

AffordanceLLM: Grounding Affordance from Vision Language Models

Shengyi Qian* Weifeng Chen Min Bai Xiong Zhou Zhuowen Tu Li Erran Li
AWS AI, Amazon

<https://jasonqsy.github.io/AffordanceLLM>




Figure 1. **Illustration for the affordance grounding task.** The input is a single image and the corresponding action (e.g., “hold”). The output is a heatmap which highlights regions one can interact. We aim to enhance the generalization capability of affordance grounding to in-the-wild objects that are unseen during training, by developing a new approach, AffordanceLLM, that takes the advantage of the rich knowledge from large-scale vision language models [40] beyond the supervision from the training images.

Abstract

Affordance grounding refers to the task of finding the area of an object with which one can interact. It is a fundamental but challenging task, as a successful solution requires the comprehensive understanding of a scene in multiple aspects including detection, localization, and recognition of objects with their parts, of geo-spatial configuration/layout of the scene, of 3D shapes and physics, as well as of the functionality and potential interaction of the objects and humans. Much of the knowledge is hidden and beyond the image content with the supervised labels from a limited training set. In this paper, we make an attempt to improve the generalization capability of the current affordance grounding by taking the advantage of the rich world, abstract, and human-object-interaction knowledge from pre-trained large-scale vision language models [40]. Under the AGD20K benchmark, our proposed model demonstrates a significant performance gain over the competing methods for in-the-wild object affordance grounding. We further demonstrate it can ground affordance for objects from random Internet images, even if both objects and actions are unseen during training.

*The work was done during an Amazon internship.

1. Introduction

Grounding affordance from a single image is a fundamental problem in computer vision. It forms the stepping stone to downstream tasks such as understanding human-object interaction [8, 17, 60], visual navigation [31], and object manipulation [1, 23]. Past approaches generally use human demonstrations as supervision to advance this field with tremendous success [34, 44, 45, 50]. While such approaches perform well on objects and actions seen during training, they struggle when generalizing in the wild, i.e. on novel objects unseen during training (Fig. 1).

The difficulties in generalization arise from the fact that affordance grounding is a challenging task that requires comprehensive understanding of an image from multiple aspects. A successful solution requires an understanding of the 3D geometry and functionality of objects and parts, of the actions and intentions of the executing agent, of the potential interaction between object and human, as well as of the spatial configuration of the environment. Much of these knowledge lies beyond the ground-truth localization/recognition of objects and parts provided as heatmaps in a limited training set.

In this paper we make attempts to improve affordance grounding in the wild by leveraging the rich world, abstract, and human-object-interaction knowledge embedded

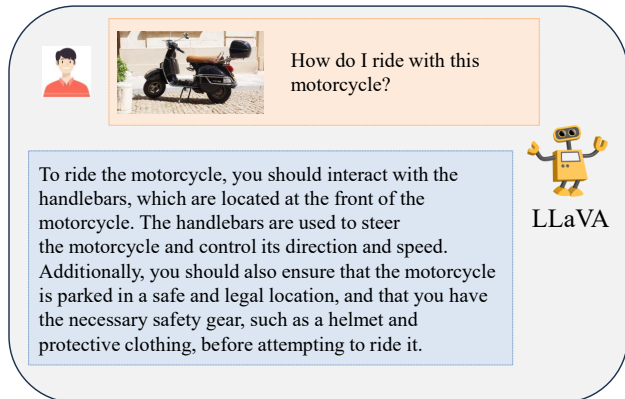


Figure 2. State-of-the-art vision language models, such as LLaVA [40], has rich human-object-interaction knowledge, thanks to the large-scale text pretraining. Given a question about how to interact with an object, it typically gives a reasonable solution.

in large-scale Vision Language Models (VLMs). With large-scale text pretraining, modern VLMs such as GPT-4 [52], LLaVA [40] and Blip-2 [35] have a rich reservoir of world knowledge, as demonstrated by their extraordinary capabilities in answering visually grounded common sense questions [4]. World knowledge is instrumental to affordance reasoning — when presented with an image of a motorcycle and questioned about “How do I ride with this motorcycle?”, LLaVA answers “To ride the motorcycle, you should interact with the handlebars...” (Fig. 2), which exhibits commonsensical understanding of affordance. Affordance models equipped with similar world knowledge have a better chance generalizing to the wild than a model that purely learns from limited affordance demonstration data.

