Seeing Beyond Classes: Zero-Shot Grounded Situation **Recognition via Language Explainer**

Anonymous Authors

ABSTRACT

2

8

9

10

11

13

14

15

16

17

18

19

20

21

22

23

24

25

26

27

28

29

30

31

32

33

34

35

36

37

38

39

40

41

42

43

44

45

46

47

48

49

Benefiting from strong generalization ability, pre-trained visionlanguage models (VLMs), e.g., CLIP, have been widely utilized in zero-shot scene understanding. Unlike simple recognition tasks, grounded situation recognition (GSR) requires the model not only to classify salient activity (verb) in the image, but also to detect all semantic roles that participate in the action. This complex task usually involves three steps: verb recognition, semantic role grounding, and noun recognition. Directly employing class-based prompts with VLMs and grounding models for this task suffers from several limitations, e.g., it struggles to distinguish ambiguous verb concepts, accurately localize roles with fixed verb-centric template¹ input, and achieve context-aware noun predictions. In this paper, we argue that these limitations stem from the model's poor understanding of verb/noun *classes*. To this end, we introduce a new approach for zero-shot GSR via Language EXplainer (LEX), which significantly boosts the model's comprehensive capabilities through three explainers: 1) verb explainer, which generates general verb-centric descriptions to enhance the discriminability of different verb classes; 2) grounding explainer, which rephrases verb-centric templates for clearer understanding, thereby enhancing precise semantic role localization; and 3) noun explainer, which creates scene-specific noun descriptions to ensure context-aware noun recognition. By equipping each step of the GSR process with an auxiliary explainer, LEX facilitates complex scene understanding in real-world scenarios. Our extensive validations on the SWiG dataset demonstrate LEX's effectiveness and interoperability in zero-shot GSR.

CCS CONCEPTS

• Computing methodologies → Scene understanding.

KEYWORDS

Zero-Shot GSR, Vision-Language Model, Large Language Model

INTRODUCTION 1

Traditional recognition methods heavily rely on the quality of vast dataset annotations, often limited in generalizing across diverse or unseen scenarios [35]. Zero-shot learning emerges to enable models to identify classes they have never been directly trained

50 Unpublished working draft. Not for distribution. 51 52 53



Figure 1: Illustration of the straightforward pipeline for Zero-Shot GSR. 1) Verb Recognition: utilizing VLMs to identify verbs via verb class-based prompts. 2) Semantic Role Grounding: employing grounding models to localize semantic roles based on the verb-centric template¹. 3) Noun Recognition: applying VLMs for identifying entities within localized semantic roles through noun class-based comparison.

on, thus greatly enhancing their universality in real-world settings [38, 43, 44]. Recently, advancements in pre-trained visionlanguage models (VLMs) [19, 20, 27], e.g., CLIP [37], have shown excellent generalization capabilities across various recognition tasks, setting new benchmarks in zero-shot learning. Particularly, CLIP incorporates dual encoders: an image encoder and a text encoder. The former processes visual input into visual features, and the latter translates texts into semantic features. This architecture facilitates the alignment of visual and text data within a unified semantic space. Leveraging class-based prompts like "A photo of [NOUN CLASS]" or "A photo of [VERB CLASS]", CLIP effectively compares images against prompts in the learned semantic space, enabling the zero-shot recognition of both unseen objects [37] and actions [29].

Nevertheless, training-free zero-shot grounded situation recognition (ZS-GSR) presents a more intricate challenge than basic object or action recognition [36]. It requires not only identifying actions (verbs) depicted in images but also discerning and localizing the semantic roles involved [36, 42]. Such a comprehensive task contributes to effectively understanding what is happening (e.g., buying in Figure 1), who is involved (e.g., AGENT and SELLER), where it is taking place (e.g., PLACE), etc. Consequently, ZS-GSR extends beyond simple class recognition to a structured understanding of scenes, requiring accurate modeling of the relationships among objects within an event. As illustrated in Figure 1, a typical GSR approach typically consists of three steps:

• Verb Recognition. This step aims to identify the verb category of the entire image. The most common zero-shot method for this sub-task uses a set of class-based prompts "a photo of [VERB

114

115

116

59 60

61 62

63 64

¹The class-contained/verb-centric template contains both verb class and its associated semantic roles, which is provided in the dataset.

⁵⁴

⁵⁵

⁵⁶

⁵⁷

⁵⁸

Anonymous Authors



Figure 2: The limitations of class-based prompts for zero-shot GSR. (a) Ambiguous Action Concepts: the verb "studying" is mistakenly identified as "coloring" due to an unclear verb meaning. (b) Constrained Role Grounding: Rigid templates misguide the grounding of "TOOL" in a complex scene. (c) Context-agnostic Noun Prediction: the noun "woman" is incorrectly classified as "hand" without considering semantic role context.

CLASS]" for each verb category [37]. These prompts are then fed into the text encoder to obtain semantic embedding, while the whole image is passed through an image encoder to get visual embedding. By comparing these two embeddings in the same semantic space, CLIP can achieve zero-shot verb recognition.

- Semantic Role Grounding. Given the predicted verb class and its predefined template¹, this step aims to precisely localize each semantic role within the image. Existing methods for this subtask often adopt an open-world grounding model, *e.g.*, Grounding DINO [30], with an object-contained grounding language as input [17, 49, 54]. As for GSR, the input can be the provided templates in the dataset, *e.g.*, the template of *buys*: "the AGENT buys GOODS with PAYMENT from the SELLER in a PLACE".
- Noun Recognition. This step aims to identify the noun category for each semantic role. Similar to the verb, a most direct zero-shot method can use noun class-based prompts like "A photo of {NOUN CLASS}" to compare with visual features of grounded semantic role for classification [37]. For example in Figure 1, the grounded semantic role *AGENT* is classified as *woman*.

However, these simple methods that rely entirely on class-based prompts at every step are limited by: 1) Ambiguous Action Con-cepts: Foundational VLMs primarily concentrate on understanding images or bag of entities, often overlooking the semantics and structures inherent in actions [10, 26, 28]. Directly classifying verbs with class-based prompts may not fully capture the nuanced meanings of actions, leading to CLIP misunderstanding the verb's embedding and generating inaccurate predictions. For instance in Figure 2 (a), the model projects the two embeddings into adjacent locations in semantic space and recognizes an activity as "coloring" due to the colorful background. It overlooks the actual activity of "studying" that entails a complex interaction with stationery and books. 2) Constrained Role Grounding: The effectiveness of semantic role grounding within visual grounding models critically depends on the quality of the text prompt [39, 46]. The employment of fixed and

simplex templates¹ for semantic role grounding is inherently constrained by the simplified sentence structure and ambiguous verb. Once the model encounters unfamiliar or complex verb classes, it easily leads to misalignment and localization errors for the involved roles. As depicted in Figure 2 (b), with templates centered on the unfamiliar "climb" class, grounding models have difficulty in locating the "TOOL". 3) Context-agnostic Noun Prediction: Noun prediction is often conditioned on the specific verb and semantic role present within a scene [36, 42]. However, when employing class-based prompts for noun classification, CLIP tends to concentrate on the category with the highest prediction confidence. This method overlooks the crucial consideration of whether the noun accurately fits the specific role defined in the scene's template, leading to contextually inappropriate predictions. For example, in Figure 2 (c), VLMs directly recognize the "hand" that occupies the main area of the image, ignoring the characteristics of the AGENT.

