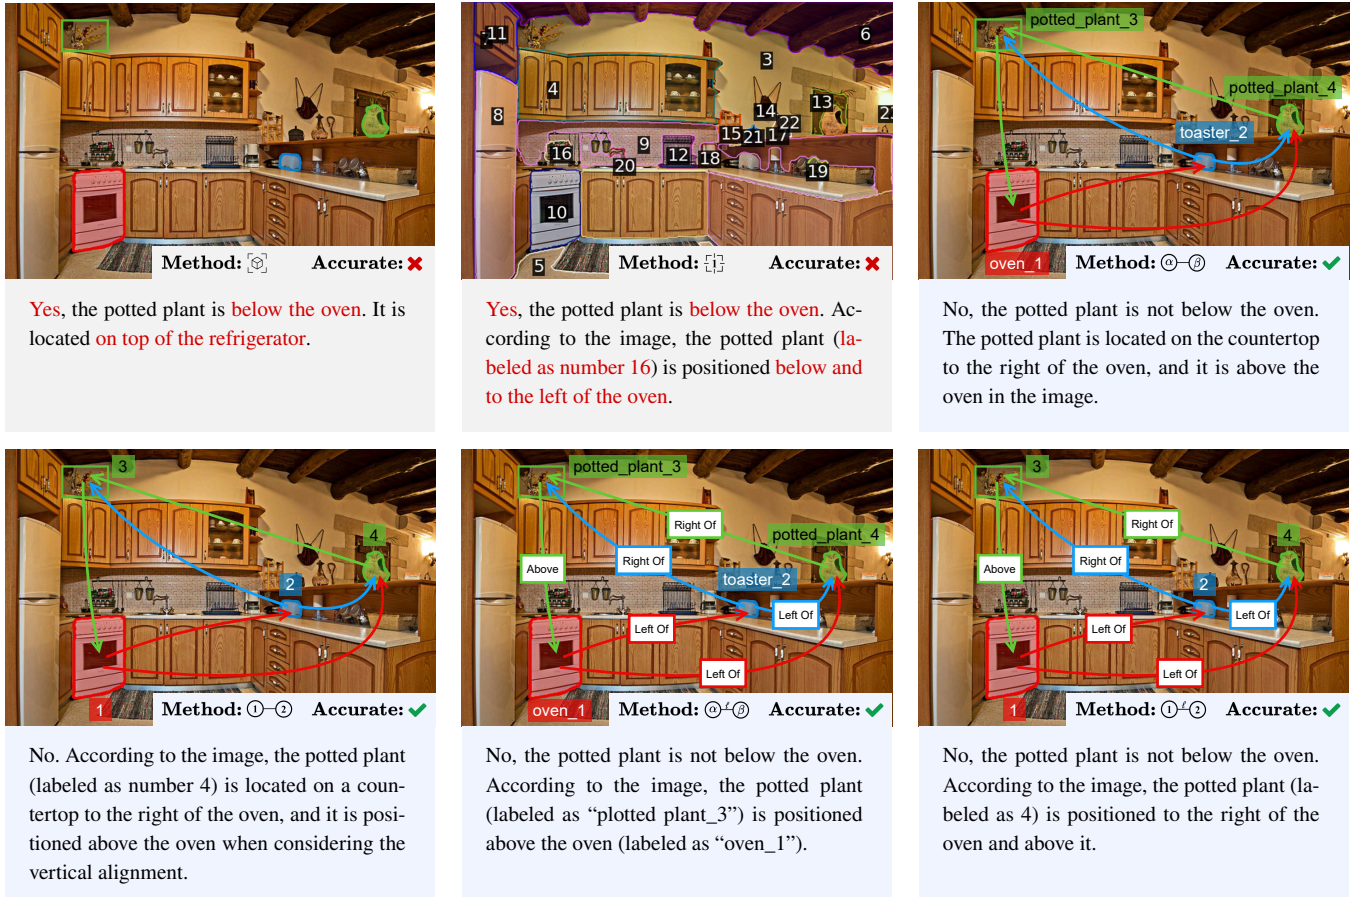


Figure 2: **Qualitative example illustrating the impact of image preprocessing on VQA performance.** The same question from VQA-v2 is posed to Qwen2.5-7B using 6 different hard visual prompts, highlighting how pixel transformations can influence the model’s responses. For figure readability, the font size and line thickness have been increased compared to their actual values. Gray boxes denote baseline outputs, while blue boxes indicate those from our proposed GOM. See icon legend in Table 2.

