

MAOAM: Unified Object & Material Selection with Vision-Language Models

JADEN PARK, University of Wisconsin-Madison, USA
VALENTIN DESCHAINTE, Adobe Research, United Kingdom
JASON KUEN, Adobe Research, USA
KANGNING LIU, Adobe Research, USA
ILIJAN GEORGIEV, Adobe Research, United Kingdom
KRISHNA KUMAR SINGH, Adobe Research, USA
YONG JAE LEE, Adobe Research, USA
MICHAEL FISCHER, Adobe Research, United Kingdom

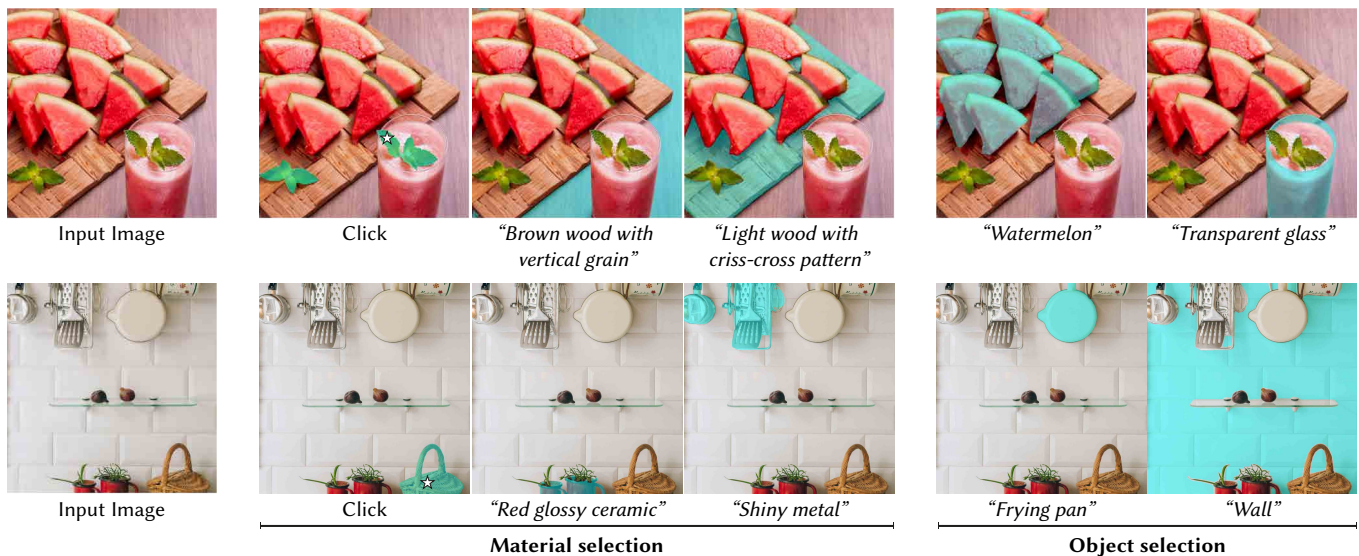


Fig. 1. Our method, MAOAM, enables click- and text-based selection of both materials and objects via single model with a unified interface. Given an input image, users can interact with the selection model via clicks (second column, denoted as overlaid white star) or text queries (third column onwards).

Selection is a core operation in interactive image editing, enabling tasks such as composition or manipulation. To be practically useful, a user should be able to specify and disambiguate the desired selection region through either text- or click-based interactions, and the system should support selecting not only objects but also other criteria, such as materials. Material-based selection can be particularly valuable for tasks like re-texturing surfaces or consistently editing all instances of a specific material in a scene. However, existing

vision-language-model (VLM) based selection methods are largely object-centric and typically support only a single interaction modality, limiting their applicability in real editing workflows.

In this work, we thus present Mask Any Object And Material (MAOAM), a unified selection framework that enables precise object- and material-level selection across both text- and click-based interactions. MAOAM leverages a VLM with a segmentation head to produce pixel-accurate masks from user prompts: the VLM interprets the user's selection intent — object- or material-level — and encodes visual entities, attributes, and spatial relations, while the segmentation head decodes the VLM's output token into a mask.

A key challenge is that material selection datasets with text annotations are unavailable. We therefore propose a scalable data generation pipeline: we collect real and synthetic images with material masks, then leverage VLMs to generate material descriptions with rich visual-semantic information. Using the generated data, we train MAOAM with a multi-task objective over click- and text-based selection, along with an auxiliary VQA task derived from the material descriptions to facilitate deeper material understanding.

Despite being trained with uni-modal prompts, our model exhibits an emergent improvement in selection quality when combining text and clicks

Authors' Contact Information: Jaden Park, University of Wisconsin-Madison, USA, jadenpark@cs.wisc.edu; Valentin Deschaintre, Adobe Research, United Kingdom, deschain@adobe.com; Jason Kuen, Adobe Research, USA, kuen@adobe.com; Kangning Liu, Adobe Research, USA, kangningl@adobe.com; Iliyan Georgiev, Adobe Research, United Kingdom, igeorgiev@adobe.com; Krishna Kumar Singh, Adobe Research, USA, krishsin@adobe.com; Yong Jae Lee, Adobe Research, USA, yongl@adobe.com; Michael Fischer, Adobe Research, United Kingdom, mifischer@adobe.com.



This work is licensed under a Creative Commons Attribution 4.0 International License. SIGGRAPH Conference Papers '26, Los Angeles, CA, USA © 2026 Copyright held by the owner/author(s). ACM ISBN 979-8-4007-2554-8/2026/07 https://doi.org/10.1145/3799902.3811186

at inference time, enabling more flexible image editing workflows. Experiments demonstrate accurate and coherent selections across diverse objects, materials, and interaction scenarios, highlighting robustness in practice.

CCS Concepts: • **Computing methodologies** → **Image segmentation**.

ACM Reference Format:

Jaden Park, Valentin Deschaintre, Jason Kuen, Kangning Liu, Iliyan Georgiev, Krishna Kumar Singh, Yong Jae Lee, and Michael Fischer. 2026. MAOAM: Unified Object & Material Selection with Vision-Language Models. In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers (SIGGRAPH Conference Papers '26)*, July 19–23, 2026, Los Angeles, CA, USA. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3799902.3811186>

1 Introduction

Selection — producing pixel-accurate segmentation masks under user-specified criteria — is a fundamental operation in interactive image editing, enabling downstream tasks such as compositing, re-lighting, and appearance manipulation. In practice, users vary both in *how* they interact with the system — e.g., clicks, which are precise and local, or text prompts, which can describe the desired region using complex visual qualities or spatial relations — and in *what* they would like to select (e.g., objects or materials). Importantly, a single interaction modality can be insufficient to fully disambiguate the intended selection criterion. For example, consider *selecting all ceramic plates* in a kitchen with both ceramic and plastic plates, and a ceramic pot. Uni-modal queries such as “*select the plates*” would include the plastic ones, while “*select the ceramic*” would also include the pot. Clicking each ceramic plate quickly becomes impractical when many instances are present. The most natural and efficient query “*select all ceramic plates*” requires joint reasoning over object and material within a single model. Hence, an ideal interactive selection model should support multiple interaction modalities (clicks, text, or both) and diverse selection criteria beyond objects alone.