Beside world knowledge, another novel factor we introduce to improve affordance reasoning is 3D geometry, as it holds rich information of object functionality. A cylindrical area, for example a handle or a stick, is closely related to the action of grabbing or holding, regardless of its color or texture. Similarly, a flat part, for example the surface of a chair or a bench, might indicate areas suitable for sitting or lying. Relating 3D geometries to actions allows us to bypass the difficulties in handling variations in visual appearances, and thus facilitates generalization.

We propose a novel approach, *AffordanceLLM*, that reflects the above intuitions. Our approach builds upon a VLM backbone (LLaVA [40]) to tap into its world knowledge. We achieve it by extending the backbone with a mask decoder and a special token `<mask_token>`, which are used to predict an affordance map. The whole model can be trained end-to-end. Additionally, we introduce depth maps as 3D information in parallel to RGB images as input to our network, with the goal of eliciting geometric reasoning capability from the network. We found both designs significantly improve performance.

We evaluate our method on the AGD20K [45] bench-

mark, as this is the only large-scale affordance grounding dataset with accurate action and object labels. We re-split the benchmark to test capability of models to generalize to object categories unseen during training. We show that our approach outperforms all state-of-the-art baselines by a large margin. We take a further step to validate the generalization ability by testing our approach on random Internet images. It produces reasonable affordance maps on object categories very different from the ones in training set. Moreover, it even possesses some capability of generalizing to completely novel actions.

In summary, our contributions are as follows:

1. We introduce the first-ever affordance grounding approach that leverages the rich world knowledge embedded in pretrained VLMs, enabling the model to generalize beyond training data;
2. We demonstrate the importance of 3D information in affordance grounding;
3. Our proposed approach generalizes to novel objects and outperforms all state-of-the-art approaches on AGD20K. It even shows evidence that it could generalize to novel actions.

2. Related Work

Eliciting World Knowledge from Vision Language Models. Foundational Vision Language Models that bridge images and language have a rich reservoir of world knowledge, and recent researches have been tapping into it to make advancement in vision tasks. The joint visual-language embedding space learnt from simple image-text pair [13, 57] has made it possible to improve open-world detection [42, 47, 48, 61], and segmentation [33, 43, 64, 67]. The world knowledge here is the correspondence between visual and language concepts.

Large language models (LLMs) trained on trillions of tokens contain even richer world knowledge and are capable of answering common-sense questions. Coupled with vision inputs, the resulting multi-modal LLMs are brought in to solve complex vision problems. For example, Kosmos-2 and Groundhog incorporate the reasoning skill of LLMs to generate bounding boxes and segmentation masks [32, 53, 75]. 3D LLM further extend LLMs to reason about 3D scenes, including visual grounding and navigation [12, 22, 24, 25, 37, 68, 77]. For robotics, PaLM-E and RT2 transfer the knowledge from visual-language domains into motion planning and manipulation [3, 14]. Our approach embodies the same idea to transfer the world knowledge from VLMs, but applies it on a novel setting – solving visual affordance grounding.

Affordance Grounding. Understanding object affordance from a single image is an important step towards embodied visual intelligence, and researchers have built many dif-