visually similar images with different answers to identical questions. For REC, we analyze RefCOCOg (Kazemzadeh et al. 2014), which involves localizing a specific object instance based on a textual description, often including spatial details to disambiguate the referent. We approach the task by prompting the MLM to generate the ID of the target object; a prediction is considered correct if the region associated with the predicted object has an  $\text{IoU} \geq 0.9$  with the ground truth bounding box. We underline that this evaluation setup is not applicable when the model receives only  $I$  or its segmentation-augmented version. In line with established practices (Yang et al. 2023), we randomly sample 1K images from each dataset to manage computational and time constraints. In contrast to prior work that filters specific queries, we retain all questions associated with each selected image, ensuring robust and statistically meaningful evaluation. On average, each image in GQA, VQAv1, VQAv2, and RefCOCOg is paired with 3, 4, 4, and 1 queries, respectively.

## Implementation Details

**Object Detectors.** Modern object detection has evolved to support both closed-vocabulary (CV) and open-vocabulary (OV) scenarios, with models increasingly expected to handle unknown classes and generalize across domains (Li et al. 2025). While OV detectors offer strong flexibility through text-conditioned recognition and self-supervised training, they often trade off precision or require specialized prompt design. To maximize object recall in our pipeline without additional training, we combine the outputs of three complementary, training-free detectors. These include OWL-V2 (Minderer et al. 2023), an OV model capable of detecting arbitrary categories via textual prompts; YOLOv8-X (Ultralytics 2023), a fast and high-confidence CV detector for common classes; and Mask R-CNN R101-FPN from Detectron2 (Wu et al. 2019), CV, which supplements robust region proposals, especially for persons.

Phase	Hyperparam Setting
	$\tau_{\text{OD-min-conf}} = \{0.4, 0.5 *, 0.8\}$ $\tau_{\text{overlap-IoU}} = 0.9$
<b>1/2/3</b>	$\tau_{\text{dir-margin}}^{\ddagger} = \{10, 20 *, 50\}$ , $\tau_{z\text{-diff}} = \{0.1 *, 0.15, 0.20\}$ , $\tau_{\text{near}}^{\ddagger} = 5,000$
	$\tau_{\text{touch-IoU}}^{\ddagger} = 0.1$ , $\tau_{\text{touch-gap}}^{\ddagger} = 3$ , $\tau_{\text{v-close}}^{\ddagger} = 0.05$ , $\tau_{\text{close}}^{\ddagger} = \{0.12 *, 0.15\}$
	$\tau_{\text{query-obj}} = \{0.5 *, 0.7, 0.8\}$ , $k = \{1, 2, 3 *, 4, 5, 6\}$ seed = {42, 123, 456},
	temp = {0.1, 0.3, 0.5}, top-p = {0.7, 0.9, 0.95}

<sup>‡</sup> Normalized pixel-level distance. **Symbols:** = Object detection; **1/2/3** = Directional/Depth/Proximity; = Modifiers; = Object/Relation filtering; = Decoding.

Table 1: **Hyperparameter sweep.** Top: algorithm. Bottom: decoding (for best algorithm config). \* = algorithm values, picked after preliminary runs with Qwen-2.5-VL on GQA.

**Segmenter.** Transformer-based segmentation models have become the standard for dense prediction tasks, offering strong performance through unified, query-driven architectures (Li et al. 2024b). We adopt Segment Anything SAM-HQ (Ke et al. 2023), OV, which produces high-resolution instance masks without retraining. Its promptable design and use-case robustness make it suitable for isolating fine object boundaries from unconstrained image sources.

