

Learning to Localize Objects Improves Spatial Reasoning in Visual-LLMs

Kanchana Ranasinghe^{1,2} Satya Narayan Shukla¹ Omid Poursaeed¹
 Michael S. Ryoo² Tsung-Yu Lin¹

¹Meta ²Stony Brook University

{kranasinghe, mryoo}@cs.stonybrook.edu, {satyanshukla, opoursaeed, tsungyulin}@meta.com

Abstract

Integration of Large Language Models (LLMs) into visual domain tasks, resulting in visual-LLMs (V-LLMs), has enabled exceptional performance in vision-language tasks, particularly for visual question answering (VQA). However, existing V-LLMs (e.g. BLIP-2, LLaVA) demonstrate weak spatial reasoning and localization awareness. Despite generating highly descriptive and elaborate textual answers, these models fail at simple tasks like distinguishing a left vs right location. In this work, we explore how image-space coordinate based instruction fine-tuning objectives could inject spatial awareness into V-LLMs. We discover optimal coordinate representations, data-efficient instruction fine-tuning objectives, and pseudo-data generation strategies that lead to improved spatial awareness in V-LLMs. Additionally, our resulting model improves VQA across image and video domains, reduces undesired hallucination, and generates better contextual object descriptions. Experiments across 5 vision-language tasks involving 14 different datasets establish the clear performance improvements achieved by our proposed framework.

1. Introduction

Holistic visual understanding requires learning beyond simply content of an image to encompass awareness on spatial locations of objects and their relations [41]. In the context of visual question answering (VQA), such spatial awareness allows better reasoning involving structural and contextual information contained within an image [8].

Since the introduction of powerful large-language models (LLMs) such as GPT-3 [5], Chat-GPT [44], Vicuna [11], and LLaMA [55, 56] that are capable of human style conversation, their visual counterparts such as BLIP-2 [33], LLaVA [38] have enabled novel tasks within the vision modality. However, despite their [33, 38] highly generic visual understanding, these models exhibit poor language-based spatial reasoning [8]. In fact, they fail at simple tasks such as distinguishing whether an object lies to the left or right of another object (see Tab. 4).



Figure 1. We illustrate one unique ability of our model: contextual region description (top). Note the contextual information used in describing the selected region in each image. Explicitly teaching localization to Visual-LLMs also improves spatial awareness in VQA settings (bottom). Color boxes only for illustration purposes.

In the case of contrastive language image models (such as CLIP [46], ALIGN [26]), recent works explore how injecting explicit spatial awareness [39, 43, 48, 74] can enable more holistic visual understanding. In fact, [48] shows how such improved spatial awareness benefits model robustness in adversarial domains. This raises the question of how generative language image models, particularly those connecting LLMs to visual encoders [33, 38] can benefit from such spatial awareness specific training. We refer to models of this category that generate textual outputs given joint image-text inputs (e.g. [33, 38]) as visual-LLMs (V-LLMs).

In this work, we explore location specific instruction fine-tuning objectives that explicitly enforce V-LLMs to

meaningfully process and generate textual image-space coordinates. We hypothesize that such training would lead to improved spatial awareness in these V-LLMs, therein improving performance on VQA tasks. To this end, we propose three instruction fine-tuning objectives that unify location representation with natural language. We also explore optimal representation forms for image-space locations and how pseudo-data generation can be leveraged for efficient scaling of our framework. We name our resulting model as LocVLM.

While the idea of adapting V-LLMs to perform localization related tasks (e.g. detection, segmentation) using V-LLMs has been explored in multiple recent works [30, 45, 58, 66, 68, 73, 75], these approaches depend on task specific architectural modifications or treat localization inputs / outputs differently from natural language. In contrast, our LocVLM focuses on a unified framework treating location and language as a single modality of inputs with the goal of complementing performance in each task. We intuit that processing location represented in textual form would enforce the LLM to select appropriate image regions as opposed to relying on region level features provided by the architecture. At the same time, textual form location outputs promote spatial awareness at language level in a human interpretable manner, in contrast to using secondary heads or specialized tokens for location prediction. Concurrent work in [8] also explores textual location representation with a generic V-LLM architecture similar to our work. Our proposed LocVLM differs with focus on optimal location representation forms, data-efficient pseudo-labelling, and video domain operation.

Our proposed framework exhibits improved spatial awareness in VQA style conversation demonstrated through experimentation on 14 datasets across 5 vision-language tasks: Spatial Reasoning, Image VQA, Video VQA, Object Hallucination, and Region Description. We summarize our key contributions as follows:

- Inject textual spatial coordinate awareness into V-LLMs
- Propose three novel localization based instruction fine-tuning objectives for V-LLMs
- Discover optimal coordinate representation forms
- Pseudo-Data generation for improved region description and scaling to video domain

2. Related Work

Localization in Contrastive Vision Language Models: Foundation vision language models (VLMs) such as CLIP [46] resulted in extensive exploration into language-tied localization in images both under dense (pixel / bounding-box) supervision [16–18, 21, 27, 31, 32, 35, 69, 70] and weak supervision [14, 39, 43, 48, 60, 61, 65, 74, 76]. Recovering explicit localization information within model rep-

Method	Kosmos [45]	Ferret [66]	Shikra [8]	Ours
Unified Arch.	✗	✗	✓	✓
Purely Textual	✗	✗	✓	✓
Pseudo Data	✗	✗	✗	✓
Video Domain	✗	✗	✗	✓

Table 1. Related Work Comparison: A unified architecture, purely textual inputs, pseudo data for scalable learning, and video domain operation distinguishes our work from these prior methods.

resentations has enabled more robust operation for certain tasks [48]. While our work differs from this contrastive setting given our use of LLM based generative predictions, we similarly explore how explicit location information within the language modality can improve V-LLMs.

Visual Large Language Models (V-LLMs): The advent of powerful large language models (LLMs) such as GPT-3 [5], Chat-GPT [44], and PaLM [13], as well as their open-source counterparts BLOOM [52], Vicuna [11], and LLaMA [55, 56], has resulted in direct use of these LLMs for computer vision tasks [22, 53]. Alternate lines of work explore how LLMs can be connected to existing visual foundation models [2, 3, 33, 38, 42, 47], in particular to CLIP visual backbones [46]. While earlier models explored large-scale (millions to billions of samples) image-text training [2, 3], later models [33, 38, 42] scale down on data dependency. LLaVA [38] in particular scales down on pre-training data to under 1 million image-text pairs, and use instruction fine-tuning [59] to enable human-style conversation with visual awareness. This is extended to video domain in [42, 49]. A shortcoming of these models is their lack of spatial awareness or location understanding in image space [8, 12, 20]. Spatial reasoning limitations in generative VLMs are studied in [12, 20]. Similar failures in captioning (and VQA) models are explored in [28]. A solution in [24] proposes code-generation based reasoning. Our work tackles these same limitations but follows an alternate direction of spatial-aware instruction fine-tuning. Another line of recent works [30, 45, 58, 66, 68, 73, 75] tackle this by introducing architectural modifications to explicitly extract region level features that are injected to the LLM as special tokens. While introducing extra tokens and layers, this also separates the localization task from language. In contrast, we use a generic architectures with purely textual location information (i.e. image space coordinates as text). Concurrent work in [8] explores this same idea, but we differ in 3 ways with, a) focus on optimal coordinate representation forms, b) data-efficient pseudo-labelling strategies, and c) video domain operation (see also Tab. 1).