In this paper, we argue that these limitations stem from the model's insufficient understanding of **individual classes**. When humans encounter an unfamiliar word, such as "soliloquy", we often turn to a dictionary to serve as an explainer, providing a clear and accessible definition that illuminates the term's meaning from a common understanding perspective. For instance, the term "soliloquy" in a theatrical script refers to "a character speaking their thoughts aloud when alone or regardless of any hearers". Fortunately, recent advancements in large language models (LLMs), *e.g.*, GPT [4], have benefited from extensive training across diverse datasets, endowing them with broad world knowledge.

Inspired by human reliance on external sources for deeper understanding, we propose the method for zero-shot grounded situation recognition via Language EXplainer (LEX). It employs LLMs to serve as explainers at crucial steps of the GSR process. Specially, for verb recognition, we devise a **verb explainer** D to prompt LLMs to generate general verb-centric descriptions. This approach provides CLIP with multiple enriched "explanations" of each verb, *e.g.*,

"books or notes displayed" in Figure 2 (a), leading to more accurate 233 verb classification. Additionally, we implement an offline descrip-234 235 tion weighting strategy that takes into account the discriminability of the verb-centric descriptions without any training data. As for 236 semantic role localization, we design a **grounding explainer** 237 to prompt the LLMs to rephrase the original verb-centric template 238 from multiple perspectives. These reformulated sentences, as seen 239 in Figure 2 (b), serving as text inputs for the grounding model, are 240 241 capable of enhancing the model's understanding of the intricate 242 relationships between semantic roles. Lastly, for noun recognition, we propose a **noun explainer** b to prompt LLMs to generate noun 243 descriptions conditioned on both verb and semantic roles. As is 244 illustrated in Figure 2 (c), noun explainer prevents the misclassi-245 fication of a noun such as "hand" when the contextual role and 246 action pertain to "woman" applying lotion, thus mitigating the risk 247 248 of contextually inappropriate classifications.

To verify the effectiveness of our LEX, we conduct extensive experiments and ablation studies on SWiG dataset. Each of our explainers can be utilized as a plug-and-play module to improve the overall performance of zero-shot GSR.

In summary, we make the following contributions in this paper:

- We propose a novel training-free zero-shot GSR framework LEX that incorporates LLMs as explainers within each process, enhancing understanding of complex visual scenes.
- (2) We introduce three explainers: verb explainer, grounding explainer, and noun explainer. Each serves as a plug-andplay module to enhance the performance of zero-shot GSR.
- (3) We devise a strategy to weight descriptions by discriminability, independent of training data.
- (4) Extensive experiments on the SWiG dataset demonstrate the effectiveness and interpretability of LEX.

2 RELATED WORK

249

250

251

252

253

254

255

256

257

258

259

260

261

262

263

264

265

266

290

Grounded Situation Recognition (GSR). Despite deep learning's 267 268 notable achievements in image classification [7, 31, 34], object detection [55], and image segmentation [6, 22], its comprehension of 269 complex scenes remains limited. Techniques, e.g., scene graph gen-270 271 eration [25, 45, 52] and human-object interaction detection [1, 5, 16], aim to parse scene contents via relation graphs. Admittedly, these tasks can provide structured visual scene representations to assist 273 downstream tasks [23]. However, they only focus on modeling with 274 dyadic relationships between a subject and an object, ignoring the 275 diversity of roles in visual events [8]. GSR emerges as a solution 276 277 for holistic scene understanding by predicting actions (verbs) and 278 detecting noun entities of each semantic role [36]. Pratt et.al. [36] 279 first proposed an RNN-based two-stage framework, which detects verbs in the first stage and predicts the nouns with bounding boxes 280 in the second stage. Subsequent methods[8, 10, 42] improved this 281 two-stage pipeline by employing a transformer-based model and 282 considering semantic relations among verbs and semantic roles. 283 284 However, previous efforts that depend on annotated training samples may face challenges like sample noises [23, 24] and long-tailed 285 distribution [9, 40, 48], potentially limiting real-world applicability. 286 Our approach predicts verbs, nouns, and their associated bounding 287 288 boxes without training data, showcasing robust generalization and interpretability in real-world scenarios. 289

Foundation Models. Foundation models are typically pre-trained on extensive training data. Owing to their strong generalization capabilities, foundation models are often applied to a variety of downstream tasks [3]. Particularly in the natural language processing (NLP) area, large language models (LLMs) [11, 12, 18], such as GPT-3 [14], PaLM [2], OPT [53], and LLaMA [41], trained on extensive web-scale text datasets, exhibit formidable capabilities ranging from text prediction to contextually relevant text generation. In the field of cross-modal research, vision-language models (VLMs), e.g., CLIP [37] and ALBEF [21], have been trained on extensive image-text pairs via contrastive learning, facilitating the cohesive integration of information between these modalities. Subsequent models, such as BLIP [20] and BLIP-2 [19], leverage diverse datasets and Q-Former's trainable query vectors to further refine this alignment. Despite the remarkable transferability of VLMs, prevailing approaches often rely on class-based hard prompts as input to the text encoder, leading to a narrow focus on specific visual features and disregarding contextual information in the surroundings. In this paper, we employ various prompts to guide LLMs as different explainers to generating "class explanations", significantly enhancing VLM's understanding of intricate scenes.

3 ZS-GSR VIA LANGUAGE EXPLAINER

Formulation. Given an image *I*, zero-shot GSR aims to identify a structured frame $\mathcal{F}_v = \{y_v, \mathcal{FR}_v\}$, where $y_v \in \mathcal{V}$ denotes the salient verb category, $\mathcal{FR}_v = \{(r, y_n, \mathbf{b}) \mid r \in \mathcal{R}_v\}$ represents the set of filled semantic roles. The \mathcal{R}_v represents the set of predefined semantic roles associated with verb-centric template TP_v . Each role $r \in \mathcal{R}_v$ is filled with a noun $y_n \in \mathcal{N}$ grounded by a bounding box $\mathbf{b} \in \mathcal{B}$. For instance in Figure 1, the frame can be detected as $\mathcal{F}_v =$ {*buying*, {(*AGENT*, *woman*, \Box), (*GOODS*, *book*, \Box), (*PAYMENT*, *cash*, \Box), (*SELLER*, *man*, \Box), (*PLACE*, *store*, \Box)}.

