



浙江大學
ZHEJIANG UNIVERSITY



蚂蚁集团
ANT GROUP



THE HONG KONG
POLYTECHNIC UNIVERSITY
香港理工大學



Osprey

Pixel Understanding with Visual Instruction Tuning

Yuqian Yuan*, Wentong Li*, Jian Liu, Dongqi Tang, Xinjie Luo, Chi Qin,

Lei Zhang, Jianke Zhu[†]

*Equal Contribution †Corresponding author

<https://github.com/CircleRadon/Osprey>



Motivation



Object Category: person

Part Taxonomy: body

Attribute: color, position ...

Caption: region short / detailed
description

SAM "Segment Everything" Predictions



No semantic information

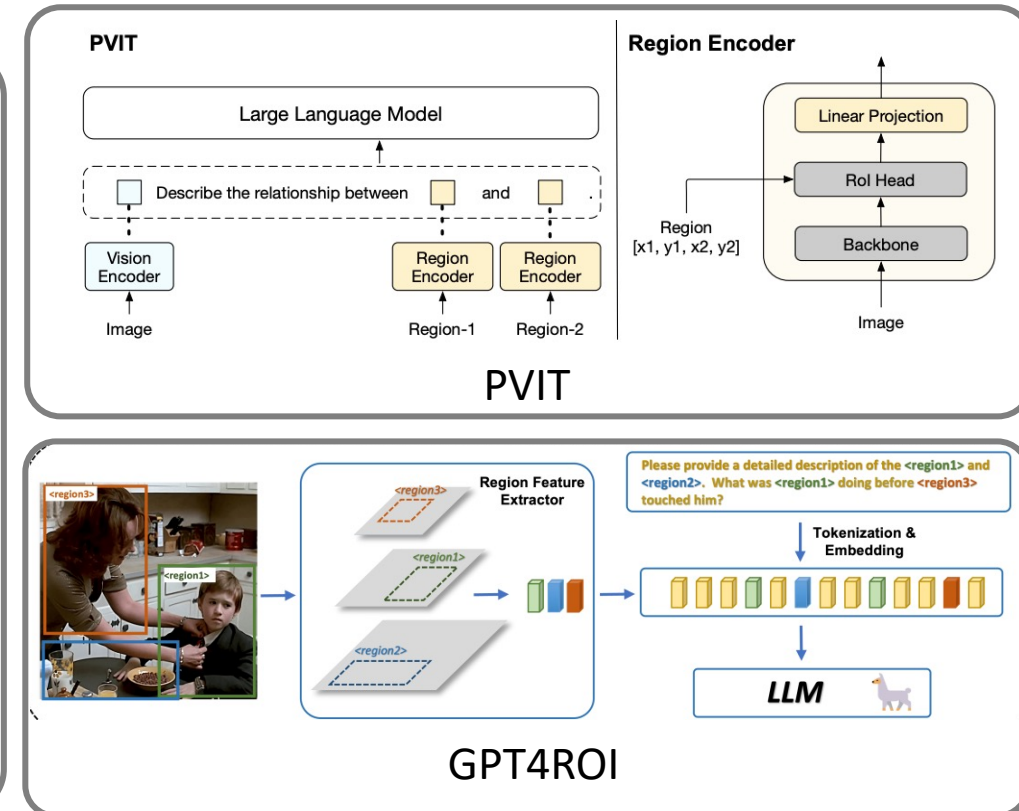
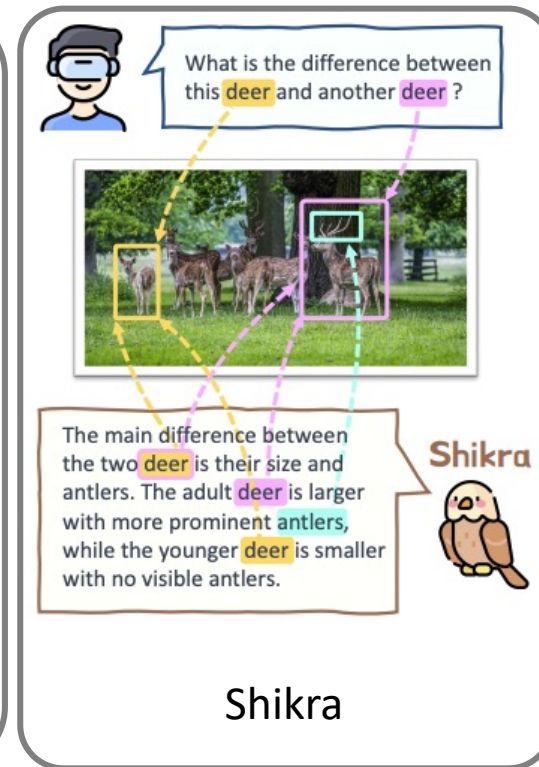
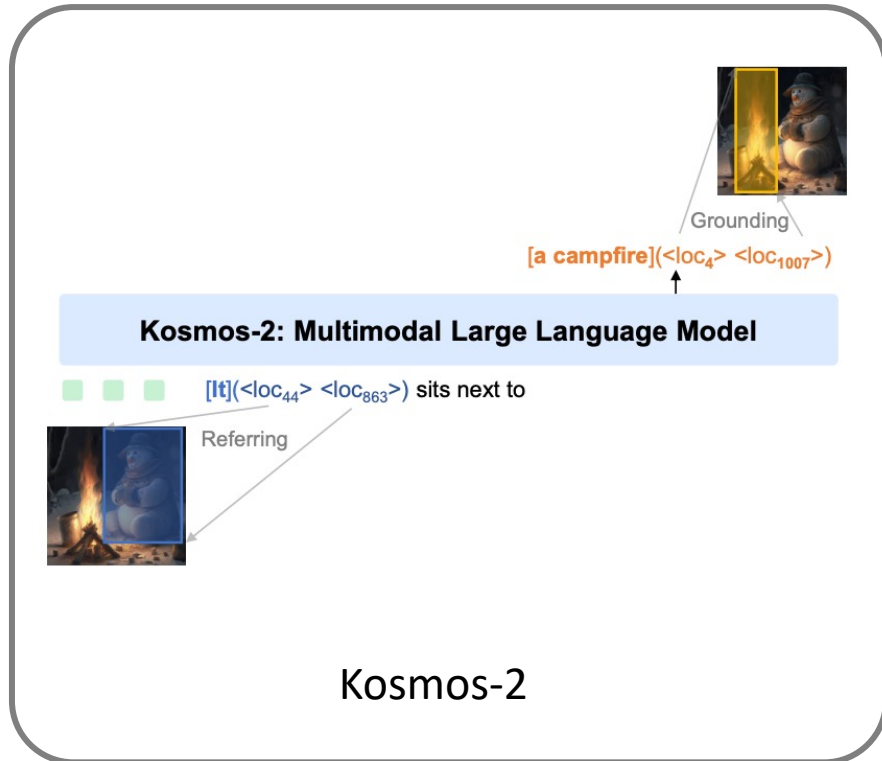
Fine-grained Region/Pixel Understanding