Location Representations: Selecting regions within an image has a rich history in computer vision [40, 57] with greater focus on location outputs since the popularity of object detection [6, 9, 10, 19, 50, 54, 58]. Early anchor-based

methods regress locations from anchor centers [19, 50], followed by direct location regression from object-level features [6, 54]. Recent works explore generative location predictions with diffusion processes [9] and sequence-generation [10, 58]. Ours resembles the latter given our use of an LLM, next token prediction objective, and sequential generation of textual location representations. However, [10, 58] utilize 1000 specialized location tokens (introduced to the LLM vocabulary) corresponding to 1000 bins uniformly spread across image space. While we explore similar binning strategies, in contrast we introduce no additional tokens, focus on purely textual representation of locations, and explore multiple textual location representation forms.

3. Method

Current V-LLMs [33, 38] exhibit weak understanding of spatial locations within images [8]. We explore and benchmark such shortcomings, and propose three novel instruction fine-tuning objectives aimed at overcoming these drawbacks of existing V-LLMs. We build these objectives based on spatial-coordinate based prompting and demonstrate how LLMs can directly both process and generate meaningful numerical coordinates in image-space after suitable training. In the rest of this section we describe our architecture and training framework, followed by coordinate processing & generation, instruction fine-tuning objectives, pseudo-data generation, and video domain operation.

3.1. Architecture and Training

The focus of our work is to explore how spatial localization related training can improve a generic V-LLM such as LLaVA [38]. Therein, our architecture and training framework is inspired from [38]. We use a visual encoder, adapter layer, and LLM stacked sequentially (illustrated in Fig. 2), and follow a multi-stage training strategy similar to [38].

Consider an image $X \in \mathbb{R}^{H,W,C}$ where H, W, C ($= 3$) denote height, width, channels of image and a textual prompt T composed of natural language (asking a question about the image). We define two variants of our model, LocVLM-B and LocVLM-L for better comparison with prior work. We first describe LocVLM-B that processes images with $H = W = 224$. Our visual encoder, ViT-L/14 from CLIP [46], processes the image X to produce a set of 256 visual tokens in \mathbb{R}^{1024} , which are in turn projected to \mathbb{R}^{4096} by an adapter layer (implemented as a linear layer). The LLaMA [55] text tokenizer processes the textual prompt T to produce textual tokens in \mathbb{R}^{4096} . The joint set of visual and textual tokens (of dimension \mathbb{R}^{4096}) are processed by a LLaMA [55] LLM to produce the final set of textual tokens which are in turn untokenized to convert to natural language. The final natural language output is expected to be a suitable response to the input textual prompt, T . In variant LocVLM-L, we use images sized

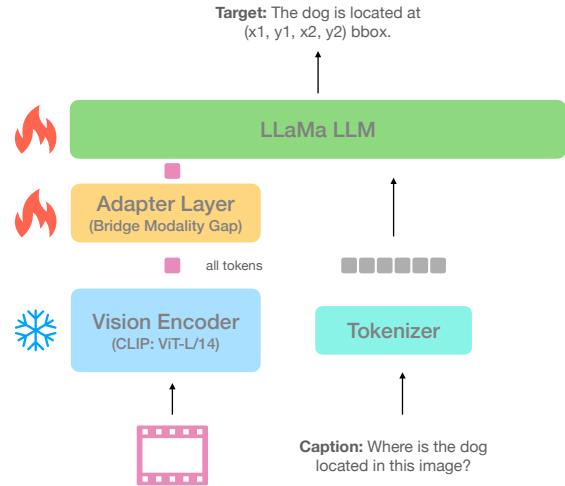


Figure 2. **Architecture:** We present the overall model architecture of our framework which is inspired from LLaVa [38].

$H = W = 336$ resulting in 576 visual token, an adapter layer implemented as an MLP, and the LLM from LLaMA-2 [56]. All other design choices remain unchanged.

We also highlight the BPE tokenization that is employed in our setup. This learned tokenization scheme may split a single word into sub-parts (that alone can appear meaningless to humans) and handles numerical text (including decimal point) as individual tokens (e.g. 12.34 would be split into 5 separate tokens).

In terms of training, we follow a two-stage strategy. Inspired by LLaVA [38], we adopt an initial pre-training stage that only updates weights of the intermediate adapter layer to align the visual encoder outputs with LLM inputs. Next, we jointly instruction fine-tune the adapter layer and LLaMA LLM with our proposed objectives and template-based localization datasets (see Sec. 4.1). Video domain operation introduces an additional phase (see Sec. 3.5).

3.2. Coordinate Processing and Generation

Humans contain the ability to reason about images using image-space coordinates. This is in contrast to existing V-LLMs that can describe the contents of an image elegantly, but lack spatial awareness regarding image contents. We hypothesize that injecting LLMs with additional spatial awareness, through coordinate based reasoning could improve their generic reasoning ability as well. To this end, we introduce our first goal of directly using textual coordinate based image locations in both natural language prompts and LLM generated outputs. For textual coordinates, we explore three different representation forms:

1. Normalized Floating Point Values
2. Integer Valued Binning (across image dimensions)
3. Deviation from Image-Grid based Anchors

CR	GQA (Acc)	RD (METEOR)	A-QA (Acc)
NFP	46.1	19.6	37.1
IVB	47.3	20.7	37.4
DIGA	47.0	20.8	37.3

Table 2. Ablation on Coordinate Representation (CR) methods: we compare each of the three proposed CR variants, namely normalized floating point values (NFP), integer valued binning (IVB), and deviation from image-grid based anchors (DIGA).

For image locations, we explore point based (e.g. center coordinates $[cx, cy]$ of object) and bounding box based (e.g. top-left and bottom-right extreme coordinates of object region $[x1, y1, x2, y2]$) forms. We next discuss the three representations for coordinates used for either location.