Recent segmentation models — ranging from the Segment Anything Model (SAM) series [Kirillov et al. 2023a; Ravi et al. 2024] to VLM-based approaches [Lai et al. 2024; Rasheed et al. 2024; Yuan et al. 2025; Zhang et al. 2024] — mostly support a single interaction modality (clicks or text) and exclusively focus on object-level selection. Material selection, in contrast, has different semantics and structure: a single material may span multiple objects (e.g., metal across fixtures) or appear as disjoint sub-regions within an object (e.g., the wooden legs of a chair). This capability is particularly valuable for tasks like re-texturing surfaces or consistently editing all instances of a specific material in a scene. Prior works on material selection [Fischer et al. 2026; Guerrero-Viu et al. 2025; Sharma et al. 2023] only allow click-based interactions, limited to local spatial cues and thus cannot explicitly express global or relative semantic criteria (e.g., *all glossy metal*; *the fabric in the back*) or disambiguate whether the intended criterion is an object or a material.

We propose a unified selection framework supporting both object- and material-level segmentation across both text- and click-based interactions. Given a natural-language prompt or a click, our system produces a mask consistent with the user’s selection criterion. We leverage VLMs to handle both modalities — text lets users describe complex visual semantics and attributes, while clicks specify precise locations. We build on prior VLM-based approaches [Lai et al. 2024;

Table 1. Supported criteria and input modalities across methods. Only MAOAM supports both object and material selection under both click- and text-based interaction.

	LISA	GLaMM	Sa2VA	SAM3	SAMa	Materialistic	MAOAM
Objects	✓	✓	✓	✓	✗	✗	✓
Materials	✗	✗	✗	✗	✓	✓	✓
Text Input	✓	✓	✓	✓	✗	✗	✓
Click Input	✗	✗	✗	✓	✓	✓	✓

Rasheed et al. 2024; Yuan et al. 2025] and extend it to unified object and material segmentation. Conditioned on the user prompt, the VLM processes the image and emits a segmentation token, which a segmentation head decodes into the mask, encoding the visual and spatial information needed for the target selection. To strengthen material understanding, we utilize a multi-task objective combining segmentation with a VQA-based reasoning task.

A central challenge in training our unified selection model is the lack of text-annotated material datasets. Existing segmentation datasets are object-centric and do not generalize to materials, as one object may comprise multiple materials, and multiple objects may consist of the same material. In addition, the material assignments in existing datasets [Fischer et al. 2026; Guerrero-Viu et al. 2025] are semantically inconsistent (e.g., a plate made of fabric) which limits their use in teaching a model about real-world material appearance.

To address these limitations, we first collect real-world and rendered synthetic sets of images with highly precise material masks. We then propose a scalable data generation pipeline that leverages advanced VLMs to densely annotate materials with text descriptions rich in visual semantics and spatial information. With the generated descriptions, we carefully formulate VQA questions to encourage fine-grained understanding of material qualities in the text space, and jointly train our model along with the selection task.

We show our data generation and training strategy improves material selection and understanding while maintaining competitive object selection performance to several strong baselines. Notably, due to the diversity in our material description data, the model learns to handle text input of varying complexity (e.g., from “*select the wood*” to “*select the red-brown wood with vertical grains and gentle weathering*”), demonstrating semantic grounding of material descriptions in the image space. Although the model is trained with uni-modal data consisting of *either* click- or text-based prompts, we observe an emergent improvement in selection when combining text and clicks during inference, enabling more flexible image editing workflows (see Fig. 1). Our contributions are as follows:

- We propose a unified model that produces object or material selection masks from both click- and text-based interactions.
- We collect selection data with material-level annotations and design a scalable VLM-based data generation pipeline that generates semantically rich and grounded descriptions of materials.
- We empirically validate that our text-data generation pipeline leads to generalization to diverse user interaction patterns, enabling flexible behavior in real editing workflows.

We release our model and test code alongside evaluation data to facilitate further research in this direction here.

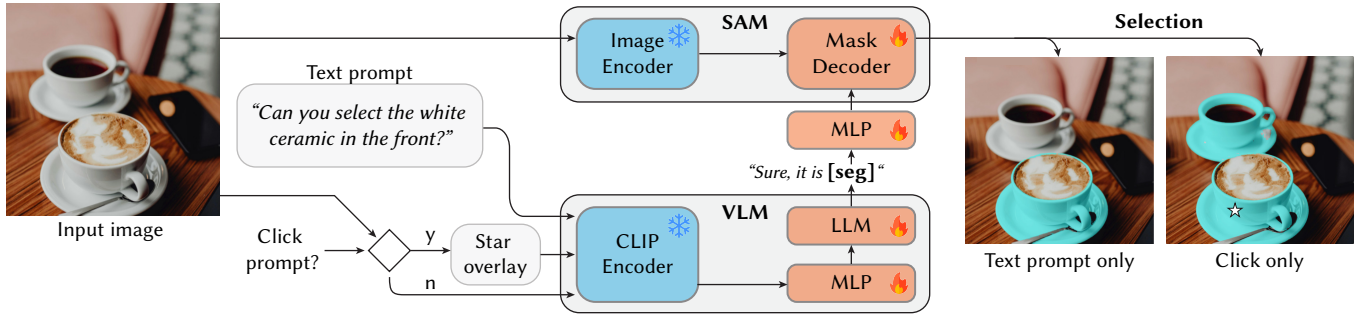


Fig. 2. MAOAM architecture overview. Given an input image, MAOAM takes a task prompt specifying the selection criteria (i.e. objects or materials) alongside a user prompt to specify the desired selection in click or text. If a click is provided, stars are overlaid onto the image as visual cues. The VLM’s CLIP-encoder and projection layer encode the image features into the language-space. The LLM processes the features and produces a segmentation-token, which is projected via another MLP as a prompt for the mask decoder. We train our model with a multi-task objective on click- and text-based selection alongside VQA, where blue and red denote frozen and trainable parameters, respectively. In the above example, notice how with click-only, all areas made of white ceramic are selected, while with the text prompt, only the front object with white ceramic is selected.

2 Previous Work

Segmentation and Selection. Image segmentation has long been a central problem in computer vision (for a survey, see [Minaee et al. 2022]). Recent advances include bipartite matching with object queries [Carion et al. 2020], unified decoder-only formulations for panoptic, semantic, and instance segmentation [Cheng et al. 2021], and multi-scale feature processing for improved accuracy [Cheng et al. 2022]. Beyond automatic segmentation, interactive selection has gained significant attention, particularly with the introduction of SAM [Kirillov et al. 2023b], which enables user-guided mask generation via clicks, boxes, or points. Subsequent works have improved mask quality [Ke et al. 2023] and extended the framework to video [Ravi et al. 2024] and text [Carion et al. 2025].

However, these methods are fundamentally object-centric. In concurrent work with SAM, Materialistic [Sharma et al. 2023] introduced selection based on material similarity, allowing joint selection of image regions sharing the same material. Follow-up work further improved selection granularity [Guerrero-Viu et al. 2025] and extended material-based selection to video and 3D [Fischer et al. 2026]. Unlike prior approaches that separate object and material selection, we propose a single, unified model that supports both object- and material-level selection with text and/or click-based interactions, enabling more flexible and expressive user interaction.

VLM-based segmentation. Recent work has augmented VLMs and multi-modal LLMs (MLLMs) with pixel-level outputs to enable grounded image understanding and segmentation from natural language. Early generalist models like X-Decoder [Xueyan et al. 2023a] and SEEM [Xueyan et al. 2023b] established unified interfaces bridging pixel-level masks and vision-language semantics for segmentation tasks. LISA [Lai et al. 2024] introduced reasoning segmentation for implicit, knowledge-intensive queries, while GLaMM [Rasheed et al. 2024] advanced dense grounding by generating natural language responses intertwined with segmentation masks. Subsequent work has expanded language-guided segmentation along multiple axes: multi-target referring and explicit rejection

handling (GSVA [Xia et al. 2024]), unified training across heterogeneous tasks (PSALM [Zhang et al. 2025]), efficient vision-language feature fusion (EVF-SAM [Zhang et al. 2024]), and reasoning-centric approaches using chain-of-thought guidance (ThinkFirst [Kao et al. 2025]) or reinforcement learning (Seg-Zero [Liu et al. 2025]). In contrast to these object-centric frameworks, our work is the first unified VLM-driven approach to support both object- and material-level selection across text and click inputs.