**Rich semantic information
containing different granularities**

Background — Multimodal Large Language Models

Region-level understanding



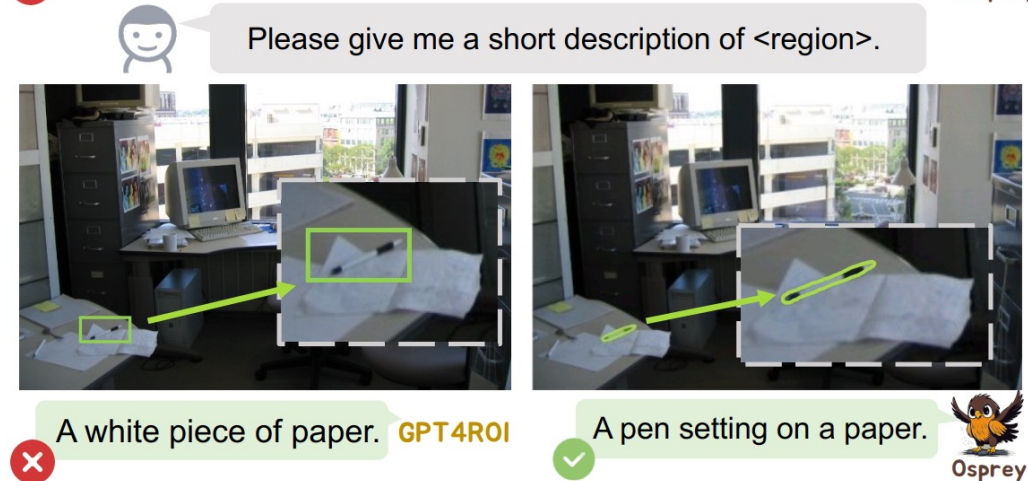
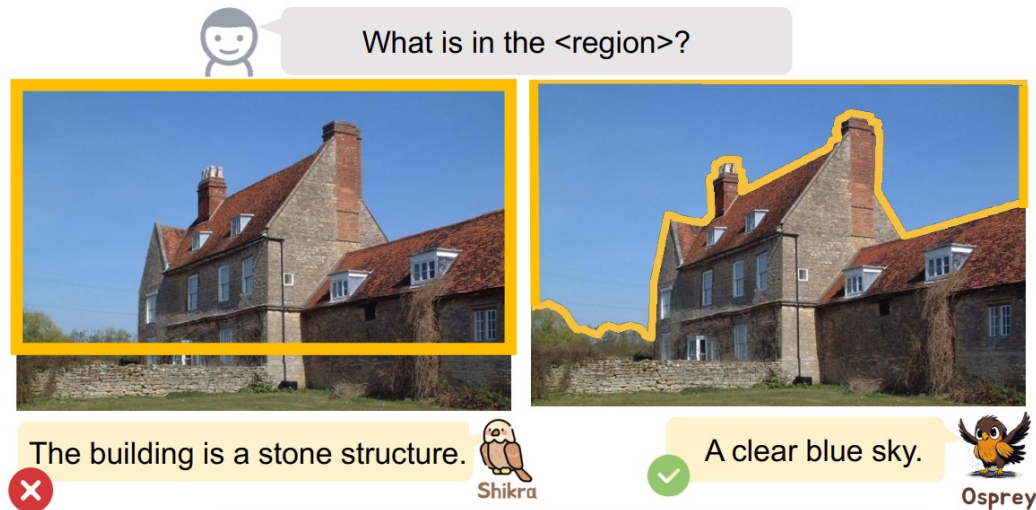
[Peng et al. 2023] Kosmos-2: Grounding multimodal large language models to the world.

[Chen et al. 2023] Shikra: Unleashing multimodal llm's referential dialogue magic.

[Chen et al. 2023] Position-enhanced visual instruction tuning for multimodal large language models.

[Zhang et al. 2023] Gpt4roi: Instruction tuning large language model on region-of-interest.

Comparisons with box-level referring



Box-level referring

- ☹️ involve **irrelevant background features** and lead to **inexact region-text pair alignment** for visual instruction tuning on LLM.
- ☹️ not be able to precisely indicate the object, resulting in **semantic deviation**.

Ours (Mask-level referring)

- 😊 Precise regional representation.
- 😊 Higher resolution for input images.
- 😊 Understand regions with various granularity.