Normalized Floating Point Values calculates absolute image coordinates and normalizes with image dimensions to a $(0, 1)$ range. We use a 4 decimal point representation for these floating point values. While this representation is simple and generic, given the nature of BPE tokenization, each individual coordinate will be represented by up to 6 tokens.

Integer Valued Binning discretizes the absolute image coordinates to one of n_b ($=224, 336$ for variant B & L respectively) bins spread uniformly across the two image dimensions. Based on the binning parameter, n_b , each coordinate will be represented some number of tokens, in our case up to 3 (less than the floating point variant).

Deviation from Image-Grid based Anchors is motivated from prior object detection works that estimate an initial anchor followed by deviation from that anchor center to estimate bounding box coordinates. We follow a similar setup, where one of n_a anchors is predicted by the model, followed by deviation of coordinate from that anchor center. Our intuition is that, given the sequential next-token prediction setup of LLMs, such a two-stage strategy would lead to faster learning and more accurate coordinates.

We refer to Appendix A for further details on each variant. In Tab. 2, we ablate each representation format on three different tasks (see Sec. 4.7 for more details) of image VQA (GQA), region description (RD), and video VQA (A-QA). Our experiments indicate optimal performance for integer valued binning (IVB). In all following experimentation, we fix our coordinate representation to IVB.

3.3. Instruction Fine-Tuning Objectives

Given suitable coordinate representations, we now have a mechanism to directly prompt LLMs with image locations in textual form. Our second goal is to build training objectives using these image coordinates that directly inject spatial awareness into V-LLMs. We propose three instruction fine-tuning objectives for this purpose.

Let us first revisit the visual instruction fine-tuning methodology in [38]. Building off the COCO dataset,

Objective	Prompt	Target
LocPred	Where is obj1?	It’s at $(x1,y1,x2,y2)$.
NegPred	Where is obj2?	There’s no obj2.
RevLoc	Describe (cx,cy)	<i>Detailed description</i>

Table 3. We summarize our three distinct instruction fine-tuning objectives. Refer to appendix (Appendix B) for exact natural language prompts and targets used for training. For illustration, we use both point and bounding-box based image locations here.

they construct a VQA dataset containing conversation style question-answer pairs relevant to each COCO image. Question-answer pairs are generated using an LLM that is fed with the ground-truth bounding-box annotations for each image. Inspired by this setup, we build a similar spatial-VQA dataset using images and annotations of the COCO dataset, but instead of LLM prompting, we utilize hand-crafted templates and pseudo-captions (discussed in Sec. 3.4) to generate conversations.

We propose three types of question-answer pairs that relate to our three instruction fine-tuning objectives: Location Prediction (LocPred), Negative Prediction (NegPred), and Reverse-Location Prediction (RevLoc). See Tab. 3 for examples. Considering the LLM based final text generation in our architecture, we utilize next-token prediction loss to achieve each objective during our training.

Location Prediction: Given an object category, we query the model to generate a point or bounding box localizing that object in the image. The object category and bounding box are derived from the COCO train set annotations. To avoid object mismatches (i.e. multiple object of same category), we first filter images containing only a single object of a given class.

Negative Prediction: Using the same prompt templates as in *LocPred* above, we query the model to generate a point or bounding box localizing a specified object in the image. However, in this case we select an object category not present in the image and accordingly provide a target text of “no such object in image”. For each image, we utilize COCO bounding-box annotations to discover objects (belonging to COCO classes) that are not present in that image.

Reverse-Location Prediction: We perform the reverse of *LocPred* here. Given a point or bounding box in image space, we query the model to describe the object in that location. The bounding box and object category are derived from the COCO train set annotations.

While introducing three novel train objectives aimed at injecting location information, we highlight that our proposed framework relies on training data (i.e. human annotations) identical to those used by LLaVA [38]. We do not use any additional ground-truth annotations for training. Next we explore how we could augment the generality of our framework while limiting to this same annotated data.

3.4. Pseudo-Data Generation

We introduced three train objectives, each utilizing template based conversations as prompts and targets. However, our reliance on categories of COCO dataset limits the object vocabulary seen during training. Therein, we propose a pre-trained V-LLM based pseudo-data generation strategy. In fact, we utilize our model after stage one training as the V-LLM leading to a form of self-training based learning. Given the abundance of only image-level annotated datasets (i.e. no bounding box ground-truth), we also explore how an object-detector generated pseudo-bounding boxes could augment our framework.

Self-Training: Given an image and bounding box annotations from the COCO train set, we prompt the V-LLM to caption each distinct object in the image. In order to prevent ambiguous object queries, we filter images to select only those containing at most one instance of a single category. We additionally prompt the V-LLM to describe the object using relational information (i.e. relative to other objects in image). This process provides us a dataset with object level bounding boxes and descriptive captions that are not limited to the COCO object categories (dataset details in Appendix C). In turn, we use this data generated by the V-LLM (our stage one model) to further improve performance of our framework. We modify each of our three train objectives (in Sec. 3.3) to utilize these image-specific pseudo-captions instead of the generic dataset level category labels.

Weak Supervision: We explore how datasets containing no object level annotation (e.g. video classification / VQA datasets) could be leveraged to adapt our framework into domains beyond images. Therein, we utilize an off-the-shelf panoptic segmentation framework from SEEM [77] to generate pseudo-bounding boxes for selected object categories within any image as well as exhaustive pixel level labels (enabling negative class identification). We leverage this setup to extend our introduced train objectives to the video domain as well.

3.5. Video Domain Operation

Inspired by the simple modifications to LLaVA [38] in [42] enabling video domain operation, we follow a similar strategy of modifying our LocVLM-B architecture to process videos while introducing no additional components. The visual backbone process multiple video frames individually (as images) and resulting tokens are averaged across spatial (S) and temporal (T) axes to obtain $S + T$ tokens. These are processed by the adapter layer and LLM to generate the textual outputs. Further details on our video architecture are discussed in Appendix D.

In addition to our two training phases discussed in Sec. 3.1, we introduce a third video instruction fine-tuning stage using a dataset we derive from ActivityNet [23]. Fol-

lowing [42], only the adapter layer is fine-tuned leaving all other parameters frozen. This resulting model is referred to as LocVLM-Vid-B.

We next introduce video variants of our three instruction fine-tuning objectives focused on static objects in videos. We utilize our proposed pseudo-labeling strategy to generate necessary video annotations and train both the adapter layer and LLM to obtain a second video model tagged LocVLM-Vid-B+. Further details on our video fine-tuning objectives are presented in Appendix D.