Visual Grounding. Visual grounding studies the alignment of linguistic concepts and image regions, with the goal of localizing where a model associates words or phrases within the visual domain. Early approaches typically predicted bounding boxes or coarse spatial regions corresponding to textual queries [Plummer et al. 2015; Yu et al. 2018], while more recent work analyzes cross-attention activations between vision and language [Kang et al. 2025].

Grounding provides valuable model interpretability, but is not designed to produce editing-quality masks, particularly for fine-grained structures or material boundaries. While our method implicitly employs grounding-based localization, our goal is high-precision, interactive selection by combining VLM-based semantic understanding with explicit user input to support accurate object- and material-level selection suitable for image editing workflows.

3 Method

We describe our approach to training a unified selection model supporting click- and text-based interactions for both object- and material-level selection. Fig. 2 shows an overview of our method.

3.1 Architecture

We build upon state-of-the-art VLM-based object segmentation architectures [Lai et al. 2024; Rasheed et al. 2024; Yuan et al. 2025], which couple a vision-language model (VLM) that processes the input image and selection prompt with a SAM-based mask decoder [Kirillov et al. 2023b], and extend them to support unified object and material segmentation. Many such systems further incorporate additional modules (e.g., 4-level region encoders or localized feature extractors) to inject spatially focused visual features directly

into the decoder. In contrast, inspired by recent work that provides explicit visual cues to the VLM [Cai et al. 2024a], we provide the VLM with the full input image either (i) with a star overlay indicating the click location or (ii) paired with a referring text prompt. This allows us to preserve a clean, unified input interface across different architectures. The visual overlay is also required to ground the VLM in the click location for the VQA task (Section 3.2). We note the choice of star shape is not important — any marker easily recognizable and unlikely to appear in natural images suffices (see Section 6 for ablation of alternative click representations.) We then enforce all relevant information — selection intent (object vs. material), visual attributes (e.g., color, reflectance, roughness), and spatial relations — to be encoded into a [SEG] token. An MLP projects this token from the textual embedding space into the visual feature space before passing it to the mask decoder, receiving no other conditioning regarding the intended selection. Finally, the decoder outputs a high-resolution (1024×1024) dense selection mask.

Our model is trained jointly on material and object data (Section 4), allowing a rich, material-aware feature representation while retaining strong object selection capabilities. Importantly, we adapt the SAM-based mask decoder, trained on object-centric data, to also segment material-specific selections. As a result, our model can select spatially disjoint regions that share the same material (e.g., all metal fixtures in a scene) while also supporting object selection.

During inference, the VLM encodes the user’s selection intent based on the prompt, allowing the same image to yield different selections depending on the interaction, while the decoder simply consumes the [SEG] token and produces the mask. Fig. 2 illustrates this flexibility: for click-based material selection (using a standard prompt such as “segment everything made of the same material as where the star is”), the model selects both cups. For text-based selection with a prompt such as “the white ceramic in the front,” the model selects the specific cup referenced by the user.

3.2 Training Objective

Let x_{img} and x_{txt} denote the input image and the user text prompt, respectively. For click-based selection, we generate a visual prompt by overlaying a star marker at the click location, yielding x_{img}^* . We intentionally represent clicks as a visual overlay rather than additional coordinate tokens, as this keeps the input modality uniform across tasks and allows the VLM to reason about the click in the same pixel space as other visual cues (e.g., nearby boundaries, texture, and context), which is important for material selection.

The input image (x_{img} or x_{img}^*) is first encoded by a CLIP vision encoder [Radford et al. 2021]. The visual embedding is mapped into the LLM token embedding space via MLP, producing a fixed-length sequence z_{img} . This projection serves two purposes: (i) it enables seamless fusion of visual and textual information within the LLM, and (ii) it avoids introducing a cross-modal module, keeping the architecture simple and compatible with pretrained LLM backbones.

The LLM processes the concatenated sequence ($z_{\text{img}}, x_{\text{txt}}$) and is trained to emit a newly introduced special token [SEG] whose hidden representation summarizes the full selection specification: the user’s intent (object vs. material), relevant visual attributes (e.g., color, reflectance, roughness), and spatial relations (e.g., “in the front”, the click position). By bottlenecking all selection information

into a single [SEG] embedding, we enforce that the VLM produces an explicit, task-relevant representation that is compatible with the interaction type (click vs. text) and can be consumed by the decoder without additional hand-engineered prompts.

To generate the final mask, we feed the [SEG] embedding to SAM’s prompt encoder [Kirillov et al. 2023b], and condition the SAM mask decoder to predict a dense, high-resolution selection mask. Importantly, the mask decoder receives no other task-specific conditioning beyond the [SEG] embedding; this design isolates language/intent understanding within the VLM, while leveraging SAM’s strong mask prior for accurate boundary delineation.

We train the model with a multi-task objective that combines click-based selection, referring-text selection, and VQA:

$$\mathcal{L}(x) = \lambda_1 \mathcal{L}_{\text{click}}(x_{\text{img}}^*) + \lambda_2 \mathcal{L}_{\text{ref}}(x_{\text{img}}, x_{\text{txt}}) + \lambda_3 \mathcal{L}_{\text{vqa}}(x_{\text{img}}^*, x_{\text{txt}})$$

where $\mathcal{L}_{\text{click}}$ and \mathcal{L}_{ref} denote click-based and text-based selection training, respectively, and \mathcal{L}_{vqa} denotes the VQA loss. We use the star-overlaid image x_{img}^* for click-based selection to explicitly convey the interaction point and align supervision with the user-provided click. We additionally use x_{img}^* for VQA to train the VLM to attend to the clicked region when answering questions, thereby aligning its visual reasoning with the same interaction cue used for click-based selection. In contrast, referring-text selection uses the original image x_{img} , encouraging purely language-driven grounding without relying on auxiliary spatial markers.

The selection task losses ($\mathcal{L}_{\text{click}}$ and \mathcal{L}_{ref}) use the same combination of (i) token-level cross-entropy loss for language modeling (including the [SEG] token) and (ii) per-pixel mask supervision using binary cross-entropy (BCE) loss and DICE loss [Milletari et al. 2016]. We include both BCE and DICE to balance per-pixel optimization with robust region-level overlap under class imbalance. For VQA, we use token-level cross-entropy loss. Further details on the network architecture, training time, GPU requirements, model size, or inference latency are reported in the Suppl. Section 2.

4 Dataset

Training a unified selection model requires material datasets with dense mask annotations and rich text descriptions. While existing material datasets provide high-quality mask annotations, they either lack textual descriptions [Fischer et al. 2026; Guerrero-Viu et al. 2025; Sharma et al. 2023], are available only as flat material maps [Vecchio and Deschaintre 2024], or are tied to a specific domain, e.g., fabrics [Deschaintre et al. 2023]. On the other hand, grounding and reasoning datasets (e.g., RefCOCOg [Mao et al. 2016]) have textual annotations but are not object centric and typically too short for fine-grained material selection. To address this, we collect new material data with dense mask annotations and develop a VLM-based pipeline to generate detailed text descriptions.

Unlike recent reasoning segmentation datasets that use allusive queries (e.g., “select the food with the highest protein” when a steak is shown), our descriptions directly refer to visual material attributes with semantically aligned information — we argue that users are unlikely to issue such indirect queries for selection tasks. Moreover, because our descriptions capture fine-grained visual details, the model learns associations between text and appearance, enabling generalization to diverse user prompts at inference time.




