

# Griffon: Spelling out All Object Locations at Any Granularity with Large Language Models

Yufei Zhan<sup>1,2</sup>, Yousong Zhu<sup>1</sup>, Zhiyang Chen<sup>1,2</sup>, Fan Yang<sup>1,3</sup>, Ming Tang<sup>1,2</sup>, and Jinqiao Wang<sup>1,2,3,4</sup>

<sup>1</sup> Foundation Model Research Center, Institute of Automation,  
Chinese Academy of Sciences, Beijing, China

<sup>2</sup> School of Artificial Intelligence,  
University of Chinese Academy of Sciences, Beijing, China

<sup>3</sup> Peng Cheng Laboratory, Shenzhen, China

<sup>4</sup> Wuhan AI Research, Wuhan, China

{zhanyufei2021,yangfan\_2022}@ia.ac.cn

{yousong.zhu,zhiyang.chen,tangm,jqwang}@nlpr.ia.ac.cn


**Abstract.** Replicating the innate human ability to detect all objects based on free-form texts at any granularity remains a formidable challenge for Large Vision Language Models (LVLMs). Current LVLMs are predominantly constrained to locate a single, pre-existing object. This limitation leads to a compromise in model design, necessitating the introduction of visual expert models or customized head structures. Beyond these constraints, our research uncovers LVLMs’ capability for basic object perception, allowing them to accurately identify and locate objects of interest. Building on this insight, we introduce a novel Language-prompted Localization Dataset to fully unleash the capabilities of LVLMs in fine-grained object perception and precise location awareness. More importantly, we present Griffon, a purely LVLm-based baseline, which does not introduce any special tokens, expert models, or additional detection modules. It simply maintains a consistent structure with popular LVLMs by unifying data formats across various localization-related scenarios and is trained end-to-end through a well-designed pipeline. Comprehensive experiments demonstrate that Griffon not only achieves state-of-the-art performance on the fine-grained RefCOCO series and Flickr30K Entities but also approaches the capabilities of the expert model Faster RCNN on the detection benchmark MSCOCO. Data, codes, and models are released at <https://github.com/jefferyZhan/Griffon>.


**Keywords:** Large Vision Language Model · Object Localization · Unified Representation

## 1 Introduction

Large Language Models (LLMs) [44, 62] have shown state-of-the-art performances on a wide range of NLP tasks for capturing more nuanced relationships

---

Corresponding Author: Yousong Zhu.



**Fig. 1: Four Types of Localization-Related Scenarios.** The overall localization task is partitioned into four scenarios based on the number of labels and the number of objects involved. Current LVLMs fail to refuse non-existing objects and detect multiple objects from one or multi-target descriptions (referents, categories, phrases *etc.*).

and contexts and better generalization from limited samples. Subsequently, Large Vision Language Models (LVLMs) [10, 31, 64] have been naturally proposed as a general assistant to employ a LLM as the central processor to understand human intents and finish corresponding tasks. Though superior in various vision understanding tasks [2, 7, 18], they are limited in fine-grained object perception and spatial localization. Additionally, concerns have been raised about the occurrence of hallucinations [30] thought to be attributed to the lack of fine-grained supervision [43].

On the one hand, the limitation primarily stems from the LVLMs’ weak understanding of multiple objects. As depicted in Fig. 1, current popular LVLMs like Shikra [5], Qwen-VL [3], and KOSMOS-2 [36], equipped with basic object perception capabilities, are constrained to localize a single, pre-existing object and struggle in more complex scenarios. Despite increasing the volume to million level [3, 36, 48, 53], the scenario diversity is still limited to the Referring Expression Comprehension (REC). In essence, the crux is not the quantity of data but the lack of comprehensive and well-arranged training benchmarks.

On the other hand, several studies attempt either to invoke offline visual expert models [27, 42, 61] through LLMs or to integrate specific visual detection heads [23, 46, 59] to achieve fine-grained localization. The former is a scheduler-like approach, managing each visual expert model [21, 32] through a specific Application Programming Interface (API). The latter generates specific task outputs through specialized head structures such as the detection head and segmentation head. While these models allow for nuanced predictions, they cannot enjoy a unified structure with LVLMs, increase the parameters exponentially, and fall short in utilizing the knowledge learned by LLMs to generalize to unseen objects. Therefore, the optimal solution lies in enhancing LVLMs with fine-grained perception and localization capabilities intrinsically, without relying on external aids.

In this paper, we introduce a pioneering language-prompted localization benchmark, advancing significantly in bridging the existing data gap in this domain. This benchmark consists of nearly 6M prompted pre-training data and 900K localization instruction-following data constructed with the help of GPT-4V [35]. It covers all four localization-related scenes and over 76K categories as

diverse as possible in open-source datasets. It serves as a crucial foundation for LVLMs to comprehensively enhance their capabilities in simultaneously locating multiple objects in complex scenarios. Leveraging this benchmark, we present **Griffon**, an innovative, purely LVLM-based baseline localizing objects at any granularity according to the input texts. It is characterized by its streamlined architecture and unified input-output representation, which adopts a conversational design without any special token, priors or extra detection heads. The training of **Griffon** is strategically divided into two stages: Stage I focuses on the foundation pre-training and aims to enhance its fine-grained multiple object localization ability, and Stage II involves further instruction tuning for full scenarios, significantly improving the comprehension of users’ intention. Additionally, we introduce a novel, training-free scoring mechanism to rank object outputs, enhancing the model’s ability to prioritize more confident detections.

In order to validate the object perception capability with accurate localization of our framework, we conduct comprehensive experiments on public datasets. **Griffon** achieves the state-of-the-art results on RefCOCO series [33,54] and Flickr30K Entities [37] and the results on par with the expert model on detection dataset MSCOCO [29]. To summarize, the contributions are as follows:

1. We first introduce the Language-prompted Localization Dataset that consists of nearly 6M basic pre-training data and 900K instruction-following data, encompassing all four possible localization-related scenarios and over 76K object categories.
2. We present **Griffon**, a unified LVLM-based baseline without extra structures. It is capable of localizing all objects at any granularity based on free-form input texts.
3. We have conducted comprehensive experiments on REC datasets RefCOCO series, phrase grounding dataset Flickr30K Entities, and object detection dataset MSCOCO 2017. **Griffon** achieves new state-of-the-art results on RefCOCO series and Flickr30K Entities and approaches the performance of the detection expert model.