Baseline for Zero-Shot GSR. As mentioned above, the baseline involves three steps: 1) Verb Recognition, the class-based prompts is fed into text encoder $T(\cdot)$ of CLIP to obtain verb text features $\{\mathbf{t}_v\}$. Similarly, the entire image is fed into image encoder $V(\cdot)$ to extract visual feature \mathbf{v}_v . Subsequently, the cosine similarity between \mathbf{t}_v and \mathbf{v}_v is computed across different verb categories for final prediction. 2) Semantic Role Grounding, Grounding DINO integrates a verb-centric template TP_v as text input and image I as visual input to generate a series of candidate bounding boxes with semantic role labels and scores. The bounding boxes with the highest score for each semantic role are selected. 3) Noun Recognition, CLIP compares cropped region feature $\{\mathbf{v}_n^r\}$ of each role r with the text features $\{\mathbf{t}_n\}$ of class-based prompts for noun classification.

To address the constraints of the baseline, we propose a new framework LEX for ZS-GSR. As illustrated in Figure 3, it comprises three components: verb recognition via verb explainer, role localization via grounding explainer, and noun recognition via noun explainer. LEX enhances the model's understanding of different classes by adding an auxiliary explainer at each step.

3.1 Verb Recognition via Verb Explainer

This component aims to recognize the salient verb category within an image. It comprises three steps: verb-centric description generation, description weighting, and verb classification.

Anonymous Authors



Figure 3: The framework of LEX. 1) Verb Recognition via Verb Explainer: generate general verb-centric descriptions to recognize verbs. 2) Role Localization via Grounding Explainer: generate a rephrased verb-centric template to localize semantic roles. 3) Noun Recognition via Noun Explainer: generate scene-specific noun descriptions to predict nouns.

3.1.1 Verb-Centric Description Generation. To enhance the discriminative capacity of the CLIP model in distinguishing verb categories, we introduce a verb explainer to "explain" basic class-based prompts for each category. Inspired by zero-shot image classification [32], we introduce a verb-centric description generation prompt to make LLMs as verb explainer to generate general verb-centric descriptions $\mathcal{D}_v = \{d_v\}$ for verb class $y_v \in \mathcal{V}$ from multiple perspectives:

$$\mathcal{D}_v = LLM \text{ (in-context examples, } y_v \text{, instruction)},$$
(1)

prompt input

where $LLM(\cdot)$ is the decoder of the LLM. The verb-centric description generation prompt input consists of three parts²: 1) *In-context* examples, the description instruction for the verb, along with some examples of the generated results. This part adopts in-context learning [13, 33] to make LLM generate analogous results. 2) *Verb class* y_v , the target category that needs to generate the description. 3) *Instruction*, the sentence used to command the LLM to generate the description, *e.g.*, "what are the useful visual features for the event of 'rehabilitating': AGENT rehabilitates ITEM at a PLACE".

These generated descriptions highlight various unique visual features associated with specific action categories, for example, "looking around for searching goods", "pushing a shopping cart or carrying a basket", contribute to enhancing the distinguishability between "*shopping*" and other similar actions like "*buying*".

3.1.2 Description Weighting. Considering the variances in description quality and its different impact on verb classification, we devise a description weighting strategy. Recent methods try to select descriptions from the perspective of the discrimination of image



Figure 4: An example of using "biting"'s scene text as "image" to calculate the discriminability score.

features [47] and the coverage among concepts [50]. However, these methods require a large number of annotated training images and are prone to overfitting to seen scenarios, which are not suitable for training-free zero-shot GSR. Thanks to the property of VLMs to align visual and text modes in a shared space, we can *replace an annotated image with a complex scene description that contains the verb category* [15].

Specifically, we first utilize the LLMs to generate a set of complex scene descriptions $\mathcal{D}_{s|v} = \{d_s\}$ for each verb $y_v \in \mathcal{V}$. These scene descriptions are fed into the text encoder $T(\cdot)$ to generate scene text embeddings $\{t_s\}$, used as "image" for subsequent description weighting². For instance in Figure 4, the verb *biting*'s scene description is used instead of an annotated image.

Intuitively, the greater the distinctiveness of a verb description in differentiating between its own class of scenes and other classes, the more important it is. Following [47], we apply the **discriminability score** $Dis(\cdot)$ to measure the distinctiveness of a verb description. The correlation $\rho(\cdot, \cdot)$ between a verb description embedding and a

²The detailed description generation prompts are in the Appendix.

Seeing Beyond Classes: Zero-Shot Grounded Situation Recognition via Language Explainer

set of scene text embeddings with class y_v is denoted as:

$$\rho(y_v, d_v) = \frac{1}{|\mathcal{D}_s|_v|} \sum_{d_s \in \mathcal{D}_{s|v}} \phi(\mathbf{t}_s, \mathbf{t}_{d_v}), \tag{2}$$

where $\phi(\cdot, \cdot)$ denotes the cosine similarity, \mathbf{t}_{d_v} is the text embedding of a verb description d_v . The conditional likelihood of aligning scenes of a class given d_v can be written as:

$$\overline{\rho}(y_v, d_v) = \frac{\rho(y_v, d_v)}{\sum\limits_{\substack{y'_v \in \mathcal{V}}} \rho(y'_v, d_v)}.$$
(3)

The discriminability score $Dis(d_v)$ can be measured by the entropy over all verb classes, written as:

$$Dis(d_v) = -\sum_{u_v' \in \mathcal{V}} \overline{\rho}(y_v'|d_v) \log(\overline{\rho}(y_v'|d_v)).$$
(4)

The smaller the discriminability score, the greater the distinction d_v offers for its category. Hence, the weight of a general verb description can be denoted as:

$$w_v(d_v) = \frac{\exp\left(1/Dis(d_v)\right)}{\sum\limits_{\substack{d'_v \in \mathcal{D}_v}} \exp\left(1/Dis(d'_v)\right)},\tag{5}$$

where $w_v(d_v)$ represents the weight of verb description d_v .

3.1.3 Verb Classification. In this step, we compute the similarity score between visual and textual features to obtain the probability distribution of verbs. We calculate the final distribution as follows:

$$Score(y_v) = \underbrace{(1-\lambda)\phi(\mathbf{v}_v, \mathbf{t}_v)}_{\text{class-based}} + \underbrace{\lambda \sum_{d_v \in \mathcal{D}_v} w_v(d_v)\phi(\mathbf{v}_v, \mathbf{t}_{d_v})}_{\text{description-based}}, \quad (6)$$

where λ is utilized as a balance factor between classed-based prompts and description-based prompts.