Fig. 3. Example annotations. We show the images’ dense, per-pixel material annotations overlaid in the second row. We additionally show two versions of our generated text descriptions for these materials; one including the object (top row) and one solely focusing on the material (bottom row).

4.1 Material Mask Data

We collect material mask data from both real and synthetic sources to capture natural diversity and precise, controlled annotations.

For real images, we collect $\sim 8\text{K}$ images from Pexels.com and hand-annotate them, with the help of external users resulting in $\sim 49\text{K}$ material masks. We denote this dataset as REALMAT.

For synthetic images, the subset of Guerrero-Viu et al. [2025] is not suitable as materials are assigned irrespective of semantics – e.g., a sofa with a checkerboard pattern of wood and stone which is unrealistic and incompatible for reasoning about real-world images. We instead render images with semantically correct materials – e.g., a sofa made of leather or fabric – using Blender and 132 scenes from Evermotion.com with pre-defined camera paths. We gather $\sim 5.5\text{K}$ images along with $\sim 55\text{K}$ material masks, and denote it SYNMAT.

Finally, we use SAMa [Fischer et al. 2026], consisting of $\sim 1.3\text{K}$ images and $\sim 3.3\text{K}$ material masks, which we denote as SAMa. Although its material assignments are not semantically meaningful, they are consistent across object parts. Both SYNMAT and SAMa consist of video frames across multiple viewpoints, providing varying mask shapes and material-light responses.

This amounts to $\sim 104\text{K}$ material annotations from $\sim 15\text{K}$ images. We show in Table 6 that the real and synthetic datasets are complementary: training on both significantly improves performance on both evaluation sets compared to training on either subset. For further details on train and test splits, refer to Suppl. Section 1.

4.2 Material Data Generation Pipeline

Human annotation provides high-quality labels, but is prohibitively expensive at the scale required for our setting. To efficiently curate material descriptions, we use VLMs to generate candidate annotations and incorporate quality control through model-based verification and targeted human review. This hybrid pipeline enables scalable annotation while maintaining high-quality supervision.

We describe the annotation and VQA generation processes below, followed by verification. Unless indicated otherwise, we use Qwen3-VL-235B-A22B-Thinking [Bai et al. 2025a] as our annotation model.

Description generation. We adopt Set-of-Marks (SoM) prompting [Yang et al. 2023] to improve spatial grounding in vision–language reasoning. Given an input image, we provide the VLM with both a SoM-overlaid image (for an example, see Fig. 4) and a mask-overlaid image indicating which regions share the same material.

For each marked region, we generate three types of descriptions. The first consists of a short material description augmented with an entity label, such as an object name or category (e.g., “chair”). The second combines a short material description with explicit spatial information – either absolute image position (e.g., “bottom right corner”) or relative to other objects (e.g., “above the table”). The third is a longer, self-contained material description that does not rely on contextual cues. We generate 6 variants of different lengths (10–50 words) which are randomly sampled during training.

Verification. The generated descriptions occasionally suffer from incorrect grounding or instruction following – e.g., including entity names in descriptions that should only contain material attributes. To address this, we perform verification using Qwen3-VL-235B-A22B-Thinking model as a verifier. In a manual audit of 500 sampled descriptions, a substantial amount of observed failures was fixed in the verification stage. Hence, we use the verified descriptions by default. However, the validation set is further inspected manually for accuracy in evaluating models. The filtering criteria are accuracy and unambiguity for both text-based segmentation and VQA. After manual filtering, we retain 1,797 out of 2,216 REALMAT samples, 2,458 out of 3,072 SYNMAT samples, and 258 out of 352 SAMa samples, resulting in 4,513 out of 5,640 total ($\sim 80.0\%$).

Visual Question Answering (VQA) Generation. Training on VQA encourages fine-grained material knowledge through reasoning in text. We formulate two variants of a four-way multiple-choice task.

We select a material in the image and retrieve its description from the previous stage as the answer. The distractors are constructed by sampling descriptions from other regions within the same image; if fewer than three distinct materials are present, we sample from other

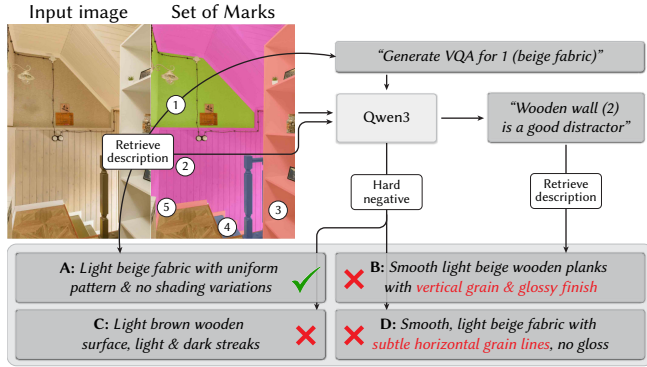


Fig. 4. Example of VQA generation with hard negative mining. The VLM receives the SoM-overlaid image (denoted with numbered circles) in order to select the distractor and paraphrase the material description for both the answer and the negative candidate to create three distractors. The paraphrased parts are encoded in color (see Section 4.2 for details).

images. This requires the model to distinguish between visually present materials based on their descriptions.

The second variant introduces hard negative mining by generating visually plausible but incorrect alternatives — e.g., changing "brown wood with dark streaks" to "horizontally grained light wood," as shown in Fig. 4. The negatives are generated for the answer and one sampled distractor. In Section 6, we show that training on VQA questions improves model performance and understanding.

4.3 Training Data Composition

Finally, to train a unified model for both object- and material-level selection, we incorporate publicly available object segmentation datasets. Notably, we show in Table 5 that mixing in the object data does not deteriorate the model’s performance on material selection.

We use the RefCOCO, RefCOCO+, and RefCOCOG [Kazemzadeh et al. 2014; Mao et al. 2016] referring segmentation datasets for text-based object selection. For click-based object selection, we use EntitySeg [Qi et al. 2023] which consists of high-quality object selection masks from real-world images. This results in a total of ~190K training samples with an approximate 1:1 ratio of material- and object-centric data, spanning diverse selection prompts and criteria. For more details on the datasets, see Suppl. Section 1.

5 Evaluation

In this section, we evaluate baselines along with MAOAM and report both quantitative and qualitative results across object and material selection tasks in both click- and text-based interactions.

5.1 Quantitative Evaluation

We evaluate MAOAM on material- and object-selection tasks using both text- and click-based inputs, where applicable. We compare against recent state-of-the-art VLM-based segmentation models, including GLaMM [Rasheed et al. 2024] (GranD-pretrained), Sa2VA [Yuan et al. 2025], LISA [Lai et al. 2024] (LISA-13B-llama2-v0-explanatory), SAM3 [Carion et al. 2025], and the material-selection

baseline Materialistic [Sharma et al. 2023]. We report mean Intersection over Union (mIoU) and F_1 , the harmonic mean of precision and recall; higher is better for both.

Material Selection. The first two columns of Table 2 report material selection performance on REALMAT, SYNMAT, and SAMA across text- and click-based selection. MAOAM achieves substantial mIoU improvements over baselines (e.g., 67.5% avg. mIoU over Sa2VA) in text-based selection, all of which do not perform well on material reasoning. For click-based selection, the performance gap is more pronounced as most baselines are not trained to process visual click prompts. MAOAM outperforms Materialistic which is trained for click-based material selection, by 35.5% average mIoU. MAOAM performs strongly across all material datasets and interaction modalities, with high F_1 scores indicating balanced selections.

Object Selection. The third and last block of Table 2 reports performance on RefCOCO, RefCOCO+, RefCOCOG and EntitySeg for text- and click-based object selection. Albeit by a smaller margin (e.g., 33.9%, 14.6% and 3.2% average mIoU improvement over GLaMM, LISA, and Sa2VA), MAOAM outperforms pretrained baselines, indicating joint training on materials does not degrade object selection.