4. Experiments

In this section, we present experimental results to highlight existing weaknesses of SOTA V-LLMs and how our proposed framework addresses these issues. We also evaluate on standard VQA benchmarks across domains to showcase the better reasoning abilities of our model and highlight novel abilities of our framework.

4.1. Experimental Setup

Datasets: We utilize the COCO dataset [37] and our model (post stage one training) to construct a localization related VQA dataset as outlined in Sections 3.3 and 3.4. We name this dataset *Localize-Instruct-200K*. In detail, this contains LocPred and RevLoc question-answer pairs that use pseudo-captions instead of COCO categories as well as NegPred. We define a second video dataset, *Localize-ActivityNet*, containing question-answer pairs constructed from Activity-Net pseudo-bounding boxes following Section 3.5. Our models are primarily trained on our Localize-Instruct-200K dataset. Our stage one training uses CC3M dataset [7]. Additionally, our Localize-ActivityNet dataset and ActivityNet dataset [23] are used for video domain training.

Training: We train our models on 8xA100 GPUs (each 80GB) following a two-phase training schedule. Our first phase trains on CC3M [7] following the setup in [38]. The second phase uses our Localize-Instruct-200K dataset and trains for 10 epochs with a batch size of 64, ADAM-W optimizer with initial learning rate $2e - 5$, 0.3 warm-up ratio, and cosine-decay learning rate schedule. Both the training phases we conduct use standard next-token-prediction loss used in LLM training.

Evaluation: During evaluation, following standard protocol [38], we iteratively generate next tokens, given visual and textual inputs. The LLM output is a distribution across the entire token vocabulary. The next token is selected through multinomial sampling of this output using a softmax temperature term of 0.2 during normalization.

4.2. Spatial Reasoning: A Toy Experiment

We investigate spatial reasoning abilities of two SOTA V-LLMs, LLaVA [38] and BLIP-2 [33], using a simple toy experiment. We create an evaluation dataset from COCO annotations containing images with distinct category object triplets (only one instance occurrence of each object category), where each object is entirely to the left or right half of the image and two objects are on opposite sides. The ground-truth bounding box annotation are utilized to automate this dataset creation procedure. This evaluation set, referred as *COCO-Spatial-27K*, contains 26,716 image-question pairs (see Appendix C for details). We introduce two evaluation settings, direct VQA and in-context learning (ICL) VQA to understand spatial reasoning abilities of these models. In direct VQA, given an image we query the model whether an object lies above or below another object. In ICL VQA, before a similar final query, we provide two example question-answer pairs (involving the other two objects in the image) in the same format as our query. Refer to Appendix E for further details on task. We perform the same for objects in top vs bottom halves of images.

These results are presented in Table 4. Our results indicate near random performance for existing V-LLMs. For the case of LLaVA, we perform keyword (*left* and *right*) frequency analysis on its instruction tuning dataset (LLaVA-Instruct-80K dataset) to verify the presence of terms *left* and *right* in its training corpus. These keywords are present in 0.37% and 1.13% of its conversations respectively (see Appendix F for more) indicating presence of these concepts in the image-text training corpus. In contrast to these methods, our proposed framework notably improves performance over both BLIP-2 [33] and the LLaVA baseline [38].

4.3. Image VQA

Image VQA involves correctly answering natural language questions regarding content within an image. We evaluate our model for Image VQA on two standard datasets, GQA and VQAv2. The GQA dataset focuses on questions requiring compositional reasoning, particularly involving surrounding information of objects within an image. We evaluate on its test-dev split containing 12,578 image-question pairs. The VQAv2 dataset contains open-ended questions about each image that require an understanding of vision, language and commonsense knowledge to answer. We use its validation split containing 214,354 image-question pairs for our evaluation. For each dataset, we follow standard V-LLM evaluation protocol following [33, 42] and report top-1 accuracy metric. Our results in Tab. 5 indicate clear improvements for LocVLM over our baseline and prior work, establishing the usefulness of our proposed framework. The closest to our work, Shikra [8] achieves performance competitive to our LocVLM-B, but

Method	ICL	All	Left	Right	All	Above	Below
BLIP-2 [33]	✗	45.5	86.1	4.74	49.2	50.4	48.6
LLava [38]	✗	55.1	84.5	36.5	58.9	57.8	59.3
Ours	✗	69.5	79.7	59.2	65.4	64.2	65.9
BLIP-2 [33]	✓	14.7	17.8	11.6	15.8	16.5	15.2
LLaVa [38]	✓	55.1	84.7	36.4	58.2	57.7	58.5
Ours	✓	76.5	90.4	61.5	74.1	73.5	74.4

Table 4. **Spatial Reasoning:** We report accuracy (%) on a spatial localization dataset derived from COCO annotations to highlight weak spatial awareness of existing V-LLMs. We query these models to answer whether one object is to the left or right / above or below of another object. The SOTA V-LLMs evaluated exhibit close to random performance. Our proposed setup outperforms existing methods. Ours refers to LocVLM-B variant.

Method	LLM	VS	Zero-Shot	GQA	VQA-V	VQA-T
SR [4]	-	-	✗	62.1	72.9	-
Shikra [8]	7B	224	✗	-	75.3	77.4
LLaVA-v1.5	7B	336	✗	62.0	78.1	78.4
LocVLM-L	7B	336	✗	63.5	78.2	78.6
LLaVA-v1	7B	224	✓	44.7	49.8	49.3
LocVLM-B	7B	224	✓	47.3	50.3	50.8
Viper-GPT	175B	-	✓	48.1	-	-
BLIP-2	11B	-	✓	44.7	54.3	53.9
LLaVA-v1.5	7B	336	✓	48.7	55.7	55.3
LocVLM-L	7B	336	✓	50.2	55.9	56.2

Table 5. **Image VQA Results:** We report accuracy (%) on the test-dev split of GQA dataset (GQA) and the validation / test splits of VQAv2 dataset (VQA-V / VQA-T). Our proposed LocVLM improves over prior works achieving state-of-the-art performance.

unlike ours they use VQA datasets (containing similar domain question-answer pairs) during training.

4.4. Video VQA

Our model is also applicable to video tasks following our video domain adaptation described in Sec. 3.5. We simply adopt the additional video instruction fine-tuning phase from [42] on the ActivityNet dataset after our initial two phases of training to obtain LocVLM-Vid-B. This third phase involves fine-tuning only the adapter layer of our model. We also explore video variants of our IFT objectives that train both adapter layer and LLM. The resulting model is termed LocVLM-Vid-B+.