3.2 Role Localization via Grounding Explainer

This module aims to locate each semantic role more accurately via an auxiliary grounding explainer. Specifically, we first employ a grounding explainer to rephrase the verb-centric templates in a more comprehensible manner. Subsequently, we use the Grounding DINO to generate candidate bounding boxes for the semantic roles.

3.2.1 Rephrased Template Generation. Due to the ambiguity of the original fixed verb-centric template, using it as grounding text input for the Grounding DINO model may result in limited or inaccurate semantic role localizations. To enhance the comprehensibility of grounding language, we introduce a grounding explainer to "explain" these fixed templates. To be specific, we employ a rephrased template generation prompt, utilizing LLM as grounding explainers to produce rephrased grounding templates $\mathcal{D}_g = \{d_g\}$ for each verb class $y_v \in \mathcal{V}$, expressed as:

$$\mathcal{D}_g = LLM \text{ (in-context examples, } TP_v \text{, instruction)}.$$
 (7)

prompt input

Similarly, the rephrased template generation prompt input also comprises three components: in-context examples, template TP_v , and instruction². The first two components are analogous to verb-centric description generation. As for the last component, the instruction sentences are designed to enable LLM to generate role-contained templates that are easy to understand, *i.e.*, "Please generate new sentences that detailedly rephrase this sentence to make it easier to understand {TEMPLATE}. Note that you must keep the original words: {SEMANTIC ROLES}".

3.2.2 Candidate Box Generation. Given an unlabeled input image *I* and the grounding text input d_g , the Grounding DINO model is adept at associating visual features with corresponding object labels (*i.e.*, semantic roles). More concretely, the Grounding DINO model outputs a sequence of bounding boxes $\mathcal{B}_g = \{\mathbf{b}\}$, associated object labels $\mathcal{R}_g = \{y_r\}$, and confidence scores $C_g = \{c\}$ as follows:

$$\mathcal{B}_g, \mathcal{R}_g, \mathcal{C}_g = DINO(I, d_g), \tag{8}$$

where $DINO(\cdot)$ is the decoder of Grounding DINO. It is possible to generate multiple boxes for each object. We select the box with the highest confidence score *c* for each object y_r that appeared in the predefined role set R_v as candidates. The chosen $\widehat{\mathcal{B}}_g$ of each grounding description d_g constitutes the candidate boxes $\widehat{\mathcal{B}} = \{\widehat{\mathcal{B}}_g | d_g \in \mathcal{D}_g\}$. These candidate bounding boxes of semantic roles are utilized for the next noun recognition.

3.3 Noun Recognition via Noun Explainer

In this module, we utilize the noun explainer to achieve contextaware noun classification. To be specific, it consists of four steps: noun filtering, scene-specific noun description generation, noun pre-prediction, and noun refinement, as displayed in Figure 5.

3.3.1 Noun Filtering. This step filters out noun categories that are unlikely to appear in the specific semantic roles from class-level. Leveraging the common sense contained in LLM, we can inquire about reasonable noun categories through a simple prompt that includes a specified semantic role². For instance, we can query "Given entity list {Noun CLASS}, which entities are most likely to be the result of a predicted semantic role PLACE" to filter unreasonable noun categories, *e.g., man.* This filtering step not only enhances the computational efficiency of subsequent noun classification by reducing the number of traversed categories but also improves the accuracy of classification results.

3.3.2 Scene-Specific Noun Description Generation. In particular, the cropped candidate box inevitably contains more than one object, and the object normally has more than one possible category. Reliance on class-based prompts often results in CLIP focusing on the most salient noun category directly, thereby ignoring the specific scene context. To address this, we incorporate a noun explainer to "explain" class-based prompts associated with each noun category within diverse scenes². This noun explainer adopts a scene-specific noun description generation prompt, guiding LLM to produce descriptions $\mathcal{D}_n^r = \{d_n\}$ that reflect the intricate context of each scene for each noun class $y_n \in \mathcal{N}$:

 $\mathcal{D}_n^r = LLM$ (in-context examples, scene, y_n , instruction). (9)

prompt input

Different from verb-centric description generation, the input prompt for scene-specific noun description generation is conditioned on extra scene information (verb-centric template and semantic role) to provide contextual information. The instruction is like this:



Figure 5: The architecture of noun recognition.

"Please describe the visual features that can distinguish the noun entity {NOUN CLASS} corresponding to {SEMANTIC ROLE} in the scene: AGENT writes on TARGET using a TOOL at a PLACE". The generated descriptions (*e.g.*, "holding a pen and other tools in hand") reflect the characteristic of corresponding semantic roles (*e.g.*, *AGENT*), thereby ensuring more accurate and contextually relevant noun categorization.

3.3.3 Noun Pre-prediction. In this step, we first compute the similarity between visual embeddings of all cropped boxes and text embedding to obtain the probability distribution of nouns. Similar to the verb classification process, this entails the incorporation of both class-based and description-based approaches to yield the pre-predict results for nouns:

$$Score(y_n) = \underbrace{(1-\lambda)\phi(\mathbf{v}_n^r, \mathbf{t}_n)}_{\text{class-based}} + \lambda \underbrace{\sum_{d_n \in \mathcal{D}_n^r} w_n(d_n)\phi(\mathbf{v}_n^r, \mathbf{t}_{d_n})}_{\text{description-based}}, \quad (10)$$

where \mathbf{t}_{d_n} is the text features of the noun description d_n . Due to the requirement of invoking a large number³ of LLMs for generating scene descriptions akin to description weighting in verb recognition, the weights $w_n(d_n)$ are uniformly assigned as $1/|\mathcal{D}_n^r|$.

3.3.4 Noun Refinement. After localization by Grounding DINO, a semantic role may correspond to multiple candidate boxes due to the integration of different $d_g \in \mathcal{D}_g$. Directly selecting the box with the highest confidence often relies solely on the localization model's understanding of the semantic role, overlooking the contextual relevance within the scene. To mitigate this, we propose to refine noun classification and bounding box selection via a global comparison. This process contrasts the visual features of the entire image \mathbf{v}_v against the text features $\{\mathbf{t}_{tp}\}$ of templates $\{\overline{TP}_v\}$ whose semantic roles are filled iteratively with predicted nouns $\{y_n\}$ in the previous step, *e.g.*, "**woman** is studying in **library**" in Figure 5. The noun filled in the template with the highest similarity along with its corresponding bounding box is the final prediction, written as:

$$\widetilde{TP}_{v}^{*} = \arg \max \phi(\mathbf{v}_{v}, \{\mathbf{t}_{tp}\})$$

$$\mathcal{FR}_{v} = \{(r, y_{n}, b) \mid y_{n} \in \widetilde{TP}_{v}^{*}\}.$$
(11)

As seen in Figure 5, the noun "*woman*" for the semantic role "*AGENT*" and its bounding box is selected as the final prediction result.