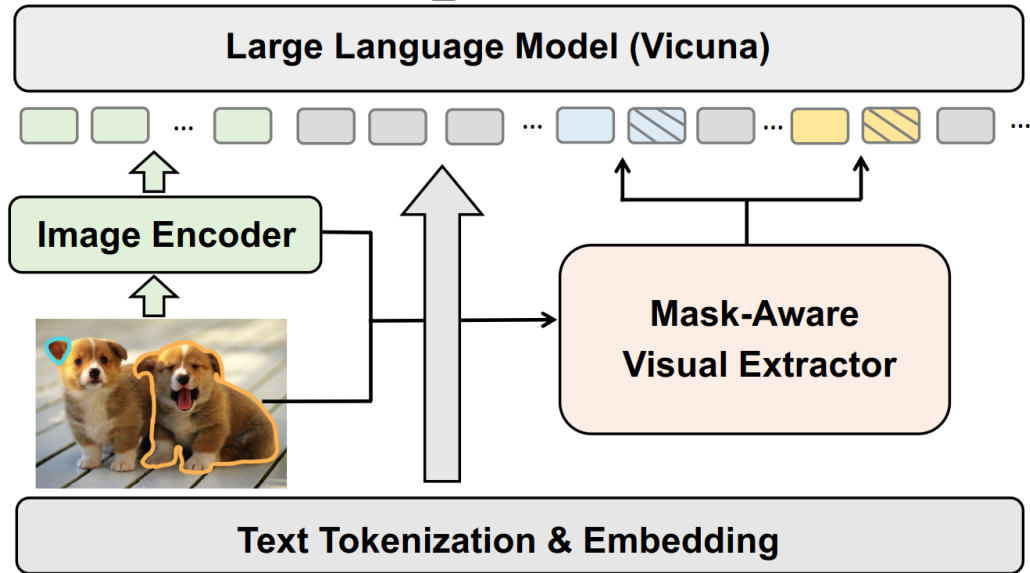
Architecture

Osprey

Answer1: A little dog sticks out its tongue and looks forward happily.
Answer2: An ear of a dog.



Osprey

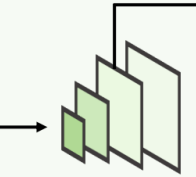


`<image>` This provides an overview of the image.
Question1: Can you give me a short description of `<region1>`?
Question2: What is in `<region2>`?

Image Encoder



Visual Encoder
(ConvNeXt-L)



Multi-level features

VL connector
(MLP)



Image token

Mask-Aware Visual Extractor

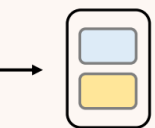


Multi-level features

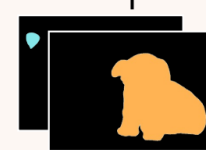
Mask Pooling

Linear Projection
&
Add

MLP

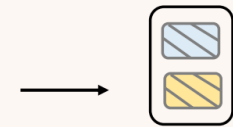


Mask token



Masks

Location Extractor
(Flatten & Projection)



Spatial token

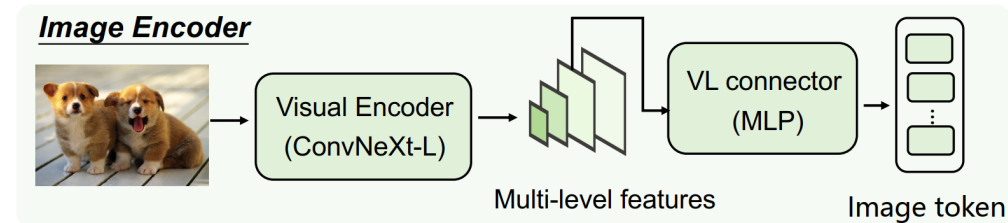
Method

Convolutional CLIP Vision Encoder

- ConvNeXt-Large CLIP model
- 512x512 input image size
- Multi-level image features
- Adopt the “res4” stage as the image-level features

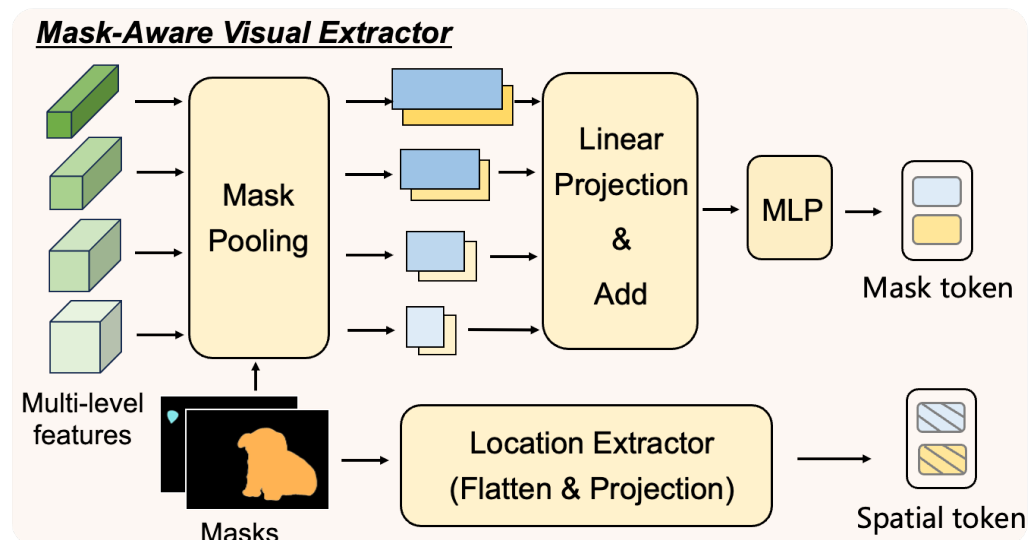
Mask-Aware Visual Exactor

- Mask Token
 - Mask Pooling: $V_{ij} = \mathcal{MP}(\mathbf{R}_i, \mathbf{Z}(x)_j)$.
 - $t_i = \sigma(\sum_{j=1}^4 \mathbf{P}_j(V_{ij}))$
- Spatial Token
 - Resize 224, flatten and project to \mathcal{S}_i .



CLIP Vision Encoder	224	448	672	896	1120
ViT-Surgery-L [27]	26.52	28.15	27.26	25.18	24.61
ConvNeXt-L [31]	23.35	34.36	40.57	43.04	43.33

Table 6. Panoptic segmentation comparisons (PQ) using different vision encoders with different input sizes on ADE20K-150 [54]. The ground truth mask is used for recognition evaluation.