**Depth Estimator.** Monocular depth estimation has seen rapid progress, with recent models targeting metric depth recovery in scale- and shift-invariant settings (Yang et al. 2024b; Wang et al. 2025). However, these approaches often rely on accurate focal length and calibration parameters, which are hard to obtain for unconstrained, in-the-wild images and add significant preprocessing overhead. For this reason, we employ MiDaS DPT-Large (Ranftl et al. 2022), which estimates relative depth without requiring camera metadata. Trained on a broad range of indoor and outdoor scenes, MiDaS provides reliable predictions in open-world settings while also offering low memory usage and latency.

**Filtering.** Aliases and synonyms for object labels are derived through WordNet (v3.0) synsets. The semantic recognition of object mentions within the query draws on pre-trained FastText embeddings (cc.en.300.vec).

**MLMs.** We test Qwen-2.5-VL-7B (Instruct) (Bai et al. 2025), Gemma-3-4B (Instruct) (Kamath et al. 2025), and LlamaV-o1-11B (Thawakar et al. 2025). These open-source models were selected to ensure empirical diversity between architectures, sizes, and usage scenarios, with a focus on common low- and mid-resource settings. Our aim is to verify whether GoM improves downstream performance without relying on instruction-heavy or domain-specialized models. LlamaV-o1, a reasoning model, allows us to inquire whether GoM adds value to step-by-step thoughts.

**Hyperparameters.** Table 1 lists the hyperparameters along with their empirical search grid. After selecting the optimal algorithmic configuration, we run MLM inference using nucleus sampling and capping generation at 512 tokens. Each run (<MLM, prompt method, dataset>) is repeated using 27 combinations of seed, temperature, and top-p.

	Method	GQA	VQA		RefCOCOg
			v1	v2	
Gemma-3-4B		56.2 $\pm$ 0.11	64.3 $\pm$ 0.14	59.9 $\pm$ 0.19	–
		53.8 $\pm$ 0.34	63.7 $\pm$ 1.15	60.2 $\pm$ 1.37	–
		56.9 $\pm$ 0.34	63.8 $\pm$ 0.07	59.0 $\pm$ 0.18	54.8 $\pm$ 0.91
		58.8 $\pm$ 0.16	65.2 $\pm$ 0.38	62.8 $\pm$ 0.36	56.3 $\pm$ 0.91
		60.3 $\pm$ 0.45	71.5 $\pm$ 0.1	70.2 $\pm$ 0.25	<b>56.4</b> $\pm$ 0.66
Qwen-2.5-VL-7B		61.6 $\pm$ 0.11	77.7 $\pm$ 0.11	73.8 $\pm$ 0.16	–
		53.2 $\pm$ 0.30	65.2 $\pm$ 1.21	67.2 $\pm$ 1.20	–
		61.8 $\pm$ 0.35	65.4 $\pm$ 0.3	68.6 $\pm$ 0.16	55.5 $\pm$ 0.39
		62.5 $\pm$ 0.32	77.9 $\pm$ 0.42	74.0 $\pm$ 0.06	56.5 $\pm$ 0.90
		61.9 $\pm$ 0.41	78.1 $\pm$ 0.37	76.4 $\pm$ 0.21	56.8 $\pm$ 0.74
LlamaV-o1-11B		60.2 $\pm$ 0.20	75.3 $\pm$ 0.11	75.1 $\pm$ 0.12	–
		61.9 $\pm$ 0.35	72.6 $\pm$ 1.22	75.3 $\pm$ 1.22	–
		62.0 $\pm$ 0.33	72.8 $\pm$ 0.42	75.5 $\pm$ 0.38	55.3 $\pm$ 1.03
		62.4 $\pm$ 0.3	76.0 $\pm$ 0.49	75.9 $\pm$ 0.23	57.2 $\pm$ 0.79
		62.5 $\pm$ 0.3	77.9 $\pm$ 0.42	80.9 $\pm$ 0.33	57.3 $\pm$ 0.71
GOM (Ours)		<b>67.0</b> $\pm$ 0.44	79.8 $\pm$ 0.31	83.4 $\pm$ 0.02	56.9 $\pm$ 0.94
		65.0 $\pm$ 0.37	<b>83.2</b> $\pm$ 0.08	<b>83.6</b> $\pm$ 0.04	<b>57.6</b> $\pm$ 0.26