4 EXPERIMENT

4.1 Experimental Setup

4.1.1 Dataset. We evaluated our method on the SWiG dataset [36], which extends the imSitu dataset [51] by incorporating additional bounding box annotations for all visible semantic roles (69.3% of semantic roles have bounding boxes). In the SWiG dataset, each image is annotated with a verb, followed by a set of semantic roles ranging from 1 to 6 (3.55 on average), each verb is annotated with three verb frames by three separate annotators. SWiG contains 25200 testing images with 504 verb categories, 190 semantic role categories, and 9929 noun entity categories.

4.1.2 Evaluation Metrics. We employed the same five evaluation metrics as [36], which include: 1) verb: the accuracy of verb prediction. 2) value: the noun prediction accuracy for each semantic role. 3) value-all (val-all): the noun prediction accuracy across the entire set of semantic roles. 4) grounded-value (grnd): the accuracy of the bounding box for each semantic role, where the predicted bounding box must achieve an IoU value of at least 0.5 with the ground-truth bounding box. 5) grounded-value-all (grnd-all): the accuracy of bounding box across the entire set of semantic roles. Besides, these metrics were presented under three evaluation settings: 1) Top-1-verb, 2) Top-5-Verb and 3) Ground-Truth-Verb, verbs are selected based on the top-1 prediction, top-5 predictions, and corresponding ground-truth, respectively. If the predicted verbs are incorrect in the Top-1/5-verb settings, the other four metrics (value, val-all, grnd, and grnd-all) are considered incorrect.

4.1.3 *Implementation Details.* We used the vision transformer with a base configuration (ViT-B/32) as the default backbone for CLIP. We adopted GPT-3.5-turbo for LLM, recognized as a widely utilized LLM in existing works. More details are left in the **Appendix**.

4.1.4 Baselines. We compared our proposed method LEX with four strong baselines: 1) **CLS**, which uses class-based prompts to calculate the text embedding for both verb and noun classification. Leverage verb-centric template as the text input of Grounding DINO for role grounding. 2) **TEM**, which uses the verb-centric template to enhance the original class-based prompts for *verb classification*, compared with CLS. 3) **CLSDE**, which generates general descriptions of each noun category without considering scene and role by LLMs to enhance the original class-based prompts for *noun classification*, different from CLS. 4) **RECODE**, which generates semantic role descriptions as visual cues to enhance the original class-based prompts for *verb class*.

4.2 Quantitative Comparison Result

We evaluated the performance of our proposed LEX and four strong baselines on the test set of SWiG dataset. From Table 1, we have the following observations: 1) The CLS baseline relying solely on class-based prompts, demonstrates inferior performance, particularly in the prediction of bounding boxes and nouns. 2) In the Top-1 and Top-5 verb settings, TEM exhibits a slight performance advantage over CLS, attributed to the enhanced scene information

⁶³⁶ ³Since noun description is conditioned on both verb and semantic role, the call number 637 of LLMs at least is $|\mathcal{V}| \times |\mathcal{R}_v| \times |\mathcal{N}|$

Seeing Beyond Classes: Zero-Shot Grounded Situation Recognition via Language Explainer

Table 1: Evaluation results on the test set of SWiG dataset. Values in gray indicate metrics obtained by the same method as CLS.

24.1.1	Top-1-Verb					Top-5-Verb					Ground-Truth-Verb			
Method	verb	value	val-all	grnd	grnd-all	verb	value	val-all	grnd	grnd-all	value	val-all	grnd	grnd-all
CLS [37]	30.18	4.70	0.16	3.09	0.07	55.49	9.31	0.23	3.09	0.07	13.51	0.42	9.05	0.14
TEM	30.59	4.78	0.19	3.17	0.08	56.42	9.36	0.21	5.14	0.08	13.51	0.42	9.05	0.14
CLSDE [32]	30.18	4.86	0.18	3.15	0.07	55.49	9.32	0.23	5.00	0.08	13.80	0.44	9.12	0.15
RECODE [26]	30.17	4.70	0.16	3.10	0.07	55.50	9.31	0.23	5.00	0.08	13.51	0.42	9.05	0.14
LEX	32.41	9.37	1.61	7.26	0.98	58.34	17.68	3.19	13.77	2.15	29.92	4.68	23.57	3.08

Table 2: Effects of each component in verb recognition.

Verb Explainer	Weighting	Top-1-Verb verb↑	Top-5-Verb verb↑
		30.18	55.49
1		32.13	57.84
1	1	32.41	58.34

Table 3: Effects of grounding explainer.

Grounding		Ground-7	Truth-Verb	
Explainer	value↑	val-all↑	grnd↑	grnd-all↑
	13.51	0.42	9.05	0.14
1	13.85	0.48	9.86	0.16

Table 4: Effects of each component in noun recognition.

D :1/	Noun	D.C	Ground-Truth-Verb						
Filter	Explainer	Refine	value↑	val-all↑	grnd↑	grnd-all↑			
			13.51	0.42	9.05	0.14			
1			28.51	4.06	21.71	2.58			
1	1		29.16	4.33	22.11	2.67			
1		1	29.39	4.51	23.22	3.01			
 ✓ 	1	1	29.92	4.68	23.57	3.08			

for verb recognition. 3) CLSDE achieves a slight performance gain in the accuracy of noun classification because general noun descriptions often fail to apply in specific contexts. 4) RECODE fails to improve the accuracy of verb recognition, likely due to the insufficient local information provided by semantic roles (*i.e.*, not clear and detailed enough to understand actions). 5) The proposed LEX exhibits significant performance gains across all metrics compared to all baseline models, *e.g.*, **32.41**% and **58.34**% in terms of verb accuracy under Top-1-Verb and Top-5-Verb, and **29.92**% noun accuracy under Ground-Truth, respectively. This indicates the effectiveness of using explainers in zero-shot GSR.

4.3 Ablation Studies

As aforementioned, each component of our proposed LEX framework can serve as a plug-and-play module for zero-shot GSR. This section⁴ ablated all the proposed components on the test split of SWiG dataset [36].