Visual Question Answering. Table 3 reports VQA performance on the two question variants described in Section 4.2 (denoted Q1 and Q2). We compare to Qwen2.5-VL-7B [Bai et al. 2025b], a pretrained VLM with strong VQA performance. MAOAM achieves high accuracy on both, while Sa2VA and Qwen2.5-VL-7B, perform poorly, indicating that their ability to interpret visually differing stimuli from materials and their descriptions is limited.

Notably, MAOAM performs better on Q2 than Q1, opposite to the baselines, aligning with Cai et al. [2024b]’s finding that models with stronger domain understanding can recognize hard-negative variants as related. MAOAM’s strong Q2 performance suggests it has acquired fine-grained material understanding, whereas the baselines are instead confused by the similar options. In Section 6, we show that incorporating VQA during training improves downstream selection performance despite the different objectives.

5.2 Qualitative Evaluation

We further provide qualitative comparisons in Figs. 5 and 6. MAOAM produces higher-precision masks with stronger grounding quality and generalizes to a wider range of material descriptions, while existing models often struggle when the selection criterion falls outside their object-centric vocabulary. We attribute this robustness to our description generation pipeline, which provides detailed material annotations with diverse reasoning patterns, including spatial relations. Although MAOAM is trained with longer, detailed material descriptions, the qualitative examples use shorter inference prompts, demonstrating generalization to more natural user queries. More examples are shown in Suppl. Fig. 2.

We next highlight several desirable properties of our model.

Emergent multimodal interaction. Although MAOAM is trained with uni-modal prompts, combining text and click during inference improves mask quality (Fig. 7). We highlight that this behavior emerges without explicit supervision, as our training data explicitly consists of uni-modal interactions.

Table 2. Comprehensive evaluation of **MAOAM** against state-of-the-art models. We report performance across material-specific datasets (REALMAT, SYNMAT, SAMA) and object-centric datasets (RefCOCO, RefCOCO+, RefCOCOg and EntitySeg) using both text-based and click-based selection. Our model significantly outperforms existing methods on material selection datasets while retaining its object-selection capability. Materialistic does not allow text-based selection.

Method	Material (Text-based)						Material (Click-based)						Object (Text-based)						Object (Click-based)	
	REALMAT		SYNMAT		SAMA		REALMAT		SYNMAT		SAMA		RefCOCO		RefCOCO+		RefCOCOg		EntitySeg	
	mIoU ↑	F1 ↑	mIoU ↑	F1 ↑	mIoU ↑	F1 ↑	mIoU ↑	F1 ↑	mIoU ↑	F1 ↑	mIoU ↑	F1 ↑	mIoU ↑	F1 ↑	mIoU ↑	F1 ↑	mIoU ↑	F1 ↑	mIoU ↑	F1 ↑
SAM3	0.263	0.293	0.224	0.253	0.068	0.074	0.538	0.624	0.505	0.599	0.623	0.710	0.433	0.472	0.329	0.364	0.422	0.456	0.664	0.748
Materialistic	–	–	–	–	–	–	0.524	0.709	0.680	0.884	0.535	0.718	–	–	–	–	–	–	0.147	0.256
LISA	0.332	0.396	0.319	0.383	0.215	0.259	0.129	0.163	0.094	0.124	0.056	0.074	0.732	0.797	0.638	0.702	0.665	0.733	0.209	0.249
GLaMM	0.349	0.415	0.328	0.396	0.260	0.305	0.185	0.238	0.159	0.210	0.101	0.129	0.616	0.692	0.521	0.597	0.603	0.679	0.364	0.423
Sa2VA	0.473	0.552	0.431	0.502	0.471	0.538	0.260	0.317	0.242	0.289	0.378	0.452	0.781	0.840	0.729	0.782	0.749	0.810	0.435	0.495
MAOAM (Ours)	0.740	0.798	0.608	0.669	0.685	0.754	0.808	0.868	0.766	0.835	0.747	0.823	0.809	0.895	0.744	0.853	0.778	0.875	0.821	0.901

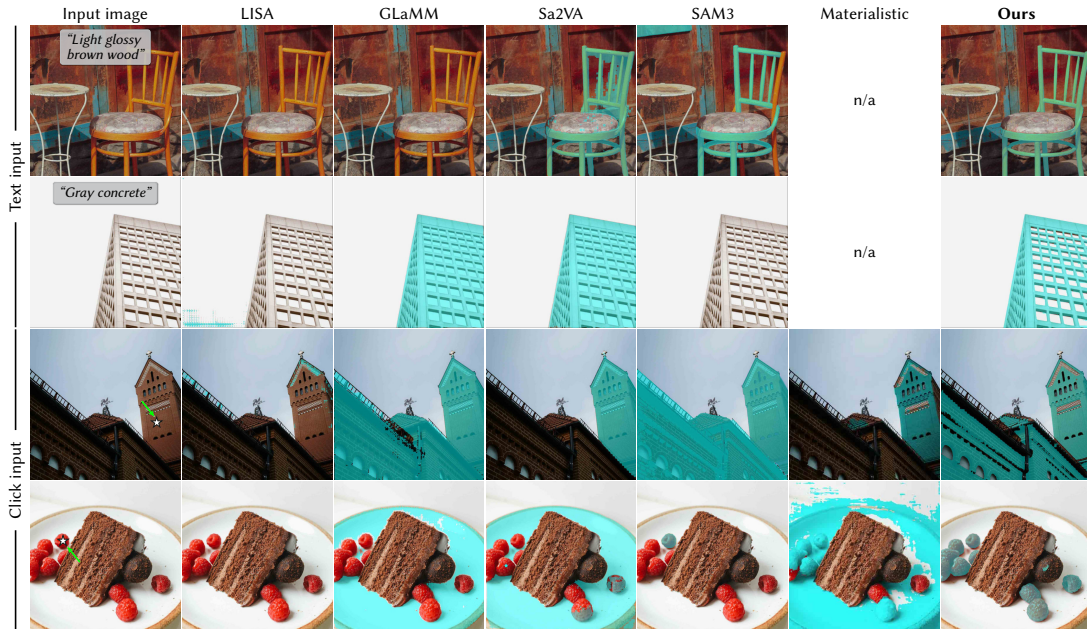


Fig. 5. Material selection. We compare our method against baselines on a material selection task, both click- and text-based (first two and last two rows, respectively). LISA, GLaMM and SAM3 occasionally produce an empty mask when the selection criterion is too complicated or foreign to their vocabulary. Materialistic does not support text-based queries and is denoted n/a.




Fig. 6. Object selection. We compare our method against several baselines on an object selection task. Materialistic neither supports text-based queries, nor object selection. Our method performs on par with the baselines, highlighting that reasoning about materials does not lead to deterioration on the object-level.

Table 3. VQA performance on material datasets (REALMAT, SYNMAT, SAMA). We report multiple-choice accuracy here.

	REALMAT		SYNMAT		SAMA	
	Q1	Q2	Q1	Q2	Q1	Q2
Qwen2.5-VL-7B	0.584	0.318	0.543	0.288	0.480	0.564
Sa2VA	0.484	0.311	0.510	0.305	0.380	0.432
MAOAM (Ours)	0.858	0.974	0.795	0.979	0.749	0.858




Fig. 7. Emergent refinement. Starting from a text prompt, we show the selection result (material top, object bottom row) with an increasing number of interactions, i.e., more explicit guidance for the selection masks, closely resembling a realistic interaction scenario. Note that our model was never explicitly trained to handle both click- and text-prompts.