Video VQA focuses on correctly answering questions regarding a given video that require spatio-temporal awareness to answer. We evaluate our video-adapted model on the task of zero-shot video VQA on four benchmark datasets, ActivityNet-QA, MSRVT-TQA, MSVD-QA, and TGIF-QA. We evaluate on the validation splits of these four datasets. ActivityNet-QA videos cover a wide range of complex human activities relevant to daily living with its question-answer pairs focusing on long-term spatio-temporal reasoning. MSRVT-TQA builds off the MSRVT-T dataset that contains web videos covering a comprehensive

Method	Zero-Shot	ActivityNet-QA	MSRVTT-QA	MSVD-QA	TGIF-QA
JustAsk [63]	✗	38.9	41.8	47.5	-
FrozenBiLM [64]	✗	43.2	47.0	54.8	-
VideoCoCa [62]	✗	56.1	46.3	56.9	-
Flamingo [2]	✓	-	17.4	35.6	-
BLIP-2 [33]	✓	-	17.4	34.4	-
InstructBLIP [15]	✓	-	25.6	44.3	-
FrozenBiLM [64]	✓	24.7	16.8	32.2	41.0
Video Chat [34]	✓	26.5	45.0	56.3	34.4
LLaMA Adapter [72]	✓	34.2	43.8	54.9	-
Video LLaMA [71]	✓	12.4	29.6	51.6	-
Video-ChatGPT [42]	✓	35.2	49.3	64.9	51.4
LocVLM-Vid-B	✓	37.4	51.2	66.1	51.8

Table 6. **Video VQA Results:** Our proposed LocVLM-Vid-B improves over Video-ChatGPT [42] and achieves state-of-the-art results (Top-1 Accuracy %) across four different video VQA benchmarks. Note the zero-shot setting of all these evaluations.

range of categories and diverse visual content. MSVD-QA is a similar dataset building off the MSVD dataset. TGIF-QA contains question-answer pairs from a dataset constructed of animated GIFs. For each dataset, we report the accuracy metric following evaluation protocol in [42]. Our results on these four datasets reported in Tab. 6 demonstrate state-of-the-art performance of our proposed LocVLM-Vid-B, with consistent improvements over the baseline from [42]. Here we use the LocVLM-Vid-B variant for fairer comparison with the baseline from [42]. We attribute the performance gains exhibited by our model to its stronger spatial awareness (see Sec. 4.2). Particularly in the case of video understanding, awareness of content at spatial level of each frame is significant to understand object motions and interactions [1, 51]. We also report more results involving additional model variants in Tab. 7.

4.5. Object Hallucination

Current state-of-the-art V-LLMs suffer from object hallucination, generating image descriptions inconsistent with the image content [36]. For example, a V-LLM would respond to “Where is the cat in this image?” with “The cat is on the table” when in reality there is no cat in the image. We evaluate the extent of hallucination in V-LLMs using three datasets we introduce (details in Appendix C) and the POPE dataset [36]. Our three datasets, Hal-COCO, Hal-ADE, and Hal-Act build off COCO, ADE-20K, and ActivityNet datasets respectively. The first two involve images and the latter videos. These datasets contain ‘Is there *obj* in image / video?’ type questions per sample, for two objects present and not present in the image / video. Hal-ADE object categories contain *no overlap* with COCO classes allowing evaluation on novel object categories unseen during our instruction fine-tuning. Results reported in Tab. 8 show clear improvements of LocVLM-B over baselines. We also evaluate LocVLM-B on the POPE benchmark [36] that builds

Method	VLT	Frames	Acc (%)
LLaVa (v1) [38]	✗	1	28.7
LLaVa (v1.5) [38]	✗	1	31.5
LocVLM-B	✗	1	29.2
LocVLM-L	✗	1	32.1
Video-ChatGPT [42]	✓	100	35.2
LocVLM-Vid-B	✓	100	37.4
LocVLM-Vid-B+	✓	8	38.2

Table 7. **Video VQA:** We report more results (Top-1 Accuracy) for ActivityNet-QA dataset including multiple baseline and LocVLM variants. Our proposed models exhibit top performance. VLT denotes video level training. More details in Appendix D.

Method	Hal-COCO	Hal-ADE	Hal-Act
Shikra [8]	86.2	58.7	-
LLaVa [38]	61.9	53.8	-
LocVLM-B	88.3	75.2	-
Video-ChatGPT [42]	-	-	50.6
LocVLM-Vid-B	-	-	68.7
LocVLM-Vid-B+	-	-	72.4

Table 8. **Hallucination Evaluation:** We report top-1 accuracy (%) for object presence type questions and showcase reduced object hallucination in our proposed framework.

off the COCO dataset object annotations and report results in Tab. 9. Our LocVLM showcases similar performance improvements on this dataset.

4.6. Region Description

A unique characteristic of our model (in contrast to V-LLMs like LLaVA [38] & BLIP-2 [33]) is its ability to reason with prompts involving coordinate based image space locations without any input modifications. Given a point or bounding box location, we prompt our model to generate an output describing that location. We refer to this unique ability of our model as *region description* (RD). We evaluate this RD capability of our model by generating object

Datasets	Metrics	BLIP-2	Shikra	LLaVA	Ours
Random	Accuracy (\uparrow)	88.6	86.9	50.4	87.9
	Precision (\uparrow)	84.1	94.4	50.2	83.6
	Recall (\uparrow)	95.1	79.3	99.1	93.9
	F1 Score (\uparrow)	89.3	86.2	66.6	88.5
	Yes	56.6	43.3	98.8	56.2
Popular	Accuracy (\uparrow)	82.8	84.0	49.9	86.0
	Precision (\uparrow)	76.3	87.6	49.9	79.7
	Recall (\uparrow)	95.1	79.2	99.3	93.9
	F1 Score (\uparrow)	84.7	83.2	66.4	86.3
	Yes	62.4	45.2	99.4	58.9
Adversarial	Accuracy (\uparrow)	72.1	83.1	49.7	78.8
	Precision (\uparrow)	65.1	85.6	49.9	76.6
	Recall (\uparrow)	95.1	79.6	99.1	93.7
	F1 Score (\uparrow)	77.3	82.5	66.3	84.3
	Yes	73.0	46.5	99.4	61.7

Table 9. **More object hallucination:** Results on POPE evaluation benchmark [36] indicate strong performance of our model.