4.3.1 Key Components in Verb Recognition. We first investigated the two major elements utilized in verb recognition: 1) Verb Explainer, which denotes using verb-centric descriptions for verb classification (Sec. 3.1.3); and 2) **Weighting**, which denotes the employment of description weighting for the text embeddings of all generated verb-centric descriptions (Sec. 3.1.2). Results are shown in Table 2. The first row corresponds to the performance of the CLS baseline. As seen, CLS achieves 30.18% top-1 and 55.49% top-5 verb accuracy. Upon applying the proposed verb explainer (the second row), we observe consistent and substantial improvements for both top-1 accuracy (30.18% \rightarrow **32.13**%) and top-5 accuracy (55.49% \rightarrow **57.84**%). This highlights the importance of our "recognition with explainer" strategy and validates the viability of our motivation. Moreover, LEX achieves better performance across all metrics with the weighting strategy. This indicates that the proposed verb explainer and weighting strategy can work synergistically.

4.3.2 *Key Component in Semantic Role Grounding.* We next studied the impact of our **Grounding Explainer**, which adopts rephrased verb-centric templates as grounding input (Sec. 3.2.1). As illustrated in Table 3, our grounding explainer proves to be effective, *e.g.*, **0.81**% accuracy improvement in terms of the grounding (grnd) metrics.

4.3.3 Key Components in Noun Recognition. We further analyzed the influence of three major components proposed for noun recognition: 1) Filter, which denotes filtering those unreasonable nouns by using LLMs (Sec. 3.3.1); 2) Noun Explainer, which denotes using scene-specific noun descriptions for noun prediction (Sec. 3.3.3); and 3) Refine, which denotes using contextual information of the scene to refine final results (Sec. 3.3.4). From the results in Table 4, we have the following findings: 1) Filtering those unreasonable noun candidates leads to significant performance gains across all noun metrics(e.g., 15% in value metrics and 3.64% in val-all metrics). 2) Substantial improvements can be made by incorporating the noun explainer which aims to provide more contextual information (e.g., $28.51\% \rightarrow 29.16\%$ in value metrics). 3) Furthermore, after incorporating the refine module, we achieve considerable gains of 0.88% in value metrics and 1.51% in grnd metrics. 4) Finally, by integrating three core components, LEX delivers the best performance across all metrics. This validates the effectiveness of our comprehensive modular designs.

4.3.4 Different Architectures of CLIP. Last, we examined the impact of the utilized CLIP's architectures. As outlined in Table 5, regardless of the visual encoder employed, our LEX demonstrates consistent and substantial improvements across all metrics. Such impressive results further verify the robustness of our approach.

4.4 Qualitative Comparison Result

We visualized CLIP's attention maps for various images and query prompts in Figure 6. As is displayed in Figure 6 (a), we can observe

⁴More ablation studies are left in the Appendix.

Table 5: Ablation studies on different architectures of CLIP.

		Top-1-Verb						Top-5-Verb				Ground-Truth-Verb			
Architecture	Method	verb	value	val-all	grnd	grnd-all	verb	value	val-all	grnd	grnd-all	value	val-all	grnd	grnd-all
	CLS	38.69	6.50	0.16	4.23	0.02	64.65	10.95	0.22	4.23	0.02	15.65	0.38	10.24	0.10
ViT-L/14	LEX	42.00	13.06	2.60	10.26	1.76	68.35	21.40	4.25	16.40	2.83	30.37	5.39	23.62	3.54
	CLS	40.04	6.69	0.17	4.41	0.05	65.67	11.03	0.19	4.41	0.05	15.38	0.31	10.49	0.11
ViT-L/14@336px	LEX	42.63	13.66	2.94	10.67	1.98	69.18	22.32	4.57	17.18	3.04	31.54	5.83	24.02	3.85
	CLS	30.18	4.70	0.16	3.09	0.07	55.49	9.31	0.23	3.09	0.07	13.51	0.42	9.05	0.14
ViT-B/32	LEX	32.41	9.37	1.61	7.26	0.98	58.34	17.68	3.19	13.77	2.15	29.92	4.68	23.57	3.08
	CLS	32.68	5.08	0.19	3.36	0.05	58.41	9.80	0.22	3.36	0.05	14.06	0.58	9.45	0.16
V11-B/16	LEX	35.11	10.03	1.93	7.74	1.19	64.18	18.29	3.71	14.05	2.44	28.96	5.01	22.72	3.16
Verb	class-based						descrip	tion-b	ased		• _				
a) GT: instructing GT: sneezing	instruct	ing: 25.0		ring: 26.3		instruct	ring 27.00		oring: 15.	34 34	tak a b scc inc as sm ha pre con hu du	cing note pook or s poring: ag licating t arms rai eezing: r ndkerchi eparedne ughing: a nched or ring the	rearby tis earby tis eff, or ob earby tis eff, or ob ess for sn agent's por r bent slig cough.	ng to son ask ques ly postur ent of sco mping. sues, sues, ojects sug eezing. osture, p ghtly forv	e poring, such ggesting ossibly vard
Noun b) GT: male child	male ch	ild: 25.2		Idler: 27.4	47	I male ch	ild 29.14		ddler: 17.	35	ma ou for too livi cu: bla ba ou lea shu	ale child: tdoors o tipping. ddler: to ng room shion fall ank: blan ckground tdoors: c ives and een on ro	male chi r in a mo ddler, end with sofils. k, genera d. butdoors, surfaces, bads.	ild, could re compl closed sp t floor cc illy pure , maybe o and a ge	l be lex indoor bace like a overings to white droplets on eneral

Figure 6: Visualization of CLIP's attention maps on input images with different prompts. (a) Examples of verb recognition. The right side shows the general verb-centric description prompts generated for each verb, which are used to visualize the corresponding attention map. (b) Examples of noun recognition. The red box indicates the bounding box corresponding to the semantic role. The right side shows the scene-specific noun description prompts generated for each noun.

that class-based approaches might focus on areas irrelevant to the query verb. For example, given the class-based prompt for the verb "*sneezing*", CLIP wrongly attends to the child's hair and arms. With the detailed guidance of verb-centric descriptions, CLIP can focus on the correct areas, which indicates the importance of providing more contextual information instead of only class names. Similar conclusions can be drawn in terms of noun recognition, as depicted in Figure 6 (b). More qualitative results are left in the **Appendix**.