Spatial and semantic reasoning. MAOAM interprets spatial relations, leveraging the spatial descriptions in our training data (Fig. 8). The model also handles prompts that span multiple objects – e.g., “select everything made out of metal” – and produces coherent masks across spatially disjoint regions.

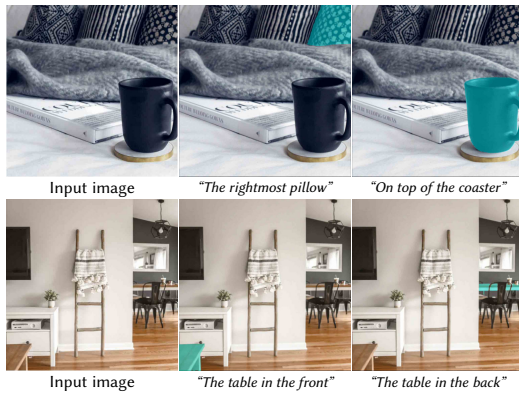


Fig. 8. Spatial reasoning. Due to our training data (Section 4.2), MAOAM understands spatial queries relative to other entities and select accordingly.

Flexible Selection. Given the same click location, different prompts yield object- or material-level selections (Fig. 9). For example, a click on a sofa can select either the entire sofa or the fabric depending on

the task prompt. Our model also disambiguates between color and material when both could apply (Fig. 10), demonstrating flexibility over various selection criteria and adherence to user instructions.



Fig. 9. Varying text prompts. The same visual prompt (star marker) is interpreted differently with varying text prompt.




Fig. 10. Disambiguation. Although both objects are yellow, MAOAM infers the masks for both click and text prompts.

Mask quality. MAOAM produces fine-grained, high-precision masks that capture detailed boundaries (Fig. 11). Unlike SAM-HQ [Ke et al. 2023], our model achieves this without explicit guidance, showing that SAM decoder can produce fine-grained masks.




Fig. 11. Mask quality. Our model performs well on intricate selection targets producing fine-grained masks. Images are from SAM-HQ [Ke et al. 2023].

Image Editing. MAOAM produces high-quality, material-aware masks from both click- and text-based interactions, enabling real-world image editing workflows (Fig. 12) such as replacing all regions of one material with another material.




Fig. 12. Editing. The selection output (masks displayed as insets) for click- and text-based queries can be used to edit materials in the image.

6 Ablation & Discussion

All ablation experiments that require training are done on the GLaMM-based MAOAM model due to lower compute requirements.

Model choices in data generation. We ablate the effect of model size and verification in our description generation pipeline, comparing descriptions generated with Qwen3-VL-8B and Qwen3-VL-235B-A22B-Thinking in Table 4. Verification consistently improves metrics for descriptions generated with the 8B model across all three datasets. Although the gain is not as consistent for the descriptions generated with the 235B model, we empirically validated that the verification leads to more grounded descriptions and better instruction following, and hence set this configuration as default.

Table 4. Description generation pipeline analysis. We report text-based material selection performance after training with the annotations produced by the respective pipeline configurations (rows). Columns Generate and Verify denote the model sizes used for description generation and verification.

Generate	Verify	REALMAT		SYNMAT		SAMA	
		mIoU ↑	F1 ↑	mIoU ↑	F1 ↑	mIoU ↑	F1 ↑
8B	–	0.644	0.712	0.532	0.604	0.626	0.703
8B	235B	0.657	0.726	0.554	0.628	0.651	0.721
235B	–	0.676	0.744	0.575	0.649	0.661	0.738
235B	235B	0.675	0.746	0.588	0.665	0.654	0.730

Training with objects helps. We ablate data composition by comparing MAOAM trained on material-only versus full mixed data (Materials + Objects). Table 5 reports material selection performance for both text- and click-based interactions. Despite joint object-material training, Materials + Objects remains competitive across all datasets and even improves some click-based results. This suggests that joint training preserves material understanding while adding object selection and joint object-material reasoning capabilities (Suppl. Fig. 3) at no cost to material performance.

Table 5. Objects training data ablation. We report mIoU for text- and click-based material selection. Joint training on materials and objects maintains competitive performance compared to model trained on material only.

Training data	REALMAT		SYNMAT		SAMA	
	Text	Click	Text	Click	Text	Click
Materials only	0.675	0.756	0.588	0.730	0.654	0.730
Materials + Objects	0.670	0.760	0.582	0.726	0.661	0.756

Training with synthetic data helps. A significant portion of our material training data (REALMAT, SYNMAT, and SAMA) is synthetic. We ablate this by training on data subsets. Table 6 shows that training on both datasets improves performance on both datasets by up to 9.15%. Full training remains strongest, improving over single-dataset training by up to 21.49%. This shows that synthetic data provide complementary supervision that transfers to real images, reducing reliance on human annotations which can be costly.

Table 6. Material training data ablation. We report mIoU for text- and click-based material selection. Joint training on REALMAT and SYNMAT improves performance on both evaluation sets. Training with all data performs best.

Training data	REALMAT		SYNMAT		SAMA	
	Text	Click	Text	Click	Text	Click
REALMAT only	0.579	0.694	–	–	–	–
SYNMAT only	–	–	0.484	0.683	–	–
SAMA only	–	–	–	–	0.595	0.665
REALMAT + SYNMAT	0.632	0.723	0.520	0.696	–	–
All combined	0.675	0.756	0.588	0.730	0.654	0.730

Robustness to input text length. We generate six referring descriptions of varying length (10~50 words) and detail. The descriptions contain compositional reasoning and attribute combinations, enabling the diverse use cases shown in Figs. 8 to 10. During inference, we observe that MAOAM generalizes to short, natural user prompts. To quantify MAOAM’s robustness to input text length, we evaluate on all six different text prompts, group them by length (short, medium, long) and report the metrics in Table 7. Results show that performance remains stable across prompt lengths, with low variance on all three material benchmarks. This suggests that MAOAM does not rely on a fixed prompt template or length, but learns to ground the relevant material cues expressed in text.

Table 7. Prompt length analysis. We report text-based material selection performance using short, medium, and long descriptions, with two prompt variants for each length. MAOAM remains robust across prompt lengths.

Prompt	REALMAT		SYNMAT		SAMA	
	mIoU ↑	F1 ↑	mIoU ↑	F1 ↑	mIoU ↑	F1 ↑
Short	0.745	0.801	0.622	0.678	0.643	0.726
Medium	0.739	0.796	0.595	0.660	0.712	0.771
Long	0.736	0.796	0.607	0.668	0.701	0.765
Mean	0.740	0.798	0.608	0.669	0.685	0.754
Std	0.004	0.002	0.011	0.007	0.030	0.020

Click representations. We compare the star-overlay with alternative visual inputs. In Table 8, “Coordinates” indicates providing the input coordinate to the mask decoder, whereas “BBox” indicates overlaying a bounding box on the target object. We note that “BBox” is how the original GLaMM model was trained. However, star-overlay achieves the best performance across all three material benchmarks. We hypothesize that the star-overlay provides the VLM with the grounding cue directly, allowing it to encode salient information in the [SEG] token. In contrast, “Coordinates” introduces the user query information at the decoding stage only, while “BBox” is less precise and may contain multiple different materials, confusing the model and leading to inferior performance. We emphasize the specific choice of a star is not essential; the marker only needs to be distinctive and easy for the model to recognize.

Table 8. Click representation analysis. We report click-based material selection performance for different click representations. Star-overlay provides the best performance while preserving a unified input interface.