Method	ZS	RefCOCO	RefCOCO+	RefCOCOg	
				Val	Test
SLR [67]	\times	-	-	-	15.4
SLR + Rerank [67]	\times	-	-	-	15.9
Kosmos-2 [45]	\times	8.67	8.82	14.3	14.1
Shikra [8]	\times	10.4	11.1	19.7	19.5
LLaVa [38]	\times	8.43	8.73	13.5	13.5
LocVLM-B	\times	14.6	15.2	26.0	26.2
Kosmos-2 [45]	\checkmark	6.34	8.25	12.4	12.2
LLava [38]	\checkmark	4.23	7.26	10.6	10.3
LocVLM-B	\checkmark	11.0	11.1	20.6	20.7

Table 10. **Region Description:** We report METEOR scores for RD task [45]. Test-B split is used for RefCOCO & RefCOCO+ datasets. Our method outperforms all prior work.

level descriptions focused on contextual information (e.g. surrounding of that object in the image). Following evaluation protocol in [45] for region description, we extend their evaluation to three standard referring localization datasets from RefCOCO [29] and report these results in Tab. 10. We select the METEOR score as the evaluation metric to account for variations in word choice in generated answers which may be acceptable in various cases (e.g. different sentence structure leading to alternate word ordering). Our results indicate clear improvements over the LLaVA baseline [38] as well as prior state-of-the-art. We attribute these improvements to our pseudo-data based training.

4.7. Ablations

Next we conduct ablative studies on separate components of our proposed setup: IFT objectives, location type, and pseudo-data. We follow the same training strategy as described in Sec. 4.1 and present these results in Tab. 11. LocVLM-B is used for all these experiments. The significance of each IFT objective is verified in Tab. 11 (top) with

LocPred	NegPred	RevLoc	GQA	RD	A-QA
\times	\times	\times	44.7	10.3	35.2
\checkmark	\times	\times	45.2	12.2	35.8
\checkmark	\checkmark	\times	46.9	12.5	37.2
\checkmark	\checkmark	\checkmark	47.3	20.7	37.4
Location Type		PD	GQA	RD	A-QA
Point		\checkmark	47.3	20.6	37.4
Bounding Box		\checkmark	47.3	20.7	37.4
Bounding Box		\times	46.5	11.6	37.1

Table 11. **Ablations:** We report top-1 accuracy (%) on GQA and ActivityNet-QA (A-QA) datasets and METEOR scores for RD task on RefCOCOg test split. (top) We ablate proposed instruction fine-tuning objectives to verify usefulness of each objective. (bottom) We first ablate point based and bounding box based location forms to showcase minimal difference across them. We next ablate use of object description pseudo-data (PD). We highlight the improvements due to pseudo-data, especially on the RD task.

LocPred	NegPred	RevLoc	A-QA
\times	\times	\times	37.4
\checkmark	\times	\times	37.6
\checkmark	\checkmark	\times	38.2
\checkmark	\checkmark	\checkmark	38.2

Table 12. **Video Ablation:** We report top-1 accuracy (%) on ActivityNet-QA (A-QA) dataset. Results indicate the generality of proposed IFT objectives for video domain training as well.

consistent performance improvements across tasks. The generality of our approach to differing location type (i.e. points vs bounding boxes) and usefulness of pseudo-data is visible in Tab. 11 (bottom). In particular, we highlight the notable performance improvement for RD task gained from using pseudo-data. We also conduct ablations for our video domain training setup and report these results in Tab. 12. The LocVLM-Vid-B+ variant is used in these experiments. Our results showcase the usefulness of proposed IFT objectives for video domain learning as well.

5. Conclusion

We introduce a simple framework that equips visual-LLMs (V-LLMs) with greater spatial understanding, termed LocVLM. We leverage the idea of encoding image coordinates within language to propose three instruction fine-tuning (IFT) objectives. This training process endows V-LLMs with the ability to reason about spatial composition of images using image space coordinates within text. A data efficient training pipeline utilizing pseudo-data allows our approach to achieve state-of-the-art results in Image VQA, Video VQA, and Region Description while improving spatial awareness and reducing object hallucination.

References

- [1] Jake K. Aggarwal and Michael S. Ryoo. Human activity analysis. *ACM Computing Surveys (CSUR)*, 43:1 – 43, 2011. [7](#)
- [2] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katie Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. *arXiv preprint arXiv:2204.14198*, 2022. [2](#), [7](#)
- [3] Anas Awadalla, Irena Gao, Joshua Gardner, Jack Hessel, Yusuf Hanafy, Wanrong Zhu, Kalyani Marathe, Yonatan Bitton, Samir Gadre, Jenia Jitsev, Simon Kornblith, Pang Wei Koh, Gabriel Ilharco, Mitchell Wortsman, and Ludwig Schmidt. Openflamingo, 2023. [2](#)
- [4] Pratyay Banerjee et al. Weakly supervised relative spatial reasoning for visual question answering. *ICCV*, 2021. [6](#)
- [5] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020. [1](#), [2](#)
- [6] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part I 16*, pages 213–229. Springer, 2020. [2](#), [3](#)
- [7] Soravit Changpinyo, Piyush Kumar Sharma, Nan Ding, and Radu Soricut. Conceptual 12m: Pushing web-scale image-text pre-training to recognize long-tail visual concepts. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3557–3567, 2021. [5](#)
- [8] Keqin Chen, Zhao Zhang, Weili Zeng, Richong Zhang, Feng Zhu, and Rui Zhao. Shikra: Unleashing multi-modal llm’s referential dialogue magic. *arXiv preprint arXiv:2306.15195*, 2023. [1](#), [2](#), [3](#), [6](#), [7](#), [8](#)
- [9] Shoufa Chen, Peize Sun, Yibing Song, and Ping Luo. Diffusiondet: Diffusion model for object detection. *arXiv preprint arXiv:2211.09788*, 2022. [2](#), [3](#)
- [10] Ting Chen, Saurabh Saxena, Lala Li, David J Fleet, and Geoffrey Hinton. Pix2seq: A language modeling framework for object detection. *arXiv preprint arXiv:2109.10852*, 2021. [2](#), [3](#), [1](#)
- [11] Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E. Gonzalez, Ion Stoica, and Eric P. Xing. Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality, 2023. [1](#), [2](#)
- [12] Jaemin Cho et al. Dall-eval: Probing the reasoning skills and social biases of text-to-image generation models. *ICCV*, 2023. [2](#)
- [13] Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. Palm: Scaling language modeling with pathways. *arXiv preprint arXiv:2204.02311*, 2022. [2](#)
- [14] Yufeng Cui, Lichen Zhao, Feng Liang, Yangguang Li, and Jing Shao. Democratizing contrastive language-image pre-training: A clip benchmark of data, model, and supervision. *arXiv preprint arXiv:2203.05796*, 2022. [2](#)
- [15] Wenliang Dai, Junnan Li, Dongxu Li, Anthony Meng Huat Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale Fung, and Steven Hoi. Instructblip: Towards general-purpose vision-language models with instruction tuning. *arXiv preprint arXiv:2305.06500*, 2023. [7](#)
- [16] Jian Ding, Nan Xue, Guisong Xia, and Dengxin Dai. Decoupling zero-shot semantic segmentation. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11573–11582, 2021. [2](#)
- [17] Zi-Yi Dou, Aishwarya Kamath, Zhe Gan, Pengchuan Zhang, Jianfeng Wang, Linjie Li, Zicheng Liu, Ce Liu, Yann LeCun, Nanyun Peng, Jianfeng Gao, and Lijuan Wang. Coarse-to-fine vision-language pre-training with fusion in the backbone. *ArXiv*, abs/2206.07643, 2022.
- [18] Golnaz Ghiasi, Xiuye Gu, Yin Cui, and Tsung-Yi Lin. Open-vocabulary image segmentation. In *ECCV*, 2022. [2](#)
- [19] Ross Girshick. Fast r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 1440–1448, 2015. [2](#), [3](#)
- [20] Gokhale et al. Benchmarking spatial relationships in text-to-image generation, 2022. [2](#)
- [21] Xiuye Gu, Tsung-Yi Lin, Weicheng Kuo, and Yin Cui. Zero-shot detection via vision and language knowledge distillation. *arXiv e-prints*, pages arXiv–2104, 2021. [2](#)
- [22] Tanmay Gupta and Aniruddha Kembhavi. Visual programming: Compositional visual reasoning without training. *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14953–14962, 2022. [2](#)
- [23] Fabian Caba Heilbron, Victor Escorcia, Bernard Ghanem, and Juan Carlos Niebles. Activitynet: A large-scale video benchmark for human activity understanding. *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 961–970, 2015. [5](#)
- [24] Joy Hsu et al. What’s left? concept grounding with logic-enhanced foundation models. *NeurIPS*, 2023. [2](#)
- [25] Drew A. Hudson and Christopher D. Manning. Gqa: A new dataset for real-world visual reasoning and compositional question answering. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6693–6702, 2019. [3](#)
- [26] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc V. Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *International Conference on Machine Learning*, 2021. [1](#)
- [27] Aishwarya Kamath, Mannat Singh, Yann LeCun, Gabriel Synnaeve, Ishan Misra, and Nicolas Carion. Mdetmodulated detection for end-to-end multi-modal understanding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1780–1790, 2021. [2](#)
- [28] Amita Kamath et al. What’s ”up” with vision-language models? investigating their struggle with spatial reasoning. *EMNLP*, 2023. [2](#)