## 2 Related Work

### 2.1 Localization Models

Traditional localization models usually refer to single-modality object detection or instance segmentation models. According to the backbone type, they can be divided into CNN-based detectors [9, 12, 13, 16, 28, 39, 40] and transformer-based detectors [4, 6, 57, 65], represented respectively by Faster R-CNN [40] and DETR [4]. However, both of them are limited to close-set categories and cannot be generalized to novel categories. With the development of vision-language models [38], open-vocabulary localization methods [14, 56, 63] can locate objects of novel categories. Building on these methods, MDETR [19], GLIP [26, 58] and Grounding DINO [32] further convert localization tasks into a grounded vision-language task incorporating object detection, phrase grounding, and Referring

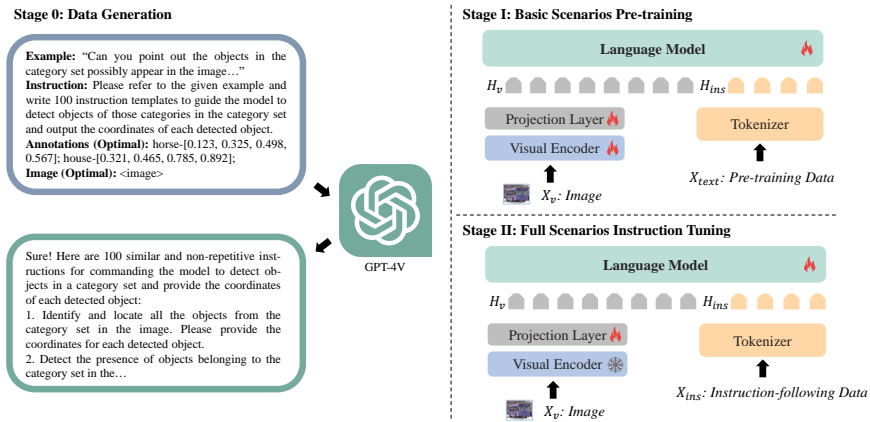
Expression Comprehension tasks into a single open-set model. Unlike the above methods, our method makes predictions in an auto-regression manner without relying on the detection head and generalizes to different granularity localization tasks and object categories based on the LLM instead of the alignment between the region proposal and the text.

## 2.2 Large Vision Language Models

The evolution of LLMs [8, 44, 60] to accommodate multimodal inputs is a burgeoning field of interest. To tune interleaved image-text data, Flaggingo [1] designed a visual perceiver that incorporates the visual information into the adapter layers of the LLM. On the contrary, BLIP-2 [24] aligns visual embeddings with textual embeddings with a Q-Former. A fully connected layer is subsequently used to feed the projected visual embeddings to LLM. Following the inclusion of instruction tuning from NLP [17, 47], models like Mini-GPT4 [64], mPLUG-OWL [52], InstructBLIP [10] have been developed. They preserve the Q-Former structure and improve the interface performance with curated instruction-following data. LLaVA [31] differs by projecting visual encoder outputs to the LLM through a trainable fully connected layer, optimized with the LLaVA-150K instruction dataset, which is commonly used for instruction tuning. More evidence of LVLMS’ remarkable visual language understanding powers and task-specific adaptability comes from the unveiling of GPT-4V [35]. Some LVLMS [3, 5, 36, 48, 53] recently have introduced REC data into their training dataset, endowing them with basic object perception skills. Nonetheless, the precision in object localization and adaptability to complex scenarios remains a hurdle. To overcome these challenges, we present an exhaustive training benchmark and formulate a robust baseline purely based on LVLMS.

## 2.3 Localization-equipped LVLMS

To rival the detection performance of specialist models, the integration of visual expert models has been a prevalent tactic. HuggingGPT [42] enhances GPT’s capabilities by connecting with a suite of visual expert models through APIs, allowing it to trigger a detection model upon recognizing an object detection prompt. AutoGPT [50] goes beyond by serving as an agent that can execute the post-processing tasks after detection. Further researches [27, 34, 61] have expanded the range of APIs and refined their proficiency in specific domains. Another approach enriches the LVLMS framework with dedicated detection heads [23, 46, 49, 59]. Accurate detection or segmentation results are produced by certain approaches [23, 49] depending on the LLM predictions. In contrast, models such as VisionLLM [46] and GPT4RoI [59] feed the output of the detection head to the LLM to enhance the detection performance or achieve particular goals. Unlike the above approaches, our methodology does not use any external specialized heads or expert models for detection. Instead, we enhance our model’s fine-grained object perception and spatial localization capabilities through our meticulously crafted dataset and unified representation.



**Fig. 2: Data Generation and Training Procedure.** Griffon follows a progressive two-stage training pipeline with the built dataset in stage 0. In different stages, distinct modules of **Griffon** are trained. The red flame represents that this module is being trained at this stage, while the gray snowflake indicates the opposite.

### 3 Methodology

In this section, we start with our constructed novel Language-prompted Localization Dataset with scenarios analysis and detailed description of data construction in Sec. 3.1. Leveraging this dataset, we present **Griffon**, a purely LVLm-based localization baseline. We detail the unified conversational input-output representation under a streamlined structure (Sec. 3.2). And we further detail our two-stage training pipeline in Sec. 3.3. Finally, for the optimization of multi-object scenarios, we present the training-free scoring mechanism in Sec. 3.4.