How to create mask-text instruction data?

Current Data

COCO、RefCOCO、RefCOCOg、RefCOCO+ ...

- Segmentation & Detection
- Categories
 - person
- Mask region captions
 - the lady with the blue shirt.
 - the back of an older woman with her hair in a barrette with a blue jacket on.
 - navy blue shirt.
 - woman back in blue.
 - a woman is wearing blue sweater.

Expected data

Question: Can you give me a detailed description of <region>?

Answer: The woman in <region> is wearing a gray shirt, jeans, and glasses. She has short hair and a watch on her wrist. She is shopping in the lively market, browsing through the fresh produce, such as fruits and bags.

Question: What color is the woman's sweater in <region>?

Answer: Blue.

Question: Please describe the atmosphere of the market.

Answer: The market has a bustling and lively atmosphere. There's a sense of excitement and energy as people gather around the fresh fruits and bags for sale. The variety of items available, as well as the interactions among customers, contribute to the vibrant ambience of the market scene.



Mask-text instruction data creation

Use strong LLM-based teacher GPT-4



Image-level description (LLaVA-115K)

The image presents a lively market scene with a group of people buying fruits and bags. There are multiple individuals in the market, all browsing through the fresh produce available...

Categories & bboxes (COCO)

person: [0.507,0.409,0.698,0.740], person: [0.243,0.496,0.558,0.746],
handbag: [0.473,0.110,0.607,0.201], orange: [0.761,0.537,0.820,0.569],
...

Mask region captions (RefCOCO、 RefCOCO+、 RefCOCOg)

<region1>: gray shirt wearing glasses.
<region2>: the lady with the blue shirt.
<region3>: a woman in glasses shops in an open air fruit market.

different prompts



GPT-4



different types of responses

Overview of Osprey-724K dataset

Object Categories

- 1 dog
- 2 dog collar
- 3 tag
- 4 chair
- 5 cushion
- 6 dinner table
- 7 cookie
- 8 cookie
- 9 cake
- 10 box

Part Attributes

11 Q: What is the category?
A: Side of a box.
Q: What is the material?
A: Paper.

12 Q: Is the inner side of the box plain in <region12> ?
A: Yes.
Q: What is the color?
A: White.

Detailed Descriptions

4 The chair behind the dog has a brown color and appears to be made of wood. It has a simple design that complements the cushions' blue color on which the dog is resting.

9 In front of the dog and on the dining table, there are several decorative dog treats and cakes. One of them looks like a big white bone, while others have dots on them. There is also a dotted item in front of the treats on the left side.

...

Robustness and Flexibility

1 Q: Cat is the category of <region1>, right?
A: No, there is no cat in the region.

2 Q: Is the category of <region2> dog?
A: No, it isn't.

4 Q: Where is the wood part of the chair located in <region4>?
Answer the question using a single word or phrase.
A: Behind the dog.

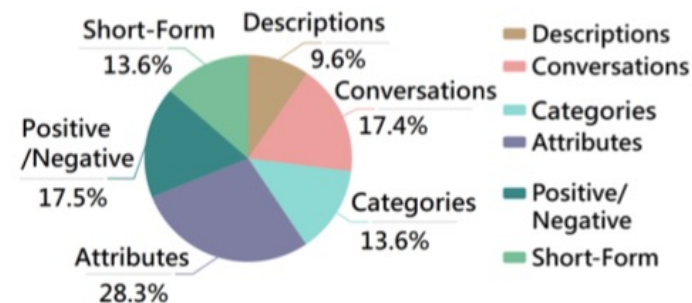
Conversations and Complex Reasoning

Question: Can you describe the pillow that the dog is resting on in <region5>?
Answer: The pillow in <region5> is a blue cushion or pillow that is part of the chair. It provides a comfortable resting place for the dog. The cushion is visible in the bottom part of the image and supports the dog's brown paws as it rests on the chair.

Question: What can you tell me about the items on the cake in <region7>, <region8> and <region9>?
Answer: In all regions, there are treats or cakes displayed. In <region9>, there is a plated item with dots or a pattern on it, placed above a big white bone. It seems to be a celebratory arrangement of doggy treats and cakes. Meanwhile, in <region7> and <region8>, there are two small dog treats or bones sitting on a plate. These treats may be specially made for dogs with some writing or decoration on them.

Question: Can you provide any more details about the dog and its surroundings based on the descriptions?
Answer: The dog in the image is black and brown and is sitting on a blue chair cushion, as seen in <region5>. The dog is surrounded by celebratory doggy treats and cakes, as mentioned in <region7> and <region8>. The presence of the treats and the dog's position on the chair suggests a special occasion or celebration. The overall theme of the image seems to be centered around the dog and its enjoyment of the treats and cakes.