**Prompt Legend**

*Baselines* Raw image, Segmented objects, Segmented objects + Object Num IDs (SoM)

*GOM (Ours)* Segmented objects + Relations + Object Text IDs / Object Num IDs / Object Text IDs + Relation labels / Object Num IDs + Relation labels

Table 2: **Accuracy results** ( $\uparrow$ ). Comparison between GoM variants (Visual SG only) and baseline prompt strategies on 4 datasets and 3 MLMs. We report mean  $\pm$  standard deviation across decoding runs. Best results are highlighted.

**Hardware Setup.** All runs were conducted on a workstation running Ubuntu 20.04.3 LTS, equipped with a single NVIDIA GeForce RTX3090 GPU (24GB VRAM), 64GB of RAM, and an Intel® Core™ i9-10900X CPU @ 3.70GHz.

## 5 Results

We evaluate GoM along three axes: (1) impact on spatial reasoning performance, (2) benefits of multimodal SG integration, and (3) efficiency in  $I^{\text{SG}}$  computation. Fair comparisons are made against vanilla MLMs and competitive hard visual prompting techniques, i.e., segmentation-only and SoM.




Figure 3: **Effect of graph density.** Performance of  $G_{\text{OM}}$  in VQA<sub>v2</sub> as a function of the number of edges in the visual scene graph. 0 edges corresponds to SoM-like prompting.

**Graph Guidance in Spatial Reasoning.** Main results are shown in Table 2.  $G_{\text{OM}}$  demonstrates superior performance across all experimental conditions, substantiating the efficacy of explicit relational encoding for spatial reasoning and scene comprehension. Zooming out, we elucidate that even lightweight open-source MLMs with  $\leq 11\text{B}$  parameters can adeptly exploit hard visual prompting, contravening precedent publications that documented success solely within commercial closed-source models such as GPT-4V (Yang et al. 2023). Gemma-3 experiences the most pronounced improvement, whereas Qwen manifests adverse sensitivity to SoM, frequently failing to reference the correct object regions. LlamaV attains the highest absolute scores—83.6 in VQA and 57.6 in REC—suggesting that reasoning models are particularly adept at leveraging  $G_{\text{OM}}$  representations. Notably, despite the absence of a universally optimal  $G_{\text{OM}}$  variant, the basic setting featuring textual object IDs without relation labels generally achieves top performance in VQA, while numeric IDs are desirable for REC. Targeted error analysis reveals that models track textual descriptors more faithfully when generating abstractive answers. When relational labels are omitted, models systematically ignore directional indicators. Figure 2 presents qualitative input-output augmentation examples. Only graph-enhanced images enable accurate spatial interpretation; SoM conversely degrades reasoning through erroneous regional attributions. Maximal  $G_{\text{OM}}$  effectiveness materializes with 3-10 entities and 4-16 relations (Figure 3). Beyond this range, surplus annotations introduce noise, diminishing the headroom.

**Scene Graph Modality.** Figure 4 systematically evaluates the contribution of SGs across distinct prompt modalities. Visual graph representations repeatedly lead to higher accuracy compared to textual encodings exclusively examined in prior work, with elevations of up to +10%. Remarkably, the combination of verbalized SGs with their visual counterparts always produces performance gains, albeit of modest magnitude. This empirical evidence corroborates that  $G_{\text{OM}}$  enhancements stem not from the superficial presence of graph structures—as postulated by other research groups—but rather from a deliberate visual graph prompting pipeline that better activates the latent reasoning faculties of MLMs.