#### 3.1 Language-prompted Localization Dataset

**Scenarios Analysis.** In the real world, a user may label the object with texts at different granularity. This granularity of texts can be clarified as fine-grained referents and coarse-grained categories. Current LVLms [5, 36] focus on the single referent localization and neglect the much more challenging multiple objects localization. Besides, there are scenarios where the target might not be present in the image. Hence, based on the quantity and type of labels specified, as depicted in Fig. 1, language-prompted localization in the real-world setting can be categorized into four distinct scenarios:

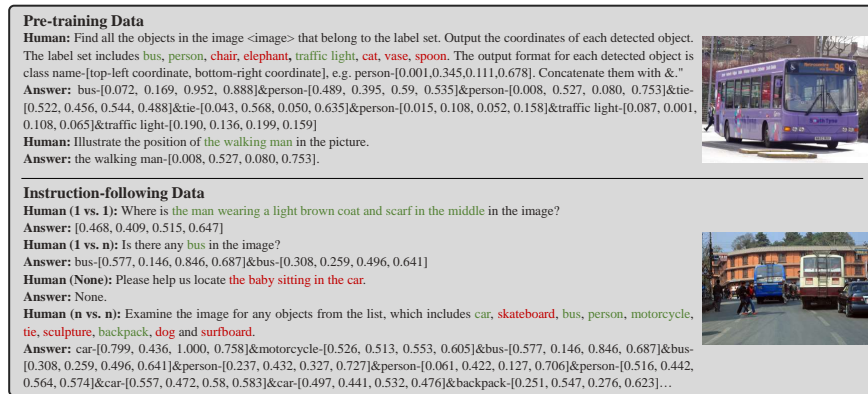
- Single Referent Localization (1 vs. 1): Models are required to distinguish the target from others and output the precise coordinates of that specific object.
- One Category with Multi Objects (1 vs. n): Upon receiving a category name or a descriptive phrase, the model locates and outputs the coordinates of all matching objects.

- An Non-existing Referent (None): If the input includes descriptions of non-existent objects, the model indicates this by returning a 'None' output.
- Multi Categories with Multi Objects ( $n$  vs.  $n$ ): When inputs comprise multiple descriptions, varying in both granularity and existence, the model identifies and outputs the coordinates of all discernible objects.

**Data Construction.** Previous datasets [10, 31] employed by LVLMs are composed of two components, including pre-training data and instruction-following data. Datasets of different visual language tasks are merged into the pre-training data and then sampled and refined to build the instruction-following data. As for language-prompted localization, it's primarily related to the Referring Expression Comprehension (REC) above. Therefore, following the above dataset construction paradigm, we respectively build base pre-training data and instruction-following data for the Language-prompted Localization Dataset. We describe the main procedure as follows and details can be found in the Supplements.

*Pre-training Data.* We first collect the public datasets of the REC task and object detection task such as Visual Genome [22] and Objects365 [41], which help with fine-grained discrimination and multiple object perception. After filtering images whose longest edge is less than 250 and whose image size is different from the size in the annotation, we totally collect 6M image-text samples, including nearly 4M in REC and 2M in object detection, respectively. As pure visual tasks, the annotations of these datasets only contain the labels without task prompts. Hence, we generate corresponding task templates to format the inputs with the help of GPT-4V [35]. As shown in the left part of Fig. 2, we first prompt GPT-4V with a template example, then ask GPT-4V to generate a large number of similar templates based on the requirements. As exemplified in Fig. 3, each task template provides an interface, which can be flexibly replaced with other labels or phrases according to the ground truth during training. We craft 60 templates for each task, which can be randomly sampled.

*Instruction-following Data.* We construct the instruction data on the basis of the pre-training data according to the concluded four scenarios. Additionally, we incorporate the Flickr30K Entities [37] for the One Category with Multi Objects scenario and the LVIS dataset [15] for its "negative categories" annotations, which aims at building instruction samples of the Non-Existing Referent scenario. On one hand, we further generate more diverse instruction templates closer to the human daily expression with the help of GPT-4V and integrate the label with a randomly selected template to build instruction-image samples of different scenarios. On the other hand, for the case of Non-Existing Referent, we further mine negatives with a more fine-grained phrase using GPT-4V, adding extra attributes or descriptions, like "a white cat" or "people wearing a scar". It can deal with scenarios where the object exists but is described with wrong attributes. We collect nearly 900K unique image-text localization instruction samples in total and provide some examples of the final data in Fig. 3.



**Fig. 3: Samples of Language-Prompted Localization Dataset.** All the images of the benchmark are collected and filtered from public datasets. Instructions are generated with the proposed method using GPT-4V [35]. The red indicates that this object does not exist in the image, while the green indicates the opposite.

### 3.2 Unified Representation With LVLMS

Locating all objects at any granularity through Large Language Models (LLMs) still remains to be explored in LVLMS. Thus, we develop **Griffon** based on the LLaVA [31] and propose a unified input-output representation. It generates predictions requiring only a single forward pass for multi-label targets without priors, structure modifications, and special tokens.

LLaVA is composed of three modules depicted in Fig. 2, including the visual encoder, projection layer, and LLM. We utilize the ViT-L/14 of CLIP [38] and LLaMA2-13B [44] as our visual encoder and LLM, respectively, while the projection layer is a lightweight fully connection layer. The input image  $X_v$  is first encoded to the visual feature  $Z_v$  and then projected to word embedding space as  $H_v$  through the projection layer. It shares the same dimension with the word embedding space and is concatenated with language embedding tokens  $H_{ins}$ , which are fed to the LLM to generate outputs.

Adopting the conversational style, the inputs of LLaVA are free-form texts ranging from commands to questions. Inheriting from the LLaVA, **Griffon** can also accept free-form inputs and further unifies inputs of different scenarios naturally. However, as LLaVA is aimed at visual language understanding tasks, it cannot handle the object coordinates exactly. Also, previous methods [5, 36] can only output the textual coordinates of the single detected object and require repeated forward passes for multiple referents. Therefore, we present a unified output representation for multiple objects. For a detected object, it is represented with the top-left and bottom-right coordinates along with its label, formatted as “label-[x1, y1, x2, y2]”. This can be adapted to multi-object scenarios seamlessly by concatenating all detected objects with “&”, without introducing connection placeholders [3].