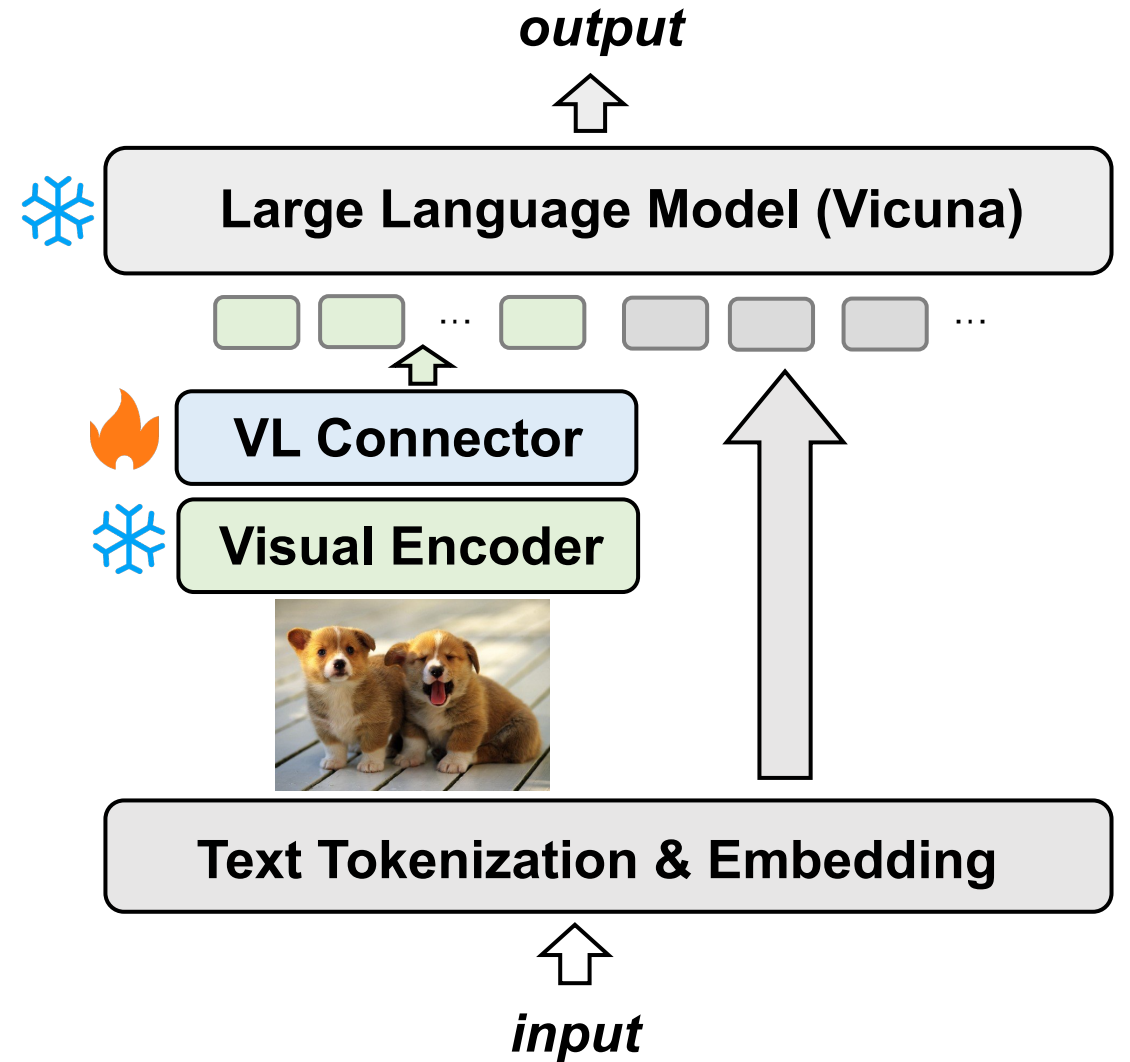
Type	Form	Raw Data	GPT-4	#Samples
Object-level	Descriptions	COCO/RefCOCO/RefCOCO+/ RefCOCOg/LLaVA-115K	✓	70K
	Conversations		✓	127K
Part-level	Categories	PACO-LVIS	✓	99K
	Attributes		✓	207K
Robustness & Flexibility	Positive/Negative	COCO/RefCOCO/RefCOCO+/ RefCOCOg/LLaVA-115K/LVIS	✗	64K/64K
	Short-Form		✓	99k



Training Stage 1: Image-Text Alignment Pre-training

Target: Train the image-level feature and language **connector** for **image-text feature alignment**.

Data: a filtered **CC3M subset (595K)**
introduced in LLaVA^[1].