5 CONCLUSION

In this paper, we proposed a novel approach LEX for zero-shot GSR utilizing large language models to provide rich, contextual explanations at each critical stage of the GSR process. By innovating beyond the constraints of conventional class-based prompt inputs, our approach leverages LLMs to clarify ambiguous actions, accurately ground semantic roles, and ensure context-aware noun identification. Furthermore, our framework enhances model interpretability and adaptability across varied visual scenes without requiring direct training on annotated datasets. Additionally, we introduced a description weighting mechanism to measure the contribution of generated descriptions and assign corresponding weights, thereby enhancing the discriminative power of the method. Extensive validation on the SWiG dataset affirmed the effectiveness and interoperability of our proposed LEX. Seeing Beyond Classes: Zero-Shot Grounded Situation Recognition via Language Explainer

ACM MM, 2024, Melbourne, Australia

987

988

989

990

991

992

993

994

995

996

997

998

999

1000

1001

1002

1003

1004

1005

1006

1007

1008

1009

1010

1011

1012

1013

1014

1015

1016

1017

1018

1019

1020

1021

1022

1023

1024

1025

1026

1027

1028

1029

1030

1031

1032

1033

1034

1035

1036

1037

1038

1039

1040

1041

929 **REFERENCES**

930

931

932

933

934

935

936

937

938

939

940

941

942

943

944

945

946

947

948

949

950

951

952

953

954

955

956

957

958

959

960

961

962

963

964

965

966

967

968

969

970

971

972

973

974

975

976

977

978

979

980

981

982

983

984

985

986

- 2022. Exploring structure-aware transformer over interaction proposals for human-object interaction detection. In CVPR.
- [2] Rohan Anil, Andrew M Dai, Orhan Firat, Melvin Johnson, Dmitry Lepikhin, Alexandre Passos, Siamak Shakeri, Emanuel Taropa, Paige Bailey, Zhifeng Chen, et al. 2023. Palm 2 technical report. arXiv preprint arXiv:2305.10403 (2023).
- [3] Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. 2021. On the opportunities and risks of foundation models. arXiv preprint arXiv:2108.07258 (2021).
- [4] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *NeurIPS* 33 (2020), 1877–1901.
- [5] Yu-Wei Chao, Zhan Wang, Yugeng He, Jiaxuan Wang, and Jia Deng. 2015. Hico: A benchmark for recognizing human-object interactions in images. In *ICCV*. 1017–1025.
- [6] Guikun Chen, Xia Li, Yi Yang, and Wenguan Wang. 2024. Neural clustering based visual representation learning. In CVPR.
- [7] Long Chen, Hanwang Zhang, Jun Xiao, Wei Liu, and Shih-Fu Chang. 2018. Zero-Shot Visual Recognition Using Semantics-Preserving Adversarial Embedding Networks. In CVPR.
- [8] Zhi-Qi Cheng, Qi Dai, Siyao Li, Teruko Mitamura, and Alexander Hauptmann. 2022. Gsrformer: Grounded situation recognition transformer with alternate semantic attention refinement. In ACM MM. 3272–3281.
- [9] Meng-Jiun Chiou, Henghui Ding, Hanshu Yan, Changhu Wang, Roger Zimmermann, and Jiashi Feng. 2021. Recovering the unbiased scene graphs from the biased ones. In ACM International Conference on Multimedia.
- [10] Junhyeong Cho, Youngseok Yoon, and Suha Kwak. 2022. Collaborative transformers for grounded situation recognition. In CVPR. 19659–19668.
- [11] KR1442 Chowdhary and KR Chowdhary. 2020. Natural language processing. (2020), 603–649.
- [12] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805 (2018).
- [13] Qingxiu Dong, Lei Li, Damai Dai, Ce Zheng, Zhiyong Wu, Baobao Chang, Xu Sun, Jingjing Xu, and Zhifang Sui. 2022. A survey on in-context learning. arXiv preprint arXiv:2301.00234 (2022).
- [14] Luciano Floridi and Massimo Chiriatti. 2020. GPT-3: Its nature, scope, limits, and consequences. Minds and Machines (2020), 681–694.
- [15] Zixian Guo, Bowen Dong, Zhilong Ji, Jinfeng Bai, Yiwen Guo, and Wangmeng Zuo. 2023. Texts as images in prompt tuning for multi-label image recognition. In CVPR. 2808–2817.
- [16] Dong-Jin Kim, Xiao Sun, Jinsoo Choi, Stephen Lin, and In So Kweon. 2020. Detecting human-object interactions with action co-occurrence priors. In ECCV. 718–736.
- [17] Fanjie Kong, Yanbei Chen, Jiarui Cai, and Davide Modolo. 2024. Hyperbolic learning with synthetic captions for open-world detection. In CVPR.
- [18] Klemens Lagler, Michael Schindelegger, Johannes Böhm, Hana Krásná, and Tobias Nilsson. 2013. GPT2: Empirical slant delay model for radio space geodetic techniques. (2013), 1069–1073.
- [19] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. 2023. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International conference on machine learning*. PMLR, 19730–19742.
- [20] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. 2022. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *ICML*. PMLR, 12888–12900.
- [21] Junnan Li, Ramprasaath Selvaraju, Akhilesh Gotmare, Shafiq Joty, Caiming Xiong, and Steven Chu Hong Hoi. 2021. Align before fuse: Vision and language representation learning with momentum distillation. Advances in neural information processing systems (2021), 9694–9705.
- [22] Kexin Li, Zongxin Yang, Lei Chen, Yi Yang, and Jun Xiao. 2023. Catr: Combinatorial-dependence audio-queried transformer for audio-visual video segmentation. In ACM MM. 1485–1494.
- [23] Lin Li, Long Chen, Yifeng Huang, Zhimeng Zhang, Songyang Zhang, and Jun Xiao. 2022. The devil is in the labels: Noisy label correction for robust scene graph generation. In CVPR. 18869–18878.
- [24] Lin Li, Long Chen, Hanrong Shi, Hanwang Zhang, Yi Yang, Wei Liu, and Jun Xiao. 2022. Nicest: Noisy label correction and training for robust scene graph generation. arXiv preprint arXiv:2207.13316 (2022).
- [25] Li Li, Wei Ji, Yiming Wu, Mengze Li, You Qin, Lina Wei, and Roger Zimmermann. 2024. Panoptic scene graph generation with semantics-prototype learning. In AAAI, Vol. 38. 3145–3153.
- [26] Lin Li, Jun Xiao, Guikun Chen, Jian Shao, Yueting Zhuang, and Long Chen. 2024. Zero-shot visual relation detection via composite visual cues from large language models. *NeurIPS* 36 (2024).