As for the representation of coordinates, there are mainly two types: extra special representation which creates new tokens to represent each coordinate in the axis, and numerical characters representation which directly uses numerals in natural languages. Shikra [5] conducts a comparison experiment on these two representations on REC task and the conclusion is the numerical approach can achieve better performance. Also, as the embeddings of the extra tokens out of the original vocabulary need to be trained from scratch, it results in slower convergence. Hence, we employ numerical characters for normalized coordinates and set the precision to 0.001. This precision strikes a balance between reducing the localization error for small objects (area less than  $32 \times 32$  pixels) and the length of the sequence.

### 3.3 Progressive Two-Stage Training Pipeline

Previous LVLMs equipped with localization capability add REC data to the vision and language alignment phrase and then are finetuned with designed localization-related instructions. By increasing the volume of data, this approach is feasible for tasks like REC that only output coordinates of the single object. However, when handling multiple objects of different granularities, requires the model to have a foundation understanding of the scene and the context of the entire image. Based on this insight and borrowing from the pre-training and instruction tuning paradigm from LLM, we propose the progressive two-stage training pipeline.

**Stage I: Basic Scenarios Pre-training.** The first stage is designed to enhance the multi-object perception capability and achieve localization in basic scenarios. We start by initialization with the weights from LLaVA to equip **Griffon** with the vision-language alignment capability and the instruction comprehension capability. Furthermore, we use the whole 6M pre-training data in the Language-prompted Localization Dataset for pre-training to create a foundation model that can accurately locate all objects in an image. To construct the input language  $X_{ins}$ , for each image  $X_v$ , a template will be sampled randomly from the all templates generated and integrated with the label for better generalization. Instead of only training the projector, we train the entire network. In this way, with fine-grained supervision, the details will be kept in the visual features, beneficial for small targets and fine-grained discrimination.

**Stage II: Full Scenarios Instruction Tuning.** To accustom to the full scenario, we conduct instruction tuning on the foundation model developed in the stage II. To be specific, following the first pre-training stage, the model has already developed the ability to localize objects based on free-form texts at any granularity. However, it fails to comprehend the instructions from specific scenarios and user intentions. Therefore, we further finetune the LLM and projection layer of **Griffon** with the 900K instruction data covering all four scenarios to refine the model’s capabilities of comprehending the user’s intention.



### 3.4 Training-Free Confidence Scoring Mechanism

In the object detection, expert models often predict an additional confidence score for each detected object to highlight high-quality bounding boxes. This approach mirrors the human cognitive process, which prioritizes attention towards objects of confidence based on existing knowledge. As a general optimization strategy, we present a training-free multi-object confidence scoring mechanism for localization with LVLMs. It generates an output sequence  $X_a$  of length  $L_a$  in an auto-regressive way, and the probability is computed by:

$$p(X_a|X_v, X_{ins}) = \prod_{j=1}^{L_a} p_{\theta}(x_{a,j}|X_v, X_{ins}, X_{a,<j}) \quad (1)$$

where  $\theta$  is the model parameters,  $X_{ins}, X_{a,<j}$  are the instructions and answer tokens before the current prediction token  $x_{a,j}$ , respectively. At each step of the generation phrase until the [EOS] token, we will get the conditional probability  $p_{\theta}(x_{a,j}|X_v, X_{ins}, X_{a,<j})$  of  $x_{a,j}$ . For each predicted object label  $R$ , it consists of  $L_r$  tokens. Supposing this label starts at  $k^{th}$  token in the sequence, the conditional probability of generating this label can be computed according to Bayes' theorem by:

$$p(R|X_v, X_{ins}, x_{a,<k}) = \prod_{j=k}^{k+L_r} p_{\theta}(x_{a,j}|X_v, X_{ins}, X_{a,<j}) \quad (2)$$

where  $X_v, X_{ins}$  are the visual and instruction inputs, and  $x_{a,<k}$  are the answer tokens up to the  $k^{th}$  token. This applies to the coordinates as well. With the conditional probability of each coordinate, we compute the localization confidence score for a bounding box starting from the  $k^{th}$  token by:

$$p_{loc} = \prod_{coord \in \{x1,y1,x2,y2\}} p(coord|X_v, X_{ins}, X_{a,<k}), \quad (3)$$

where  $x1, y2, x2, y2$  are the top-left and bottom coordinates, respectively. The final confidence score for a detected object is the geometric mean of the label score and localization score, where  $q = 0.5$ :

$$score = p(R|X_v, X_{ins}, X_{a,<j})^q \times p_{loc}^{1-q}. \quad (4)$$

## 4 Experiments

In this section, we first introduce the key implementation details briefly. Based on the implementation, We conduct quantitative experiments to evaluate the fine-grained object perception capability with spatial localization in three localization-related scenarios separately, including Single Referent, One Category with Multi Objects, and Multi Categories with Multi Objects scenarios. They correspond to REC, phrase grounding, and object detection tasks respectively. After that, ablation studies are provided to better comprehend the detailed options of **Griffon**. A qualitative analysis of all four localization-related scenarios is provided in the last.

#### 4.1 Implementation Details

To strike a balance between the training efficiency and small objects performance, we increase the image input size to 448, applying the bilinear interpolation to the position embedding of the visual encoder. After the initialization, we adopt the burn-in strategy to warm up **Griffon** with a subset of REC data for 1 epoch. The learning rate is set to 2e-5 with a global batch size of 128. Then, **Griffon** is trained according to the designed pipeline. With the increase of data in the stage I, we align the batch size with the available computation resources, set to 256 with the learning rate linear scaled to 4e-5. In stage II, we decrease the batch size to 128 together with the learning rate to 2e-5. More details can be found in the Supplements.

**Table 1:** Referring Expression Comprehension results with  $ACC@0.5$ . \* indicates the result trained with only REC grounding data for a fair comparison with the above models.