Figure 4: **Accuracy impact deriving from augmenting the visual ( $I$ ) and textual ( $T$ ) prompt with scene graphs (SG apex).** Gemma-3 results. Proposed  $G_{\text{OM}}$  solutions have  $I^{\text{SG}}$ .

**Efficiency.** The performance-efficiency tradeoff is a critical consideration habitually neglected in evaluations of visual prompting methodologies. Comparative latency analysis on our computational infrastructure reveals that  $G_{\text{OM}}$  incurs very low overhead, averaging 1.13 seconds per image compared to 0.77 and 0.92 seconds for segmentation-only and SoM approaches, respectively, attributable to relation estimation. On the other hand, this cost is offset by advances in spatial reasoning, primarily in VQA tasks.

## 6 Conclusion

We introduced Graph-of-Mark, the first visual prompting method that embeds depth-aware scene graphs into the input image to support spatial reasoning. Unlike prior techniques that treat objects as isolated units,  $G_{\text{OM}}$  exposes multimodal models to relational data without requiring retraining or architectural changes. Through extensive evaluation, we show that  $G_{\text{OM}}$  consistently enhances performance in spatially grounded tasks and offers a powerful, robust mechanism for structured visual understanding in lightweight open-source MLMs. Our method outperforms competitive approaches while demonstrating the advantages of visual design over graph-verbalized text prompting. Future directions include scene hypergraphs for complex scenes, stereo vision for improved depth reasoning, and temporal modeling for video understanding. Furthermore, we project significant potential in the clinical domain, where  $G_{\text{OM}}$  could be specialized and optimized for near-real-time augmentation. This would enable medical MLMs to achieve superior test-time performance on tasks such as diagnostic classification and surgical video analysis. As MLMs’ capabilities continue to grow, we expect the benefits of  $G_{\text{OM}}$  to amplify accordingly.

## Acknowledgements

Research partially supported by: AI-PACT (CUP B47H22004450008, B47H22004460001); National Plan PNC-I.1 DARE (PNC0000002, CUP B53C22006450001); PNRR Extended Partnership FAIR (PE00000013, Spoke 8); 2024 Scientific Research and High Technology Program, “AI analysis for risk assessment of empty lymph nodes in endometrial cancer surgery,” Fondazione Cassa di Risparmio in Bologna; Chips JU TRISTAN (G.A. 101095947). LG Solution Srl for co-funding L. Molfetta’s PhD scholarship.