- [27] Liunian Harold Li, Pengchuan Zhang, Haotian Zhang, Jianwei Yang, Chunyuan Li, Yiwu Zhong, Lijuan Wang, Lu Yuan, Lei Zhang, Jenq-Neng Hwang, et al. 2022. Grounded language-image pre-training. In CVPR. 10965–10975.
- [28] Manling Li, Ruochen Xu, Shuohang Wang, Luowei Zhou, Xudong Lin, Chenguang Zhu, Michael Zeng, Heng Ji, and Shih-Fu Chang. 2022. Clip-event: Connecting text and images with event structures. In CVPR. 16420–16429.
- [29] Wei Lin, Leonid Karlinsky, Nina Shvetsova, Horst Possegger, Mateusz Kozinski, Rameswar Panda, Rogerio Feris, Hilde Kuehne, and Horst Bischof. 2023. Match, expand and improve: Unsupervised finetuning for zero-shot action recognition with language knowledge. In *ICCV*. 2851–2862.
- [30] Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Chunyuan Li, Jianwei Yang, Hang Su, Jun Zhu, et al. 2023. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. arXiv preprint arXiv:2303.05499 (2023).
- [31] Dengsheng Lu and Qihao Weng. 2007. A survey of image classification methods and techniques for improving classification performance. *International journal of Remote sensing* 28, 5 (2007), 823–870.
- [32] Sachit Menon and Carl Vondrick. 2022. Visual Classification via Description from Large Language Models. In *ICLR*.
- [33] Sewon Min, Xinxi Lyu, Ari Holtzman, Mikel Artetxe, Mike Lewis, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2022. Rethinking the Role of Demonstrations: What Makes In-Context Learning Work?. In *EMNLP*. 11048–11064.
- [34] Zachary Novack, Saurabh Garg, Julian McAuley, and Zachary C Lipton. 2023. Chils: Zero-shot image classification with hierarchical label sets. In ICML.
- [35] Farhad Pourpanah, Moloud Abdar, Yuxuan Luo, Xinlei Zhou, Ran Wang, Chee Peng Lim, Xi-Zhao Wang, and QM Jonathan Wu. 2022. A review of generalized zero-shot learning methods. *IEEE TPAMI* 45, 4 (2022), 4051–4070.
- [36] Sarah Pratt, Mark Yatskar, Luca Weihs, Ali Farhadi, and Aniruddha Kembhavi. 2020. Grounded situation recognition. In ECCV. Springer, 314–332.
- [37] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *ICML*. 8748–8763.
- [38] Bernardino Romera-Paredes and Philip Torr. 2015. An embarrassingly simple approach to zero-shot learning. In *ICML*. PMLR, 2152–2161.
- [39] Yibing Song, Ruifei Zhang, Zhihong Chen, Xiang Wan, and Guanbin Li. 2023. Advancing visual grounding with scene knowledge: Benchmark and method. In CVPR. 15039-15049.
- [40] Kaihua Tang, Yulei Niu, Jianqiang Huang, Jiaxin Shi, and Hanwang Zhang. 2020. Unbiased scene graph generation from biased training. In CVPR. 3716–3725.
- [41] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. arXiv preprint arXiv:2307.09288 (2023).
- [42] Meng Wei, Long Chen, Wei Ji, Xiaoyu Yue, and Tat-Seng Chua. 2022. Rethinking the two-stage framework for grounded situation recognition. In AAAI, Vol. 36. 2651–2658.
- [43] Yongqin Xian, Christoph H Lampert, Bernt Schiele, and Zeynep Akata. 2018. Zero-shot learning—a comprehensive evaluation of the good, the bad and the ugly. *IEEE TPAMI* 41, 9 (2018), 2251–2265.
- [44] Yongqin Xian, Tobias Lorenz, Bernt Schiele, and Zeynep Akata. 2018. Feature generating networks for zero-shot learning. In CVPR. 5542–5551.
- [45] Danfei Xu, Yuke Zhu, Christopher B Choy, and Li Fei-Fei. 2017. Scene graph generation by iterative message passing. In CVPR. 5410–5419.
- [46] Jie Xu, Hanbo Zhang, Qingyi Si, Yifeng Li, Xuguang Lan, and Tao Kong. 2024. Towards Unified Interactive Visual Grounding in The Wild. In *ICRA*.
- [47] An Yan, Yu Wang, Yiwu Zhong, Chengyu Dong, Zexue He, Yujie Lu, William Yang Wang, Jingbo Shang, and Julian McAuley. 2023. Learning concise and descriptive attributes for visual recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 3090–3100.
- [48] Shaotian Yan, Chen Shen, Zhongming Jin, Jianqiang Huang, Rongxin Jiang, Yaowu Chen, and Xian-Sheng Hua. 2020. Pcpl: Predicate-correlation perception learning for unbiased scene graph generation. In ACM MM. 265–273.
- [49] Chen Yang and Thomas A Cleland. 2024. Annolid: Annotate, Segment, and Track Anything You Need. arXiv preprint arXiv:2403.18690 (2024).
- [50] Yue Yang, Artemis Panagopoulou, Shenghao Zhou, Daniel Jin, Chris Callison-Burch, and Mark Yatskar. 2023. Language in a bottle: Language model guided concept bottlenecks for interpretable image classification. In CVPR 19187-19197.
- [51] Mark Yatskar, Luke Zettlemoyer, and Ali Farhadi. 2016. Situation recognition: Visual semantic role labeling for image understanding. 5534–5542.
- [52] Qifan Yu, Juncheng Li, Yu Wu, Siliang Tang, Wei Ji, and Yueting Zhuang. 2023. Visually-prompted language model for fine-grained scene graph generation in an open world. In *ICCV*. 21560–21571.
- [53] Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, et al. 2023. Opt: Open pre-trained transformer language models, 2022. URL https://arxiv. org/abs/2205.01068 (2023), 19–0.

Anonymous Authors

1045	[54] Zhen Zhang, Anran Lin, Chun Wai Wong, Xiangyu Chu, Qi Dou, and KW Au.	(2023).	1103
1046	2023. Interactive Navigation in Environments with Traversable Obstacles Using Large Language and Vision-Language Models. <i>arXiv preprint arXiv:2310.08873</i>	[55] Zhong-Qiu Zhao, Peng Zheng, Shou-tao Xu, and Xindong Wu. 2019. Object detection with deep learning: A review. 30, 11 (2019), 3212–3232.	1104
1047			1105
1048			1106
1049			1107
1050			1108
1051			1109
1052			1110
1053			1111
1054			1112
1055			1115
1050			1114
1057			1115
1050			1110
1057			1117
1061			1110
1062			1120
1062			1120
1064			1122
1065			1122
1066			1124
1067			1125
1068			1126
1069			1127
1070			1128
1071			1129
1072			1130
1073			1131
1074			1132
1075			1133
1076			1134
1077			1135
1078			1136
1079			1137
1080			1138
1081			1139
1082			1140
1083			1141
1084			1142
1085			1143
1086			1144
1087			1145
1088			1146
1089			1147
1090			1148
1091			1149
1092			1150
1093			1151
1094			1152
1095			1155
1097			1154
1098			1155
1099			1150
1100			1158
1101			1159
1102			1160