Type	Model	Data Vol.	Size	RefCOCO			RefCOCO+			RefCOCOg	
				val	test-A	test-B	val	test-A	test-B	val	test
Specialists	MDETR [19]			87.5	90.4	82.7	81.1	85.5	73.0	83.3	83.3
	TransVG [11]			81.0	82.7	78.4	64.8	70.7	56.9	68.7	67.7
	G-DINO-L [32]	-	-	90.6	93.2	88.2	82.8	89.0	75.9	86.1	87.0
	UNINEXT-L [49]			91.4	93.7	88.9	83.1	87.9	76.2	86.9	87.5
Generalists	OFA-L [45]	10M	512	80.0	83.7	76.4	68.3	76.0	61.8	67.8	67.5
	Qwen-VL [3]	77M	448	89.4	92.3	85.3	83.1	88.3	77.2	85.6	85.5
	PINK [48]	5M	224	88.7	92.1	84.0	81.8	88.2	73.9	83.9	84.3
	Ferret-13B [53]	8.7M	336	89.5	92.4	84.4	82.8	88.1	75.2	85.8	86.3
	Shikra-13B [5]	4M	224	87.8	90.6	80.2	82.9	87.8	74.4	82.6	83.2
	<b>Griffon-13B</b>	7M	448	89.4	92.5	84.6	83.3	88.4	76.0	85.1	86.1
	<b>Griffon-13B*</b>	4M	448	<b>90.1</b>	<b>93.4</b>	<b>86.1</b>	<b>84.8</b>	<b>90.5</b>	<b>77.8</b>	<b>86.1</b>	<b>87.2</b>

#### 4.2 Experiments on Referring Expression Comprehension

Referring Expression Comprehension is aimed at the Single Referent scenario, reflecting the fine-grained determination capability with spatial localization. In Tab. 1, we compare **Griffon** with a group of popular frameworks on RefCOCO [54], RefCOCO+ [54], and RefCOCOg [33]. We mainly compare **Griffon** with the generalists based on LVLMs, because they share a similar structure and training approach. As these generalists are different in the scale of LLMs, input resolution, and the volume of data, we list the characteristics for comparison. Meanwhile, we also present the results of some experts to demonstrate the level **Griffon** has reached. As shown in Tab. 1, **Griffon** surpasses the state-of-the-art model Qwen-VL under the same input resolution and with quite less training data, and also outperforms other models utilizing a similar amount of data on all

metrics by a large margin. For fair comparison of the training pipeline we build, we also follow the previous models to only employ single referent data during pre-training. This makes the model focus on the dominant single referent localization, though the expression annotated might correspond to multiple objects. It is the crux of the previous high performance on REC tasks but failure in all multi-object scenarios for current generalists. In this setting, our model also achieves state-of-the-art results, proving the effectiveness of our pipeline.

### 4.3 Experiments on Phrase Grounding

Besides the Single Referent scenario, we present quantitative evaluations for the harder One Category with Multiple Objects scenario on the phrase grounding task. Existing LVLMs are typically limited to the Single Referent scenario, predicting only one box per phrase, often missing out on multiple instances. In this way, they are only evaluated with the MERGED-BOXES protocol ignoring the single instance in the group the phrase refers to. **Griffon** excels in multi-object scenarios, seamlessly localizing both fine-grained and coarse-grained phrases. We first evaluate this scenario on Flickr30K Entities [37] with both ANY-BOX protocol and MERGED-BOXES protocol. As depicted in the Tab. 2, **Griffon** achieves state-of-the-art on Flickr30K Entities, validating its success in multiple objects scenarios.

**Table 2:** Phrase grounding results on Flickr30K Entities [37]. ANY represents the ANY-BOX protocol, which focuses on the atomicity of each instance, while MERGED represents the MERGED-BOXES protocol evaluating whether the model finds all the referred objects with a merged box.

Type	Model	ANY		MERGED	
		val	test	val	test
Specialists	BAN [20]	-	67.9	-	-
	DDPN [55]	-	-	72.8	73.5
	VisualBert [25]	70.4	71.3	-	-
	MDETR [19]	82.5	83.4	82.3	83.8
Generalists	UniTAB [51]	-	-	78.8	79.6
	Ferret-13B [53]	-	-	81.1	<b>84.8</b>
	Shikra-13B [5]	-	-	77.4	78.4
	<b>Griffon-13B</b>	<b>83.7</b>	<b>84.2</b>	<b>82.0</b>	82.8

### 4.4 Experiments on Object Detection

Object detection exactly integrates all four scenarios into one task as there exist images containing one or no existing category for some images, and might

be the hardest task. It can reflect the capabilities of dealing with complex scenarios. We conduct experiments on the widely-used object detection benchmark MSCOCO [29]. Since the existing popular LVLM methods cannot locate multiple objects of specified categories once, we evaluate Qwen-VL and Ferret on object detection tasks following the model design to ask one category at a time and iterate through all categories for each image. However, the  $mAP$  and  $AP_S$  of them are approximately 0, and the  $AR_{100}$  of Qwen-VL is 3.7, which also validates the limitation of LVLMs we depict in the Sec. 1 and Fig. 1. We provide some visualizations of the test results in the Supplements for further illustration. Therefore, we don't include their results in the table and mainly compare our model with some basic expert models in Tab. 3. As illustrated in Tab. 3, **Griffon** is first capable of detecting multiple objects of various categories instead of 0 precision of existing LVLMs. It approaches the performance of the expert model Faster RCNN and outperforms Faster RCNN in  $AP_S$  under the same resolution. With the above results, it can be concluded that **Griffon** is better than the previous generalist models with stronger localization and perception capability in both multi-object and single-object scenarios.

**Table 3:** Object detection results on MSCOCO val2017 [29]. Griffon is the only LVLM capable of object detection, *i.e.* detecting multiple objects of specified multiple categories once.

Type	Model	Size	Epochs	$mAP$	$AP_{50}$	$AP_{75}$	$AP_S$	$AP_M$	$AP_L$
Specialists	Pix2Seq-R50 [6]	1333	300	43.0	61.0	45.6	25.1	46.9	59.4
	FRCNN-R50 [40]	448	12	26.3	42.1	27.5	4.6	27.7	49.9
	DETR-DC5 [57]	1333	12	15.5	29.4	14.5	4.3	15.1	26.9
Generalists	<b>Griffon-13B</b>	448	1	24.8	40.6	25.1	5.9	25.5	48.7

#### 4.5 Ablation Study

In this section, we ablate on two key designs in **Griffon**. We conduct these experiments on MSCOCO val2017 [29] for demonstration.