Click representation	REALMAT		SYNMAT		SAMA	
	mIoU \uparrow	F1 \uparrow	mIoU \uparrow	F1 \uparrow	mIoU \uparrow	F1 \uparrow
Star-overlay	0.760	0.873	0.726	0.830	0.756	0.852
Coordinates	0.755	0.843	0.714	0.812	0.727	0.822
BBox	0.700	0.818	0.668	0.789	0.707	0.785

Presence of stars in the original image. We overlay a 32×32 star at the click location, selecting its color from 10 candidates to maximize contrast with the region of interest. This naturally raises a concern when star-shaped objects appear in the image, as they could act as misleading grounding signals. Fig. 13 demonstrates MAOAM’s robustness when star-shaped objects of varying sizes appear. In all cases, MAOAM correctly interprets the user-provided click and segments the intended region of interest. In the second row, we further test the text-only prompt “select the yellow star” without any click input, and MAOAM correctly segments the yellow stars.

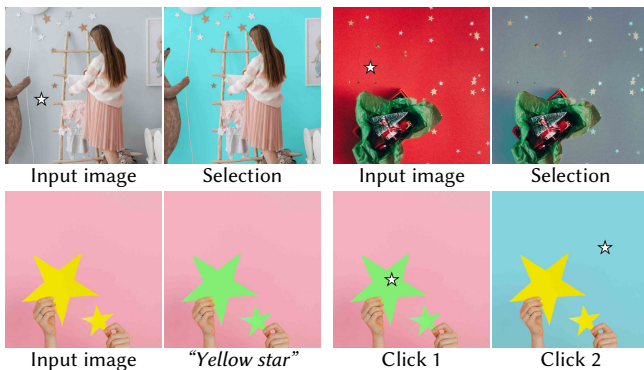


Fig. 13. Star robustness. MAOAM remains robust when star-shaped objects appear in the image. In the second row, the text prompt is “yellow star” and no click is provided; MAOAM correctly segments the yellow stars, despite using colored star overlays only as small click markers.

Joint material and object reasoning. In Fig. 14, we provide additional examples showing why a single model capable of both material and object selection is useful. Each row contains objects from the same category with varying materials. The first two rows show that MAOAM can select a material-specific subset of the objects using material queries, while also selecting the full object set with an object query (e.g., cooking utensils). In the bottom row, the model further interprets joint material-object queries, such as “brown eggs”, and adjusts its selection accordingly.

Limitations. Our method can underperform in challenging cases. Fig. 15 highlights two failure modes: VLM reasoning and mask decoding. Reasoning is limited by the VLM backbone and may benefit from additional test-time compute via chain-of-thought tokens [Kao et al. 2025]. Mask quality is limited by the SAM decoder and could be improved with refinement modules [Yao et al. 2024].

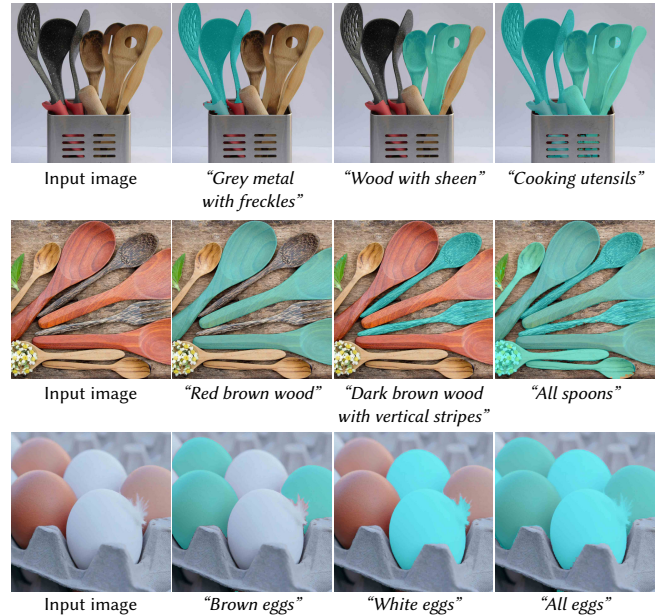


Fig. 14. Joint object-material reasoning. MAOAM can switch between selecting material-specific subsets and full objects. In scenes containing the same object with different materials, material queries select the relevant subset, object queries select all instances, and joint material-object queries (e.g. “brown eggs”) select instances satisfying both criteria.



Fig. 15. Limitations. For the first image, the model fails to distinguish the mortar from the bricks, while for the second image, the produced mask is inaccurate, presumably due to coarse resolution of the VLM image encoder.

7 Conclusion

In this work, we present MAOAM, a unified framework for material- and object-selection with click- and text-based prompts. We propose a scalable, automatic annotation pipeline that enables us to generate a large corpus of rich material text descriptions for visual grounding. We demonstrate strong material selection performance while matching or outperforming object-centric segmentation methods.

Acknowledgments

This work was supported in part by NSF IIS2404180 and Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No. 2022-0-00871, Development of AI Autonomy and Knowledge Enhancement for AI Agent Collaboration). The authors would like to thank Sudeep Katakol for his help in data generation and Zijun Wei, Yash Savani, and Soochahn Lee for helpful discussions.