## References

- Anderson, P.; et al. 2018. Bottom-Up and Top-Down Attention for Image Captioning and Visual Question Answering. In *CVPR*. CVF/IEEE Computer Society.
- Antol, S.; et al. 2015. VQA: Visual Question Answering. In *ICCV*. IEEE Computer Society.
- Bai, S.; et al. 2025. Qwen2.5-VL Technical Report. *CoRR*, abs/2502.13923.
- Cai, M.; et al. 2024. ViP-LLaVA: Making Large Multimodal Models Understand Arbitrary Visual Prompts. In *CVPR*. IEEE.
- Chen, B.; et al. 2024a. SpatialVLM: Endowing Vision-Language Models with Spatial Reasoning Capabilities. In *CVPR*. IEEE.
- Chen, J.; et al. 2023a. MiniGPT-v2: large language model as a unified interface for vision-language multi-task learning. *CoRR*, abs/2310.09478.
- Chen, K.; et al. 2023b. Shikra: Unleashing Multimodal LLM's Referential Dialogue Magic. *CoRR*, abs/2306.15195.
- Chen, L.; et al. 2024b. Are We on the Right Way for Evaluating Large Vision-Language Models? In *NeurIPS*.
- Chen, X.; et al. 2024c. R-LLaVA: Improving Med-VQA Understanding through Visual Region of Interest. *CoRR*, abs/2410.20327.
- Cheng, A.; et al. 2024. SpatialRGPT: Grounded Spatial Reasoning in Vision-Language Models. In *NeurIPS*.
- Delvecchio, G. P.; Molfetta, L.; and Moro, G. 2025. Neuro-Symbolic Artificial Intelligence: A Task-Directed Survey in the Black-Box Models Era. In *IJCAI*. ijcai.org.
- di Lena, P.; et al. 2015. GOTA: GO term annotation of biomedical literature. *BMC Bioinform.*, 16.
- Domeniconi, G.; et al. 2014a. Cross-domain Text Classification through Iterative Refining of Target Categories Representations. In *KDIR*. SciTePress.
- Domeniconi, G.; et al. 2014b. Discovering New Gene Functionalities from Random Perturbations of Known Gene Ontological Annotations. In *KDIR*. SciTePress.
- Dong, J.; et al. 2024. Modality-Aware Integration with Large Language Models for Knowledge-Based Visual Question Answering. In *ACL*. ACL.
- Doveh, S.; et al. 2023. Teaching Structured Vision & Language Concepts to Vision & Language Models. In *CVPR*. IEEE.
- Fu, X.; et al. 2024. BLINK: Multimodal Large Language Models Can See but Not Perceive. In *ECCV*. Springer.
- Goyal, Y.; et al. 2017. Making the V in VQA Matter: Elevating the Role of Image Understanding in Visual Question Answering. In *CVPR*. IEEE Computer Society.
- Gu, Q.; et al. 2024. ConceptGraphs: Open-Vocabulary 3D Scene Graphs for Perception and Planning. In *ICRA*. IEEE.
- Guo, Q.; et al. 2024. RegionGPT: Towards Region Understanding Vision Language Model. In *CVPR*. IEEE.
- Gupta, T.; and Kembhavi, A. 2023. Visual Programming: Compositional visual reasoning without training. In *CVPR2023*. IEEE.
- Heo, M.; et al. 2025. Omni-RGPT: Unifying Image and Video Region-level Understanding via Token Marks. In *CVPR*. CVF/IEEE.
- Herzig, R.; et al. 2023. Incorporating Structured Representations into Pretrained Vision & Language Models Using Scene Graphs. In *EMNLP*. ACL.
- Hu, J.; et al. 2024. Leveraging Hallucinations to Reduce Manual Prompt Dependency in Promptable Segmentation. In *NeurIPS*.
- Hudson, D. A.; and Manning, C. D. 2019. GQA: A New Dataset for Real-World Visual Reasoning and Compositional Question Answering. In *CVPR*. CVF/IEEE.
- Kamath, A.; et al. 2023. What's "up" with vision-language models? Investigating their struggle with spatial reasoning. In *EMNLP*. ACL.
- Kamath, A.; et al. 2025. Gemma 3 Technical Report. *CoRR*, abs/2503.19786.
- Kant, Y.; et al. 2020. Spatially Aware Multimodal Transformers for TextVQA. In *ECCV*. Springer.
- Kazemzadeh, S.; et al. 2014. ReferItGame: Referring to Objects in Photographs of Natural Scenes. In *EMNLP*. ACL.
- Ke, L.; et al. 2023. Segment Anything in High Quality. *arXiv preprint arXiv:2306.01567*. NeurIPS 2023; version 2.
- Kirillov, A.; et al. 2023. Segment Anything. In *ICCV*. IEEE.
- Li, F.; et al. 2024a. Visual in-Context Prompting. In *CVPR*. IEEE.
- Li, L.; et al. 2019. Relation-Aware Graph Attention Network for Visual Question Answering. In *ICCV*. IEEE.
- Li, X.; et al. 2024b. Transformer-Based Visual Segmentation: A Survey. *IEEE TPAMI*.
- Li, Y.; et al. 2025. Open World Object Detection: A Survey. *IEEE TCSVT*.
- Lin, W.; et al. 2025. Draw-and-Understand: Leveraging Visual Prompts to Enable MLLMs to Comprehend What You Want. In *ICLR*.
- Lin, Y.; et al. 2024. Rethinking Visual Prompting for Multimodal Large Language Models with External Knowledge. *CoRR*, abs/2407.04681.
- Lin, Z.; et al. 2023. SPHINX: The Joint Mixing of Weights, Tasks, and Visual Embeddings for Multi-modal Large Language Models. *CoRR*, abs/2311.07575.
- Liu, H.; et al. 2023. Visual Instruction Tuning. In *NeurIPS*.
- Lu, Y.; et al. 2024. OmniParser for Pure Vision Based GUI Agent. *CoRR*, abs/2408.00203.
- Majumdar, A.; et al. 2024. OpenEQA: Embodied Question Answering in the Era of Foundation Models. In *CVPR*. IEEE.
- Marino, K.; et al. 2021. KRISP: Integrating Implicit and Symbolic Knowledge for Open-Domain Knowledge-Based VQA. In *CVPR*. CVF/IEEE.
- Minderer, M.; et al. 2023. Scaling Open-Vocabulary Object Detection. In *NeurIPS*.
- Mitra, C.; et al. 2024. Compositional Chain-of-Thought Prompting for Large Multimodal Models. In *CVPR*. IEEE.