**Unified Representation Order.** In Sec. 3.2, we introduce the proposed unified representation which bridges the gap between single-referent scenarios and comprehensive scenarios with multiple objects for LVLMs. As **Griffon** generates outputs in the sequence-to-sequence manner, the relative positions of the coordinates and labels would affect the performance, especially for the multi-objects prediction. To compare the effectiveness of these two formats, we directly train the model with the pre-training data for 1 epoch for fast validation and evaluate them on the MSCOCO dataset. As shown in Tab. 4, the label-first order surpasses the coord-first order on all metrics, with the mAP improved by +4.4%.

We believe the difference is caused by placing the label first can provide more uncertainty, as it corresponds to multiple different objects, to increase the number of predictions and further improve the recall rate. Therefore, we finally utilize the label-first format in **Griffon**.

**Table 4:** Ablation study of unified representation order. The Coord-First order is “[x1,y1,x2,y2]-label”, and the Label-First is the opposite.

Order	mAP	$AP_{50}$	$AP_{75}$	$AR_{100}$
Coord-first	17.8	28.9	17.9	30.5
Label-first	22.2	35.4	22.6	33.0


**Confidence Scoring Mechanism.** To optimize the multi-object predictions, we present the training-free confidence scoring mechanism. It produces the confidence of this prediction from both localization and recognition aspects with calculated conditional probability, detailed in Sec. 3.4. We validate the effectiveness of the proposed scoring mechanism in Tab. 5. The localization score can improve the multi-objects performance by +1.3% mAP compared with the baseline. Composed of these two scores, the final score improves the mAP by 2.3% overall. The result is quite evident because if removing the score, the evaluation degrades into a random process leading to mismatches. With the confidence score, the prediction with the higher score, *i.e.* higher quality, will be matched with the ground truth. Therefore, it’s better to utilize the training-free confidence scoring mechanism during inference.

**Table 5:** Ablation study of Training-free Scoring Mechanism.

loc	label	mAP	$AP_{50}$	$AP_{75}$	$AR_{100}$
×	×	20.9	35.0	20.7	33.9
✓	×	22.2	35.4	22.6	33.0
✓	✓	23.2	37.7	23.4	33.9

#### 4.6 Qualitative Analysis

To gain more insights about **Griffon**, we present extensive visualization results in Fig. 4. We qualitatively validate the effectiveness of **Griffon** by comparing it with the expert model Grounding DINO and the generalist model Qwen-VL in all possible localization scenarios. As shown in Fig. 4, for Single Referent scenario, **Griffon** can not only localize objects of normal sizes but also small



**Fig. 4:** Visualization results of Qwen-VL [3], Grounding DINO [32] and **Griffon** across all four scenarios. We use abbreviations to substitute the four scenarios.

objects compared with Qwen-VL. Also, **Griffon** can refuse to output results when the input object is nonexistent and doesn't omit any objects in the One Category with Multi Objects scenario. It further approaches the performance of the expert model Grounding DINO in Multi Categories with Multi Objects scenario in the case of lower resolution (448 *vs.* 1333), accurately localizing the small objects and objects partially hidden. Therefore, our **Griffon** is capable of being applied in real-world scenarios requiring accurate localization.

## 5 Conclusion

In this paper, we present the first Language-prompted Localization Dataset across all possible localization-related scenarios. Leveraging this dataset, we establish a pure LVLM-based localization baseline, **Griffon**. It is capable of localizing objects based on the input texts at any granularity, without introducing any priors, detection modules, or special tokens. **Griffon** achieves the state-of-the-art results on the REC task and phrase grounding task and approaches the performance of visual expert model Faster RCNN on the detection benchmark. These results prove that a single LVLM can also do well in fine-grained visual perception tasks without external aids. We hope and believe our attempts at unified encoding for localization and the promising results achieved in object detection will facilitate the deep integration of visual tasks and vision-language tasks, offering potential utility to the broader community. As we mainly focus on localization and conduct experiments only in this domain, which might be a limitation, integrating more tasks into **Griffon** is our future direction.

**Acknowledgement.** This work was supported by National Science and Technology Major Project (2021ZD0114600), National Natural Science Foundation of China (No.62276260, No.62176254) and in part by the Beijing Natural Science Foundation under Grant No.4244099.