- [29] Sahar Kazemzadeh, Vicente Ordonez, Mark Matten, and Tamara Berg. Referitgame: Referring to objects in photographs of natural scenes. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 787–798, 2014. 8
- [30] Xin Lai, Zhuotao Tian, Yukang Chen, Yanwei Li, Yuhui Yuan, Shu Liu, and Jiaya Jia. Lisa: Reasoning segmentation via large language model. *arXiv preprint arXiv:2308.00692*, 2023. 2
- [31] Boyi Li, Kilian Q Weinberger, Serge Belongie, Vladlen Koltun, and René Ranfl. Language-driven semantic segmentation. *ICLR*, 2022. 2
- [32] Jiahao Li, Greg Shakhnarovich, and Raymond A. Yeh. Adapting clip for phrase localization without further training. *ArXiv*, abs/2204.03647, 2022. 2
- [33] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. *arXiv preprint arXiv:2301.12597*, 2023. 1, 2, 3, 6, 7
- [34] Kunchang Li, Yanan He, Yi Wang, Yizhuo Li, Wenhai Wang, Ping Luo, Yali Wang, Limin Wang, and Yu Qiao. Videochat: Chat-centric video understanding. *arXiv preprint arXiv:2305.06355*, 2023. 7
- [35] Liunian Harold Li, Pengchuan Zhang, Haotian Zhang, Jianwei Yang, Chunyuan Li, Yiwu Zhong, Lijuan Wang, Lu Yuan, Lei Zhang, Jenq-Neng Hwang, Kai-Wei Chang, and Jianfeng Gao. Grounded language-image pre-training. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10955–10965, 2021. 2
- [36] Yifan Li, Yifan Du, Kun Zhou, Jinpeng Wang, Wayne Xin Zhao, and Ji-Rong Wen. Evaluating object hallucination in large vision-language models. *arXiv preprint arXiv:2305.10355*, 2023. 7, 8, 2
- [37] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13*, pages 740–755. Springer, 2014. 5, 2, 3
- [38] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *arXiv preprint arXiv:2304.08485*, 2023. 1, 2, 3, 4, 5, 6, 7, 8
- [39] Huaishao Luo, Junwei Bao, Youzheng Wu, Xiaodong He, and Tianrui Li. Segclip: Patch aggregation with learnable centers for open-vocabulary semantic segmentation. *ArXiv*, abs/2211.14813, 2022. 1, 2
- [40] Jitendra Malik. Visual grouping and object recognition. In *Proceedings 11th International Conference on Image Analysis and Processing*, pages 612–621. IEEE, 2001. 2
- [41] David Marr. *Vision: A computational investigation into the human representation and processing of visual information*. MIT press, 1982. 1
- [42] Salman Khan Muhammad Maaz, Hanoona Rasheed and Fahad Khan. Video-chatgpt: Towards detailed video understanding via large vision and language models. *ArXiv 2306.05424*, 2023. 2, 5, 6, 7, 3
- [43] Jishnu Mukhoti, Tsung-Yu Lin, Omid Poursaeed, Rui Wang, Ashish Shah, Philip H. S. Torr, and Ser Nam Lim. Open vocabulary semantic segmentation with patch aligned contrastive learning. *CVPR*, abs/2212.04994, 2023. 1, 2
- [44] OpenAI. GPT-4 technical report. <https://arxiv.org/abs/2303.08774>, 2023. 1, 2
- [45] Zhiliang Peng, Wenhui Wang, Li Dong, Yaru Hao, Shaohan Huang, Shuming Ma, and Furu Wei. Kosmos-2: Grounding multimodal large language models to the world. *arXiv preprint arXiv:2306.14824*, 2023. 2, 8
- [46] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 1, 2, 3
- [47] Kanchana Ranasinghe and Michael S. Ryoo. Language-based action concept spaces improve video self-supervised learning. In *NeurIPS*, 2023. 2
- [48] Kanchana Ranasinghe, Brandon McKinzie, Sachin Ravi, Yinfei Yang, Alexander Toshev, and Jonathon Shlens. Perceptual grouping in contrastive vision-language models. In *ICCV*, 2023. 1, 2
- [49] Kanchana Ranasinghe, Xiang Li, Kumara Kahatapitiya, and Michael S. Ryoo. Understanding long videos in one multimodal language model pass, 2024. 2
- [50] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 779–788, 2016. 2, 3
- [51] Michael S. Ryoo and Jake K. Aggarwal. Recognition of composite human activities through context-free grammar based representation. *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*, 2:1709–1718, 2006. 7
- [52] Teven Le Scao, Angela Fan, Christopher Akiki, Ellie Pavlick, Suzana Ilić, Daniel Hesslow, Roman Castagné, Alexandra Sasha Luccioni, François Yvon, Matthias Galle, et al. Bloom: A 176b-parameter open-access multilingual language model. *arXiv preprint arXiv:2211.05100*, 2022. 2
- [53] D’idac Sur’is, Sachit Menon, and Carl Vondrick. Vipergpt: Visual inference via python execution for reasoning. *ArXiv*, abs/2303.08128, 2023. 2
- [54] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. Fcos: Fully convolutional one-stage object detection. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9627–9636, 2019. 2, 3
- [55] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023. 1, 2, 3
- [56] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023. 1, 2, 3
- [57] Jasper RR Uijlings, Koen EA Van De Sande, Theo Gevers, and Arnold WM Smeulders. Selective search for object recognition. *IJCV*, 104(2):154–171, 2013. 2