- Moro, G.; et al. 2018. Cross-domain & In-domain Sentiment Analysis with Memory-based Deep Neural Networks. In *IC3K*. SciTePress.
- Peng, Z.; et al. 2023. Kosmos-2: Grounding Multimodal Large Language Models to the World. *CoRR*, abs/2306.14824.
- Radford, A.; et al. 2021. Learning Transferable Visual Models From Natural Language Supervision. In *ICML*. PMLR.
- Ramakrishnan, S. K.; et al. 2025. Does Spatial Cognition Emerge in Frontier Models? In *ICLR*.
- Ranftl, R.; et al. 2022. MiDaS: A High-Accuracy Monocular Depth Estimator. <https://github.com/intel-isl/MiDaS>. Release v3.1 (DPT-Large).
- Rasheed, H. A.; et al. 2024. GLaMM: Pixel Grounding Large Multimodal Model. In *CVPR*. IEEE.
- Shiri, F.; et al. 2024. An Empirical Analysis on Spatial Reasoning Capabilities of Large Multimodal Models. In *EMNLP*. ACL.
- Solovyev, R. A.; et al. 2021. Weighted boxes fusion: Ensembling boxes from different object detection models. *IMAVIS*.
- Souza, B.; et al. 2023. SelfGraphVQA: A Self-Supervised Graph Neural Network for Scene-based Question Answering. In *ICCV*. IEEE.
- Thawakar, O.; et al. 2025. LlamaV-o1: Rethinking Step-by-step Visual Reasoning in LLMs. *CoRR*, abs/2501.06186.
- Ultralytics. 2023. Ultralytics YOLOv8: Next-Generation Real-Time Object Detector. <https://github.com/ultralytics/ultralytics>. Release 2023-01-10.
- Wan, D.; et al. 2024. Contrastive Region Guidance: Improving Grounding in Vision-Language Models Without Training. In *ECCV*. Springer.
- Wang, R.; et al. 2025. MoGe: Unlocking Accurate Monocular Geometry Estimation for Open-Domain Images with Optimal Training Supervision. In *CVPR*. CVF/IEEE.
- Wang, W.; et al. 2023a. VisionLLM: Large Language Model is also an Open-Ended Decoder for Vision-Centric Tasks. In *NeurIPS*.
- Wang, W.; et al. 2024a. The All-Seeing Project: Towards Panoptic Visual Recognition and Understanding of the Open World. In *ICLR*.
- Wang, W.; et al. 2024b. CogVLM: Visual Expert for Pre-trained Language Models. In *NeurIPS*.
- Wang, Y.; et al. 2023b. VQA-GNN: Reasoning with Multimodal Knowledge via Graph Neural Networks for Visual Question Answering. In *ICCV2023*. IEEE.
- Wang, Z.; et al. 2022. SGEITL: Scene Graph Enhanced Image-Text Learning for Visual Commonsense Reasoning. In *AAAI*. AAAI Press.
- Wu, C.; et al. 2023. Scene Graph Enhanced Pseudo-Labeling for Referring Expression Comprehension. In *Findings of EMNLP*. ACL.
- Wu, J.; et al. 2024a. VisionLLM v2: An End-to-End Generalist Multimodal Large Language Model for Hundreds of Vision-Language Tasks. In *NeurIPS*.