## References

1. Alayrac, J.B., Donahue, J., Luc, P., Miech, A., Barr, I., Hasson, Y., Lenc, K., Mensch, A., Millican, K., Reynolds, M., et al.: Flamingo: a visual language model for few-shot learning. *Advances in Neural Information Processing Systems* **35**, 23716–23736 (2022)
2. Antol, S., Agrawal, A., Lu, J., Mitchell, M., Batra, D., Zitnick, C.L., Parikh, D.: Vqa: Visual question answering. In: *Proceedings of the IEEE international conference on computer vision*. pp. 2425–2433 (2015)
3. Bai, J., et al: Qwen-vl: A versatile vision-language model for understanding, localization, text reading, and beyond. *arXiv:2308.12966* (2023)
4. Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: End-to-end object detection with transformers. In: *European conference on computer vision*. pp. 213–229. Springer (2020)
5. Chen, K., et al: Shikra:unleashing multimodal llm’s referential dialogue magic. *arXiv:2306.15195* (2023)
6. Chen, T., Saxena, S., Li, L., Fleet, D.J., Hinton, G.: Pix2seq: A language modeling framework for object detection. *arXiv preprint arXiv:2109.10852* (2021)
7. Chen, X., Fang, H., Lin, T.Y., Vedantam, R., Gupta, S., Dollár, P., Zitnick, C.L.: Microsoft coco captions: Data collection and evaluation server. *arXiv preprint arXiv:1504.00325* (2015)
8. Chowdhery, A., Narang, S., Devlin, J., Bosma, M., Mishra, G., Roberts, A., Barham, P., Chung, H.W., Sutton, C., Gehrmann, S., et al.: Palm: Scaling language modeling with pathways. *arXiv preprint arXiv:2204.02311* (2022)
9. Dai, J., Li, Y., He, K., Sun, J.: R-fcn: Object detection via region-based fully convolutional networks. *Advances in neural information processing systems* **29** (2016)
10. Dai, W., Li, J., Li, D., Tiong, A.M.H., Zhao, J., Wang, W., Li, B., Fung, P., Hoi, S.: Instructblip: Towards general-purpose vision-language models with instruction tuning (2023)
11. Deng, J., Yang, Z., Chen, T., Zhou, W., Li, H.: Transvg: End-to-end visual grounding with transformers. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 1769–1779 (2021)
12. Duan, K., Bai, S., Xie, L., Qi, H., Huang, Q., Tian, Q.: Centernet: Keypoint triplets for object detection. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 6569–6578 (2019)
13. Girshick, R.: Fast r-cnn. In: *Proceedings of the IEEE international conference on computer vision*. pp. 1440–1448 (2015)
14. Gu, X., Lin, T.Y., Kuo, W., Cui, Y.: Open-vocabulary object detection via vision and language knowledge distillation. *arXiv preprint arXiv:2104.13921* (2021)
15. Gupta, A., Dollar, P., Girshick, R.: Lvis: A dataset for large vocabulary instance segmentation. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 5356–5364 (2019)
16. He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask r-cnn. In: *Proceedings of the IEEE international conference on computer vision*. pp. 2961–2969 (2017)
17. Honovich, O., Scialom, T., Levy, O., Schick, T.: Unnatural instructions: Tuning language models with (almost) no human labor. *arXiv preprint arXiv:2212.09689* (2022)
18. Johnson, J., Hariharan, B., Van Der Maaten, L., Fei-Fei, L., Lawrence Zitnick, C., Girshick, R.: Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 2901–2910 (2017)

19. Kamath, A., Singh, M., LeCun, Y., Misra, I., Synnaeve, G., Carion, N.: Mdetr—modulated detection for end-to-end multi-modal understanding. arXiv preprint arXiv:2104.12763 (2021)
20. Kim, J.H., Jun, J., Zhang, B.T.: Bilinear attention networks. *Advances in neural information processing systems* **31** (2018)
21. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., et al.: Segment anything. arXiv preprint arXiv:2304.02643 (2023)
22. Krishna, R., Zhu, Y., Groth, O., Johnson, J., Hata, K., Kravitz, J., Chen, S., Kalantidis, Y., Li, L.J., Shamma, D.A., et al.: Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International journal of computer vision* **123**, 32–73 (2017)
23. Lai, X., Tian, Z., Chen, Y., Li, Y., Yuan, Y., Liu, S., Jia, J.: Lisa: Reasoning segmentation via large language model. arXiv preprint arXiv:2308.00692 (2023)
24. Li, J., Li, D., Savarese, S., Hoi, S.: BLIP-2: bootstrapping language-image pre-training with frozen image encoders and large language models. In: *ICML (2023)*
25. Li, L.H., Yatskar, M., Yin, D., Hsieh, C.J., Chang, K.W.: Visualbert: A simple and performant baseline for vision and language. arXiv preprint arXiv:1908.03557 (2019)
26. Li\*, L.H., Zhang\*, P., Zhang\*, H., Yang, J., Li, C., Zhong, Y., Wang, L., Yuan, L., Zhang, L., Hwang, J.N., Chang, K.W., Gao, J.: Grounded language-image pre-training. In: *CVPR (2022)*
27. Liang, Y., Wu, C., Song, T., Wu, W., Xia, Y., Liu, Y., Ou, Y., Lu, S., Ji, L., Mao, S., et al.: Taskmatrix. ai: Completing tasks by connecting foundation models with millions of apis. arXiv preprint arXiv:2303.16434 (2023)
28. Lin, T.Y., Dollár, P., Girshick, R., He, K., Hariharan, B., Belongie, S.: Feature pyramid networks for object detection. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 2117–2125 (2017)
29. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13*. pp. 740–755. Springer (2014)
30. Liu, F., Lin, K., Li, L., Wang, J., Yacoob, Y., Wang, L.: Mitigating hallucination in large multi-modal models via robust instruction tuning (2023)
31. Liu, H., et al.: Visual instruction tuning. In: *NeurIPS (2023)*
32. Liu, S., Zeng, Z., Ren, T., Li, F., Zhang, H., Yang, J., Li, C., Yang, J., Su, H., Zhu, J., et al.: Grounding dino: Marrying dino with grounded pre-training for open-set object detection. arXiv preprint arXiv:2303.05499 (2023)
33. Nagaraja, V.K., Morariu, V.I., Davis, L.S.: Modeling context between objects for referring expression understanding. In: *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part IV 14*. pp. 792–807. Springer (2016)
34. Nakano, R., Hilton, J., Balaji, S., Wu, J., Ouyang, L., Kim, C., Hesse, C., Jain, S., Kosaraju, V., Saunders, W., et al.: Webgpt: Browser-assisted question-answering with human feedback. arXiv preprint arXiv:2112.09332 (2021)
35. OpenAI: Gpt-4 technical report (2023)
36. Peng, Z., et al.: Kosmos-2: Grounding multimodal large language models to the world. In: *ICLR (2024)*
37. Plummer, B.A., Wang, L., Cervantes, C.M., Caicedo, J.C., Hockenmaier, J., Lazebnik, S.: Flickr30k entities: Collecting region-to-phrase correspondences for richer