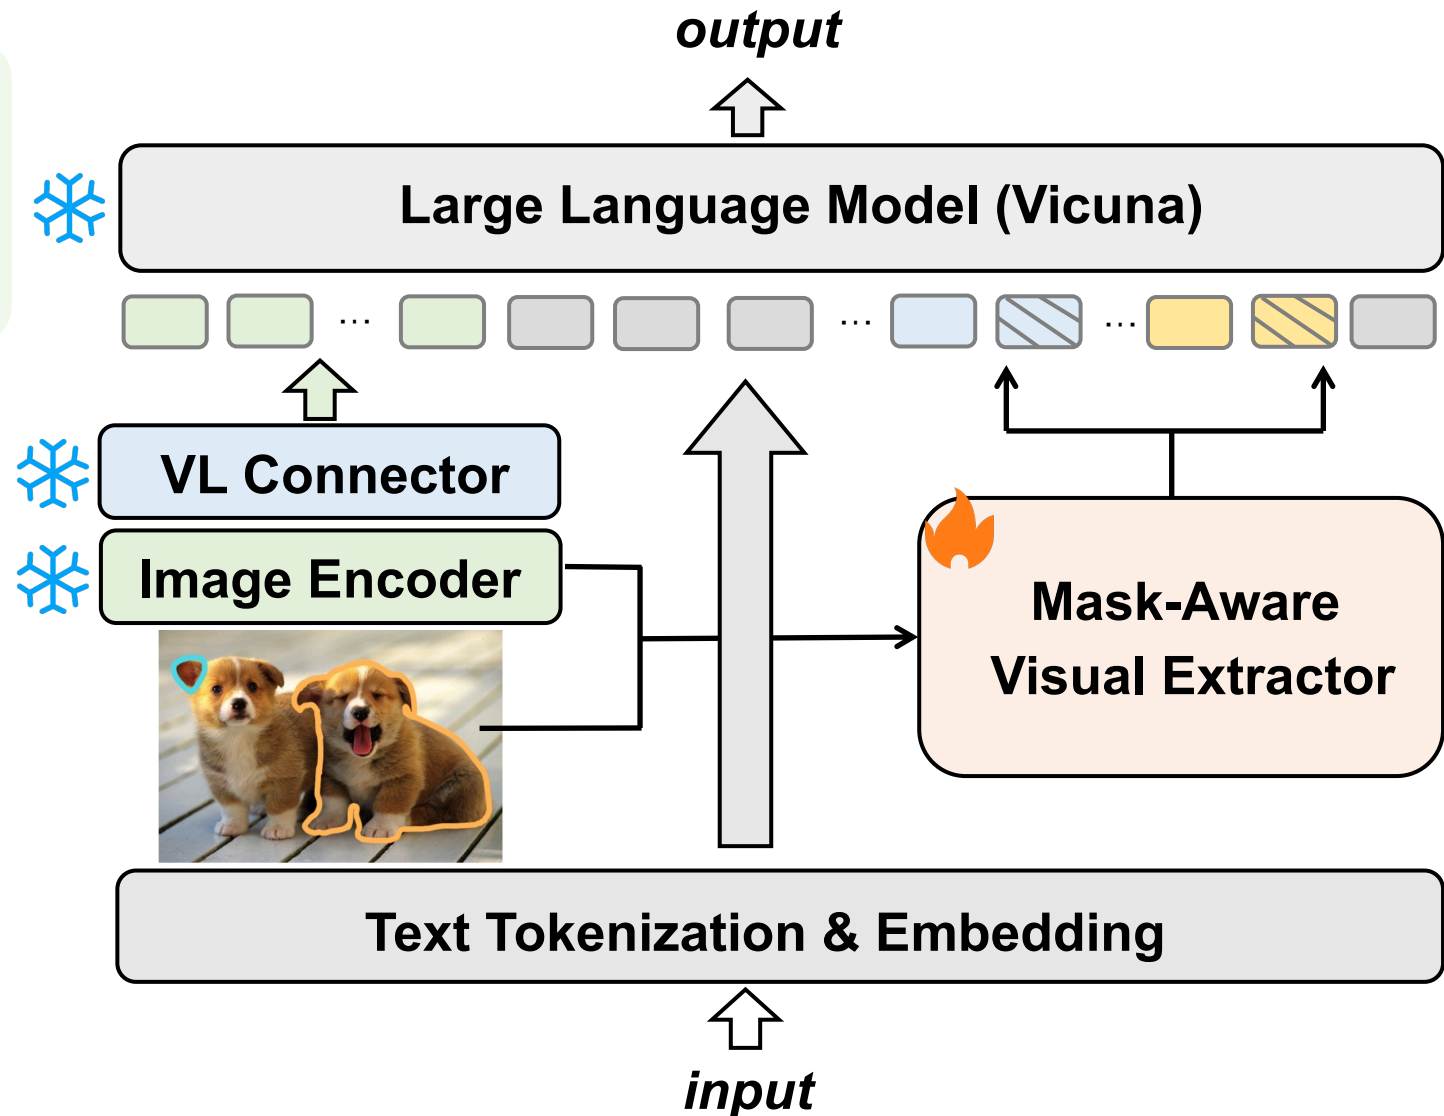
[1] [Liu et al. 2023] Visual instruction tuning.

Training Stage 2: Mask-Text Alignment Pre-training

Target: Align mask-based region features with language embeddings.

Data:

- **Object-level datasets** (COCO, RefCOCO, RefCOCO+)
- **Part-level datasets** (Pascal Part, Part Imagenet)