References

- Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, et al. 2025a. Qwen3-vl technical report. *arXiv preprint arXiv:2502.13923* (2025).
- Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, Humen Zhong, Yuanzhi Zhu, Mingkun Yang, Zhaohai Li, Jianqiang Wan, Pengfei Wang, Wei Ding, Zheren Fu, Yiheng Xu, Jiabo Ye, Xi Zhang, Tianbao Xie, Zesen Cheng, Hang Zhang, Zhibo Yang, Haiyang Xu, and Junyang Lin. 2025b. Qwen2.5-VL Technical Report. *arXiv preprint arXiv:2502.13923* (2025).
- Mu Cai, Haotian Liu, Siva Karthik Mustikovela, Gregory P. Meyer, Yuning Chai, Dennis Park, and Yong Jae Lee. 2024a. Making Large Multimodal Models Understand Arbitrary Visual Prompts. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- Mu Cai, Reuben Tan, Jianrui Zhang, Bocheng Zou, Kai Zhang, Feng Yao, Fangrui Zhu, Jing Gu, Yiwu Zhong, Yuzhang Shang, Yao Dou, Jaden Park, Jianfeng Gao, Yong Jae Lee, and Jianwei Yang. 2024b. TemporalBench: Towards Fine-grained Temporal Understanding for Multimodal Video Models. *arXiv preprint arXiv:2410.10818* (2024).
- Nicolas Carion, Laura Gustafson, Yuan-Ting Hu, Shoubhik Debnath, Ronghang Hu, Dac Suris, Chaitanya Ryali, Kalyan Vasudev Alwala, Haitham Khedr, Andrew Huang, Jie Lei, Tengyu Ma, Baishan Guo, Arpit Kalla, Markus Marks, Joseph Greer, Meng Wang, Peize Sun, Roman Rädle, Triantafyllos Afouras, Effrosyni Mavroudi, Katherine Xu, Tsung-Han Wu, Yu Zhou, Liliane Momeni, Rishi Hazra, Shuangrui Ding, Sagar Vaze, Francois Porcher, Feng Li, Siyuan Li, Aishwarya Kamath, Ho Kei Cheng, Piotr Dollár, Nikhila Ravi, Kate Saenko, Pengchuan Zhang, and Christoph Feichtenhofer. 2025. SAM 3: Segment Anything with Concepts. *arXiv:2511.16719 [cs.CV]* <https://arxiv.org/abs/2511.16719>
- Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. 2020. End-to-end object detection with transformers. In *European conference on computer vision*. Springer, 213–229.
- Bowen Cheng, Ishan Misra, Alexander G. Schwing, Alexander Kirillov, and Rohit Girdhar. 2022. Masked-Attention Mask Transformer for Universal Image Segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 1290–1299.
- Bowen Cheng, Alex Schwing, and Alexander Kirillov. 2021. Per-pixel classification is not all you need for semantic segmentation. *Advances in neural information processing systems* 34 (2021), 17864–17875.
- Valentin Deschaintre, Julia Guerrero-Viu, Diego Gutierrez, Tamy Boubekeur, and Belen Masia. 2023. The Visual Language of Fabrics. *ACM Trans. Graph.* 42, 4, Article 50 (jul 2023), 15 pages. doi:10.1145/3592391
- Michael Fischer, Iliyan Georgiev, Thibault Groueix, Vladimir G Kim, Tobias Ritschel, and Valentin Deschaintre. 2026. SAMa: Material-aware 3D Selection and Segmentation. In *2026 International Conference on 3D Vision (3DV)*. IEEE.
- J. Guerrero-Viu, M. Fischer, I. Georgiev, E. Garces, D. Gutierrez, B. Masia, and V. Deschaintre. 2025. Fine-Grained Spatially Varying Material Selection in Images. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia)* (2025).
- Seil Kang, Jinyeong Kim, Junhyeok Kim, and Seong Jae Hwang. 2025. Your large vision-language model only needs a few attention heads for visual grounding. In *Proceedings of the Computer Vision and Pattern Recognition Conference*. 9339–9350.
- Shiu-hong Kao, Yu-Wing Tai, and Chi-Keung Tang. 2025. Think Before You Segment: High-Quality Reasoning Segmentation with GPT Chain of Thoughts. *arXiv preprint arXiv:2503.07503* (2025).
- Sahar Kazemzadeh, Vicente Ordonez, Mark Matten, and Tamara Berg. 2014. Referring-ItGame: Referring to Objects in Photographs of Natural Scenes. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Lei Ke, Mingqiao Ye, Martin Danelljan, Yu-Wing Tai, Chi-Keung Tang, Fisher Yu, et al. 2023. Segment anything in high quality. *Advances in Neural Information Processing Systems* 36 (2023), 29914–29934.
- Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. 2023b. Segment anything. In *Proceedings of the IEEE/CVF international conference on computer vision*. 4015–4026.
- 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. 2023a. Segment Anything. *arXiv:2304.02643* (2023).
- Xin Lai, Zhuotao Tian, Yukang Chen, Yanwei Li, Yuhui Yuan, Shu Liu, and Jiaya Jia. 2024. Lisa: Reasoning segmentation via large language model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 9579–9589.
- Yuqi Liu, Bohao Peng, Zhisheng Zhong, Zihao Yue, Fanbin Lu, Bei Yu, and Jiaya Jia. 2025. Seg-Zero: Reasoning-Chain Guided Segmentation via Cognitive Reinforcement. *arXiv preprint arXiv:2503.06520* (2025).
- Junhua Mao, Jonathan Huang, Alexander Toshev, Oana Camburu, Alan L Yuille, and Kevin Murphy. 2016. Generation and Comprehension of Unambiguous Object Descriptions. In *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*. 11–20.
- F. Milletari, N. Navab, and S.A. Ahmadi. 2016. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *3DV*.
- Shervin Minaee, Yuri Boykov, Fatih Porikli, Antonio Plaza, Nasser Kehtarnavaz, and Demetri Terzopoulos. 2022. Image Segmentation Using Deep Learning: A Survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 44, 7 (2022). doi:10.1109/TPAMI.2021.3059968
- Bryan A Plummer, Liwei Wang, Chris M Cervantes, Juan C Caicedo, Julia Hockenmaier, and Svetlana Lazebnik. 2015. Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. In *Proceedings of the IEEE international conference on computer vision*. 2641–2649.
- Lu Qi, Jason Kuen, Tiancheng Shen, Jiuxiang Gu, Weidong Guo, Jiaya Jia, Zhe Lin, and Ming-Hsuan Yang. 2023. High-Quality Entity Segmentation. In *International Conference on Computer Vision (ICCV)*.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*. PmlR, 8748–8763.
- Hanoona Rasheed, Muhammad Maaz, Sahal Shaji, Abdelrahman Shaker, Salman Khan, Hisham Cholakkal, Rao M Anwer, Eric Xing, Ming-Hsuan Yang, and Fahad S Khan. 2024. Glamm: Pixel grounding large multimodal model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 13009–13018.
- Nikhila Ravi, Valentin Gabeur, Yuan-Ting Hu, Ronghang Hu, Chaitanya Ryali, Tengyu Ma, Haitham Khedr, Roman Rädle, Chloe Rolland, Laura Gustafson, Eric Mintun, Junting Pan, Kalyan Vasudev Alwala, Nicolas Carion, Chao-Yuan Wu, Ross Girshick, Piotr Dollár, and Christoph Feichtenhofer. 2024. SAM 2: Segment Anything in Images and Videos. *arXiv preprint arXiv:2408.00714* (2024). <https://arxiv.org/abs/2408.00714>
- Prafull Sharma, Julien Philip, Michaël Gharbi, Bill Freeman, Fredo Durand, and Valentin Deschaintre. 2023. Materialistic: Selecting Similar Materials in Images. *ACM Trans. Graph.* 42, 4, Article 154 (jul 2023). doi:10.1145/3592390
- Giuseppe Vecchio and Valentin Deschaintre. 2024. Matsynth: A modern pbr materials dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 22109–22118.
- Zhuofan Xia, Dongchen Han, Yizeng Han, Xuran Pan, Shiji Song, and Gao Huang. 2024. GSA: Generalized Segmentation via Multimodal Large Language Models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 3858–3869.
- Zou Xueyan, Zi-Yi Dou, Jianwei Yang, Zhe Gan, Linjie Li, Chunyan Li, Xiyang Dai, Jianfeng Wang, Lu Yuan, Nanyun Peng, Lijuan Wang, Yong Jae Lee, and Jianfeng Gao. 2023a. Generalized Decoding for Pixel, Image and Language. *CVPR* (2023).
- Zou Xueyan, Jianwei Yang, Hao Zhang, Feng Li, Linjie Li, Jianfeng Wang, Lijuan Wang, Jianfeng Gao, and Yong Jae Lee. 2023b. Segment Everything Everywhere All at Once. *NeurIPS* (2023).
- Jianwei Yang, Hao Zhang, Feng Li, Xueyan Zou, Chunyan Li, and Jianfeng Gao. 2023. Set-of-mark prompting unleashes extraordinary visual grounding in gpt-4v. *arXiv preprint arXiv:2310.11441* (2023).
- Jingfeng Yao, Xinggang Wang, Shusheng Yang, and Baoyuan Wang. 2024. Vitmatte: Boosting image matting with pre-trained plain vision transformers. *Information Fusion* 103 (2024), 102091.
- Licheng Yu, Zhe Lin, Xiaohui Shen, Jimei Yang, Xin Lu, Mohit Bansal, and Tamara L Berg. 2018. Mattnet: Modular attention network for referring expression comprehension. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 1307–1315.
- Haobo Yuan, Xiangtai Li, Tao Zhang, Yueyi Sun, Zilong Huang, Shilin Xu, Shunping Ji, Yunhai Tong, Lu Qi, Jiashi Feng, et al. 2025. Sa2va: Marrying sam2 with llava for dense grounded understanding of images and videos. *arXiv preprint arXiv:2501.04001* (2025).
- Yuxuan Zhang, Tianheng Cheng, Lianghui Zhu, Rui Hu, Lei Liu, Heng Liu, Longjin Ran, Xiaoxin Chen, Wenyu Liu, and Xinggang Wang. 2024. Evf-sam: Early vision-language fusion for text-prompted segment anything model. *arXiv preprint arXiv:2406.20076* (2024).
- Zheng Zhang, Yeyao Ma, Enming Zhang, and Xiang Bai. 2025. Psalm: Pixelwise segmentation with large multi-modal model. In *European Conference on Computer Vision*. Springer, 74–91.