- image-to-sentence models. In: Proceedings of the IEEE international conference on computer vision. pp. 2641–2649 (2015)
38. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: International conference on machine learning. pp. 8748–8763. PMLR (2021)
  39. Redmon, J., Divvala, S., Girshick, R., Farhadi, A.: You only look once: Unified, real-time object detection. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 779–788 (2016)
  40. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems* **28** (2015)
  41. Shao, S., Li, Z., Zhang, T., Peng, C., Yu, G., Zhang, X., Li, J., Sun, J.: Objects365: A large-scale, high-quality dataset for object detection. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 8430–8439 (2019)
  42. Shen, Y., Song, K., Tan, X., Li, D., Lu, W., Zhuang, Y.: Hugginggpt: Solving ai tasks with chatgpt and its friends in huggingface. *arXiv preprint arXiv:2303.17580* (2023)
  43. Shukang, Y., Chaoyou, F., Sirui, Z., Tong, X., Hao, W., Dianbo, S., Yunhang, S., Ke, L., Xing, S., Enhong, C.: Woodpecker: Hallucination correction for multimodal large language models. *arXiv preprint arXiv:2310.16045* (2023)
  44. Touvron, H., Martin, L., Stone, K., Albert, P., Almahairi, A., Babaei, Y., Bashlykov, N., Batra, S., Bhargava, P., Bhosale, S., et al.: Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288* (2023)
  45. Wang, P., Yang, A., Men, R., Lin, J., Bai, S., Li, Z., Ma, J., Zhou, C., Zhou, J., Yang, H.: Ofa: Unifying architectures, tasks, and modalities through a simple sequence-to-sequence learning framework. In: International Conference on Machine Learning. pp. 23318–23340. PMLR (2022)
  46. Wang, W., Chen, Z., Chen, X., Wu, J., Zhu, X., Zeng, G., Luo, P., Lu, T., Zhou, J., Qiao, Y., et al.: Visionllm: Large language model is also an open-ended decoder for vision-centric tasks. *arXiv preprint arXiv:2305.11175* (2023)
  47. Wang, Y., Kordi, Y., Mishra, S., Liu, A., Smith, N.A., Khashabi, D., Hajishirzi, H.: Self-instruct: Aligning language model with self generated instructions. *arXiv preprint arXiv:2212.10560* (2022)
  48. Xuan, S., Guo, Q., Yang, M., Zhang, S.: Pink: Unveiling the power of referential comprehension for multi-modal llms (2023)
  49. Yan, B., Jiang, Y., Wu, J., Wang, D., Luo, P., Yuan, Z., Lu, H.: Universal instance perception as object discovery and retrieval. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 15325–15336 (2023)
  50. Yang, H., Yue, S., He, Y.: Auto-gpt for online decision making: Benchmarks and additional opinions. *arXiv preprint arXiv:2306.02224* (2023)
  51. Yang, Z., Gan, Z., Wang, J., Hu, X., Ahmed, F., Liu, Z., Lu, Y., Wang, L.: Unitab: Unifying text and box outputs for grounded vision-language modeling. In: European Conference on Computer Vision. pp. 521–539. Springer (2022)
  52. Ye, Q., Xu, H., Xu, G., Ye, J., Yan, M., Zhou, Y., Wang, J., Hu, A., Shi, P., Shi, Y., Jiang, C., Li, C., Xu, Y., Chen, H., Tian, J., Qi, Q., Zhang, J., Huang, F.: mplug-owl: Modularization empowers large language models with multimodality (2023)
  53. You, H., et al: Ferret: Refer and ground anything anywhere at any granularity. In: ICLR (2024)

54. Yu, L., Poirson, P., Yang, S., Berg, A.C., Berg, T.L.: Modeling context in referring expressions. In: *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part II* 14. pp. 69–85. Springer (2016)
55. Yu, Z., Yu, J., Xiang, C., Zhao, Z., Tian, Q., Tao, D.: Rethinking diversified and discriminative proposal generation for visual grounding. *arXiv preprint arXiv:1805.03508* (2018)
56. Zareian, A., Rosa, K.D., Hu, D.H., Chang, S.F.: Open-vocabulary object detection using captions. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 14393–14402 (2021)
57. Zhang, H., Li, F., Liu, S., Zhang, L., Su, H., Zhu, J., Ni, L.M., Shum, H.Y.: Dino: Detr with improved denoising anchor boxes for end-to-end object detection. *arXiv preprint arXiv:2203.03605* (2022)
58. Zhang, H., Zhang, P., Hu, X., Chen, Y.C., Li, L.H., Dai, X., Wang, L., Yuan, L., Hwang, J.N., Gao, J.: Glipv2: Unifying localization and vision-language understanding. *arXiv preprint arXiv:2206.05836* (2022)
59. Zhang, S., Sun, P., Chen, S., Xiao, M., Shao, W., Zhang, W., Chen, K., Luo, P.: Gpt4roi: Instruction tuning large language model on region-of-interest. *arXiv preprint arXiv:2307.03601* (2023)
60. Zhang, S., Roller, S., Goyal, N., Artetxe, M., Chen, M., Chen, S., Dewan, C., Diab, M., Li, X., Lin, X.V., et al.: Opt: Open pre-trained transformer language models. *arXiv preprint arXiv:2205.01068* (2022)
61. Zhao, Y., Lin, Z., Zhou, D., Huang, Z., Feng, J., Kang, B.: Bubogpt: Enabling visual grounding in multi-modal llms. *arXiv preprint arXiv:2307.08581* (2023)
62. Zheng, L., Chiang, W.L., Sheng, Y., Zhuang, S., Wu, Z., Zhuang, Y., Lin, Z., Li, Z., Li, D., Xing, E.P., Zhang, H., Gonzalez, J.E., Stoica, I.: Judging llm-as-a-judge with mt-bench and chatbot arena (2023)
63. Zhou, X., Girdhar, R., Joulin, A., Krähenbühl, P., Misra, I.: Detecting twenty-thousand classes using image-level supervision. In: *European Conference on Computer Vision*. pp. 350–368. Springer (2022)
64. Zhu, D., Chen, J., Shen, X., Li, X., Elhoseiny, M.: Minigpt-4: Enhancing vision-language understanding with advanced large language models. *arXiv preprint arXiv:2304.10592* (2023)
65. Zhu, X., Su, W., Lu, L., Li, B., Wang, X., Dai, J.: Deformable detr: Deformable transformers for end-to-end object detection. *arXiv preprint arXiv:2010.04159* (2020)