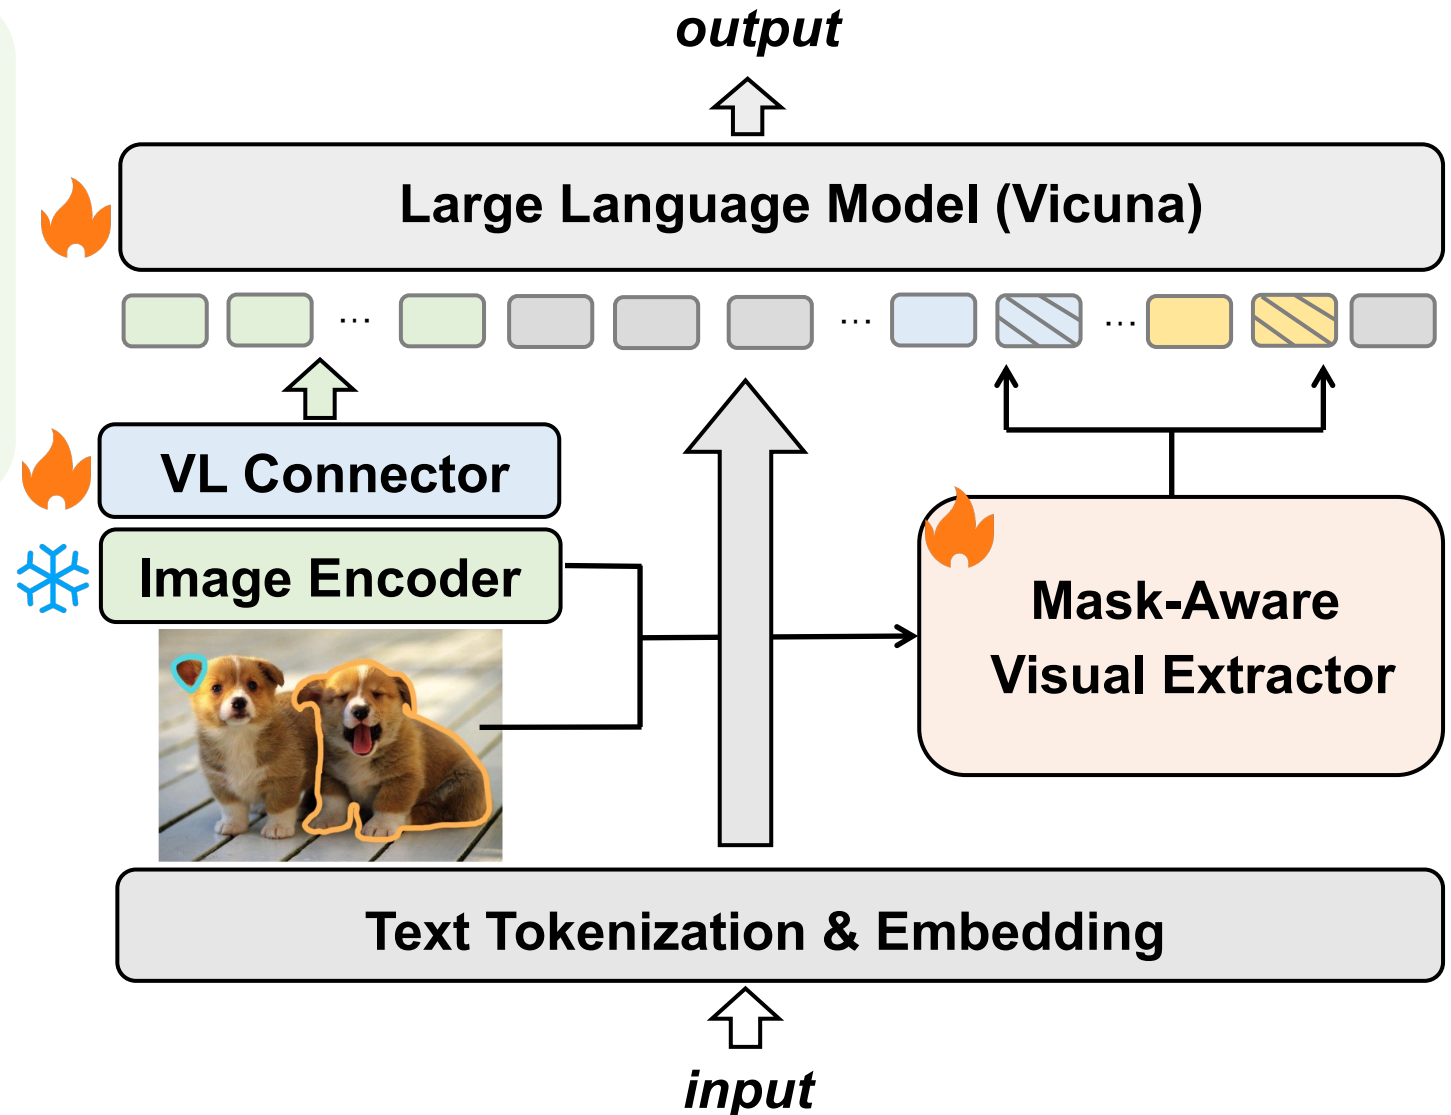


Training Stage 3: End-to-End Fine-tuning

Target: Extending the capability of Osprey to **follow user instructions** and tackle complex pixel-level region understanding tasks.

Data:

- Osprey-724K
- Visual Genome (VG)
- Visual Commonsense Reasoning (VCR)



Experiments



Strong Region Understanding ability

- Open-vocabulary Segmentation

Method	Type	Cityscapes			ADE20K-150		
		PQ	AP	mIoU	PQ	AP	mIoU
CLIP-ConvNeXt-L [43]	Mask	22.53	12.07	23.06	36.86	39.38	28.74
CLIP-Surgery-ViT-L [30]	Mask	27.24	28.35	21.92	26.55	29.70	21.42
Kosmos-2 [40]	Box	12.09	9.81	13.71	6.53	4.33	5.40
Shikra-7B [5]	Box	17.80	11.53	17.77	27.52	20.35	18.24
GPT4RoI-7B [58]	Box	34.70	21.93	36.73	36.32	26.08	25.82
Ferret-7B [54]	Mask	35.57	26.94	38.40	39.46	29.93	31.77
Osprey-7B (Ours)	Mask	50.64	29.17	49.78	41.89	41.24	29.63

- Referring object classification

Method	LVIS		PACO	
	Semantic Similarity	Semantic IoU	Semantic Similarity	Semantic IoU
LLaVA-1.5 [29]	48.95	19.81	42.20	14.56
Kosmos-2 [37]	38.95	8.67	32.09	4.79
Shikra-7B [5]	49.65	19.82	43.64	11.42
GPT4RoI-7B [53]	51.32	11.99	48.04	12.08
Ferret-7B [49]	63.78	36.57	58.68	25.96
Osprey-7B (Ours)	65.24	38.19	73.06	52.72

- Detailed region description

Method	Detailed Description
LLaVA-1.5 [32]	71.11
Kosmos-2 [40]	40.89
Shikra-7B [5]	40.97
GPT4RoI-7B [58]	49.97
Osprey-7B (Ours)	77.54
Osprey-7B* (Ours)	83.78