- Wu, M.; et al. 2024b. ControlMLLM: Training-Free Visual Prompt Learning for Multimodal Large Language Models. In *NeurIPS*.
- Wu, Y.; et al. 2019. Detectron2. <https://github.com/facebookresearch/detectron2>. Version 0.6.4.
- Xiao, J.; et al. 2024. Can I Trust Your Answer? Visually Grounded Video Question Answering. In *CVPR*. IEEE.
- Yang, C.; et al. 2024a. Improving Vision-and-Language Reasoning via Spatial Relations Modeling. In *WACV*. IEEE.
- Yang, J.; et al. 2023. Set-of-Mark Prompting Unleashes Extraordinary Visual Grounding in GPT-4V. *CoRR*, abs/2310.11441.
- Yang, J.; et al. 2025a. Magma: A Foundation Model for Multimodal AI Agents. In *CVPR*. CVF/IEEE.
- Yang, J.; et al. 2025b. Thinking in Space: How Multimodal Large Language Models See, Remember, and Recall Spaces. In *CVPR*. CVF/IEEE.
- Yang, L.; et al. 2024b. Depth Anything: Unleashing the Power of Large-Scale Unlabeled Data. In *CVPR*. IEEE.
- Yao, T.; et al. 2018. Exploring Visual Relationship for Image Captioning. In *ECCV*. Springer.
- Yao, Y.; et al. 2024. CPT: Colorful Prompt Tuning for pre-trained vision-language models. *AI Open*.
- Yin, S.; et al. 2024. A survey on multimodal large language models. *National Science Review*.
- You, H.; et al. 2024. Ferret: Refer and Ground Anything Anywhere at Any Granularity. In *ICLR*.
- Yu, R.; and Wang, X. 2024. Attention Prompting on Image for Large Vision-Language Models. In *ECCV*. Springer.
- Yuan, Y.; et al. 2024. Osprey: Pixel Understanding with Visual Instruction Tuning. In *CVPR*. IEEE.
- Yusuf, A. A.; et al. 2025. Graph-enhanced visual representations and question-guided dual attention for visual question answering. *Neurocomputing*.
- Zhang, H.; et al. 2024a. Ferret-v2: An Improved Baseline for Referring and Grounding with Large Language Models. *CoRR*, abs/2404.07973.
- Zhang, H.; et al. 2025. A Call for New Recipes to Enhance Spatial Reasoning in MLLMs. *CoRR*, abs/2504.15037.
- Zhang, J.; et al. 2024b. CounterCurate: Enhancing Physical and Semantic Visio-Linguistic Compositional Reasoning via Counterfactual Examples. In *Findings of ACL*. ACL.
- Zhang, S.; et al. 2024c. GPT4RoI: Instruction Tuning Large Language Model on Region-of-Interest. In *ECCV*. Springer.
- Zhang, Y.; and Zhao, Q. 2021. Explicit Knowledge Incorporation for Visual Reasoning. In *CVPR*. CVF/IEEE.
- Zhu, F.; et al. 2025a. Struct2D: A Perception-Guided Framework for Spatial Reasoning in Large Multimodal Models. *CoRR*, abs/2506.04220.
- Zhu, K.; et al. 2025b. Guiding Medical Vision-Language Models with Diverse Visual Prompts: Framework Design and Comprehensive Exploration of Prompt Variations. In *NAACL*. ACL.