- [58] Wenhai Wang, Zhe Chen, Xiaokang Chen, Jiannan Wu, Xizhou Zhu, Gang Zeng, Ping Luo, Tong Lu, Jie Zhou, Yu Qiao, et al. Visionllm: Large language model is also an open-ended decoder for vision-centric tasks. *arXiv preprint arXiv:2305.11175*, 2023. [2](#), [3](#), [1](#)
- [59] Jason Wei, Maarten Bosma, Vincent Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M. Dai, and Quoc V. Le. Finetuned language models are zero-shot learners. *ArXiv*, abs/2109.01652, 2021. [2](#)
- [60] Jiarui Xu, Shalini De Mello, Sifei Liu, Wonmin Byeon, Thomas Breuel, Jan Kautz, and Xiaolong Wang. Groupvit: Semantic segmentation emerges from text supervision. *CVPR*, 2022. [2](#)
- [61] Jilan Xu, Junlin Hou, Yuejie Zhang, Rui Feng, Yi Wang, Yu Qiao, and Weidi Xie. Learning open-vocabulary semantic segmentation models from natural language supervision. *ArXiv*, abs/2301.09121, 2023. [2](#)
- [62] Shen Yan, Tao Zhu, Zirui Wang, Yuan Cao, Mi Zhang, Soham Ghosh, Yonghui Wu, and Jiahui Yu. Videococa: Video-text modeling with zero-shot transfer from contrastive captioners. *arXiv*, 2022. [7](#)
- [63] Antoine Yang, Antoine Miech, Josef Sivic, Ivan Laptev, and Cordelia Schmid. Just ask: Learning to answer questions from millions of narrated videos. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1686–1697, 2021. [7](#)
- [64] Antoine Yang, Antoine Miech, Josef Sivic, Ivan Laptev, and Cordelia Schmid. Zero-shot video question answering via frozen bidirectional language models. *arXiv preprint arXiv:2206.08155*, 2022. [7](#)
- [65] Lewei Yao, Runhui Huang, Lu Hou, Guansong Lu, Minzhe Niu, Hang Xu, Xiaodan Liang, Zhenguo Li, Xin Jiang, and Chunjing Xu. FILIP: Fine-grained interactive language-image pre-training. In *ICLR*, 2022. [2](#)
- [66] Haoxuan You, Haotian Zhang, Zhe Gan, Xianzhi Du, Bowen Zhang, Zirui Wang, Liangliang Cao, Shih-Fu Chang, and Yinfei Yang. Ferret: Refer and ground anything anywhere at any granularity. *ArXiv*, abs/2310.07704, 2023. [2](#)
- [67] Licheng Yu, Hao Tan, Mohit Bansal, and Tamara L. Berg. A joint speaker-listener-reinforcer model for referring expressions. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3521–3529, 2016. [8](#)
- [68] Yuhang Zang, Wei Li, Jun Han, Kaiyang Zhou, and Chen Change Loy. Contextual object detection with multimodal large language models. *arXiv preprint arXiv:2305.18279*, 2023. [2](#)
- [69] Yan Zeng, Xinsong Zhang, and Hang Li. Multi-grained vision language pre-training: Aligning texts with visual concepts. *ArXiv*, abs/2111.08276, 2021. [2](#)
- [70] Haotian Zhang, Pengchuan Zhang, Xiaowei Hu, Yen-Chun Chen, Liunian Harold Li, Xiyang Dai, Lijuan Wang, Lu Yuan, Jenq-Neng Hwang, and Jianfeng Gao. Glipv2: Unifying localization and vision-language understanding. *ArXiv*, abs/2206.05836, 2022. [2](#)
- [71] Hang Zhang, Xin Li, and Lidong Bing. Video-llama: An instruction-tuned audio-visual language model for video understanding. *ArXiv*, abs/2306.02858, 2023. [7](#)
- [72] Renrui Zhang, Jiaming Han, Aojun Zhou, Xiangfei Hu, Shilin Yan, Pan Lu, Hongsheng Li, Peng Gao, and Yu Jiao Qiao. Llama-adapter: Efficient fine-tuning of language models with zero-init attention. *ArXiv*, abs/2303.16199, 2023. [7](#)
- [73] Shilong Zhang, Peize Sun, Shoufa Chen, Min Xiao, Wenqi Shao, Wenwei Zhang, Kai Chen, and Ping Luo. Gpt4roi: Instruction tuning large language model on region-of-interest. *arXiv preprint arXiv:2307.03601*, 2023. [2](#)
- [74] Yabo Zhang, Zihao Wang, Jun Hao Liew, Jingjia Huang, Manyu Zhu, Jiashi Feng, and Wangmeng Zuo. Associating spatially-consistent grouping with text-supervised semantic segmentation. *ArXiv*, abs/2304.01114, 2023. [1](#), [2](#)
- [75] Yang Zhao, Zhijie Lin, Daquan Zhou, Zilong Huang, Jiashi Feng, and Bingyi Kang. Bubogpt: Enabling visual grounding in multi-modal llms. *arXiv preprint arXiv:2307.08581*, 2023. [2](#)
- [76] Chong Zhou, Chen Change Loy, and Bo Dai. Extract free dense labels from clip. In *European Conference on Computer Vision*, 2021. [2](#)
- [77] Xueyan Zou, Jianwei Yang, Hao Zhang, Feng Li, Linjie Li, Jianfeng Gao, and Yong Jae Lee. Segment everything everywhere all at once. In *NeurIPS*, 2023. [5](#), [2](#)