Experiments

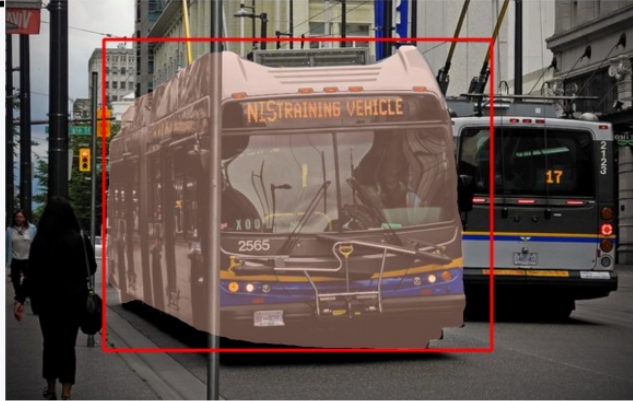


Effective relief from hallucinations

- Object Hallucination Benchmark

Sampling	Metrics	Osprey-7B*	Ferret-7B	Shikra-7B	LLaVA-1.5	InstructBLIP	MiniGPT4	MM-GPT	mPLUG-Owl
Random	Accuracy	89.47	90.24	86.90	88.73	88.57	79.67	50.10	53.97
	Precision	93.40	97.72	94.40	88.89	84.09	78.24	50.05	52.07
	Recall	84.93	83.00	79.26	88.53	95.13	82.20	100.00	99.60
	F1 Score	88.97	89.76	86.19	88.71	89.27	80.17	66.71	68.39
	Yes (%)	45.47	43.78	43.26	49.80	56.57	52.53	99.90	95.63
Popular	Accuracy	87.83	84.90	83.97	85.83	82.77	69.73	50.00	50.90
	Precision	89.94	88.24	87.55	83.91	76.27	65.86	50.00	50.46
	Recall	85.20	80.53	79.20	88.67	95.13	81.93	100.00	99.40
	F1 Score	87.50	84.21	83.16	86.22	84.66	73.02	66.67	66.94
	Yes (%)	47.37	45.63	45.23	52.83	62.37	62.20	100.00	98.57
Adversarial	Accuracy	85.33	82.36	83.10	72.10	65.17	79.20	50.00	50.67
	Precision	85.43	83.60	85.60	74.69	65.13	61.19	50.00	50.34
	Recall	85.20	80.53	79.60	88.34	95.13	82.93	100.00	99.33
	F1 Score	85.31	82.00	82.49	80.94	77.32	70.42	66.67	66.82
	Yes (%)	49.87	48.18	46.50	59.14	73.03	67.77	100.00	98.67

Comparison with other approaches Detailed and accurate region description



Osprey A bus with the **number 2565** on it is parked on the left side of the street, next to a bus with the **number 17** on it. The bus with the number 2565 is the closest to the viewer, and it has **a yellow sign** on the front.

Kosmos A bus and a truck driving down a street.

Shikra Bus with "training vehicle" written on it.

GPT4RoI A bus heading to vancouver.

LLaVA-1.5 The image depicts a city street with a large black and white bus driving down the road. The bus is in motion, and it is passing by a bus stop where a person is waiting. The bus stop is equipped with a bus stop sign, which is visible in the scene. The street is lined with buildings, creating an urban atmosphere. There is also a traffic light present in the scene, indicating that the bus is traveling on a busy street with controlled intersections.

Visualizations from online Demo



Fine-grained Understanding ability



Description(Osprey)

red light on top of police car.



Description(Osprey)

A dinosaur is seen in the background, partially hidden behind the main dinosaur. It appears to be a smaller dinosaur, possibly a child, and it's partially obscured by the larger dinosaur in the foreground.

point-prompt box-prompt segment everything

Input image

Generate segmentation and description

Clear Image

Segmentation

Description(Osprey)

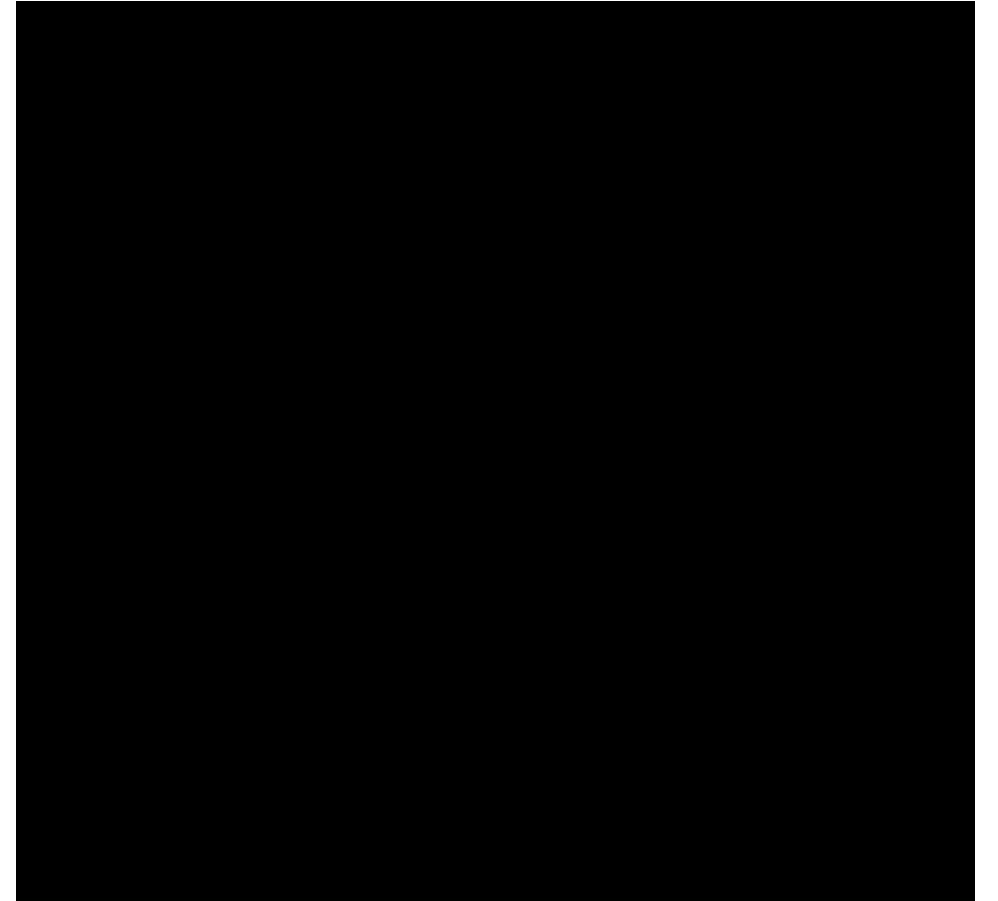
Combining with SAM

 Powerful generalization ability

Raw Image



Video Demo



A part of *Along the River During the Qingming Festival*

清明上河图

Conclusion

- Osprey
 - a novel approach to incorporate pixel-level mask region references into language instructions.
 - seamlessly integrate with SAM to generate the semantics.
- Osprey-724K
 - Contains data of different granularities and levels of detail.
 - Robustness & Flexibility.
- And open-source!



浙江大學
ZHEJIANG UNIVERSITY



蚂蚁集团
ANT GROUP



Microsoft



THE HONG KONG
POLYTECHNIC UNIVERSITY
香港理工大學

Thanks for your attention!



Demo, code, data, model

can be found at:

<https://github.com/CircleRadon/Osprey>