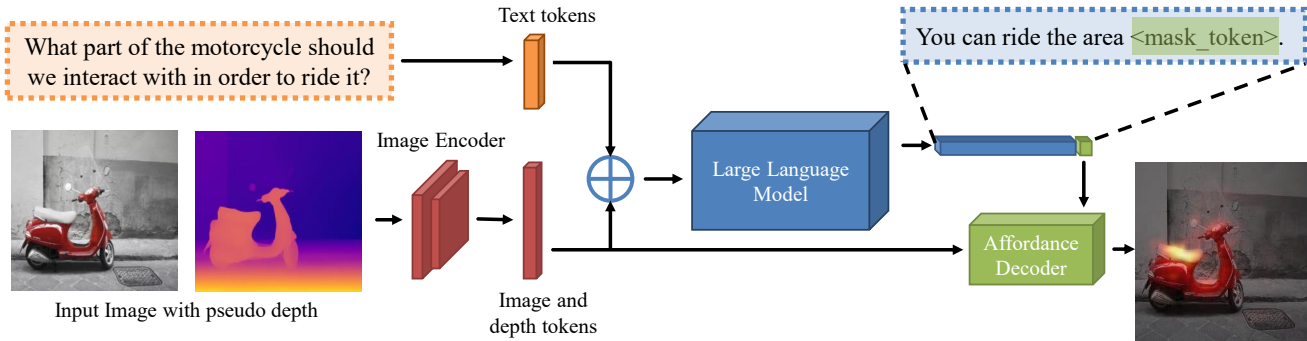


Figure 3. **Overview of AffordanceLLM.** The inputs of our model includes a single image and a text prompt related to interaction. We use OWL-ViT [48] as the image encoder to generate image features and project it into the same hidden dimension as the large language model. As well, we use a tokenizer to encode the text prompt. The text features and image features are concatenated together and feed into the LLM. The LLM is fine-tuned to predict a special token, which is used as a query to the mask decoder to generate the final affordance map.

ferent approaches to endow machines to have this ability. Nagarajan *et al.* first proposes to ground object affordance from Internet videos [50]. Fang *et al.* constructs an object affordance dataset based on product review videos [16]. Luo *et al.* annotates the first large-scale affordance dataset and call it AGD20K [45]. LOCATE [34] is the state-of-the-art approach on AGD20K. More recently, researchers further extends the scope of the affordance grounding problem, including extending it to scene understanding [9, 30, 54], 3D models [69], egocentric videos [49], hand pose generation [27, 70], or associating it with human parts [46]. We use AGD20K as our primary benchmark, and compare our approach with state-of-the-art methods [21, 34, 44, 50, 54].

Incorporating 3D Information for Vision Tasks. 3D information has been shown to be critical in certain vision and robotics tasks. For example, Zhou *et al.* [76] found that visual navigation in mobile sensorimotor systems can benefit from 3D input. Kerr *et al.* [28] found the NeRF-rendered depth map can help grasping in robotics. Similarly, grounding affordance could benefit from 3D information as well, as 3D shapes of objects and their parts hold a lot of hints on their utility and the proper ways to interact with them. While 3D information is not usually available for an arbitrary image, fortunately, researchers have built a series of robust approaches to estimate the 3D of an image, ranging from surface normal estimation [15, 65], depth estimation [10, 36, 59, 72, 73], to 3D reconstruction [18, 39, 51] and few-image NeRF [5, 41, 56, 74]. In our paper, we mainly use DPT [59] to generate pseudo depth maps to help VLMs to build 3D understanding.

Robotics Manipulation. Manipulation of in-the-wild objects is an important but challenging task in robotics due to the difficulty of data collection. Researchers have developed many methods for different objects in different scenes, such as tabletop objects [19, 26, 62] and mobile manipulation [71]. While manipulation is not our goal, learning affordance can be a solution for manipulation [2, 23, 66].

3. Approach

We now introduce our approach, AffordanceLLM, which takes a single image I and an affordance text query T , and generates an affordance map M . We use a template of “What part of the $\langle \text{object_name} \rangle$ should we interact with in order to $\langle \text{action_name} \rangle$ it?” as the text query T . We then train LLM to generate a special token $\langle \text{mask_token} \rangle$ and use its hidden state to decode a dense affordance map M . A brief overview is shown in Fig 3.

3.1. Overview

Large language model. We choose LLaVA-7B as our backbone multimodal large language model. We refer the reader for a fuller explanation in [40], but briefly, LLaVA contains an image encoder, a text tokenizer, and a large language model LLM. The image encoder is typically a CLIP pretrained ViT, with a linear layer to project the hidden dimension. It encodes the image I into image features F_I . At the same time, the tokenizer encodes the text T into text features F_T . They are concatenated and feed into the language model. The LLM produces text output A as:

$$A = \text{LLM}(F_I, F_T). \quad (1)$$

Predicting affordance. How do we perform affordance reasoning while leveraging the world knowledge embedded in LLM? We propose to treat affordance as an implicit text token predicted from the LLM, which could be further decoded into a 2D map. Specifically, we train the LLM to predict a special token $\langle \text{mask_token} \rangle$, the hidden state of which is first projected into a query embedding q and then fed into a Decoder to generate a dense affordance map. Decoder shares a similar architecture as the ones in SAM [29] and 3DOI [54]. It takes in q and image features F_I to produce an affordance map M , i.e.,

$$M = \text{Decoder}(F_I, q). \quad (2)$$

Pseudodepth as additional inputs. Besides images, the affordance reasoning task could benefit from 3D information (we will validate the benefits in Sec 4). However, modern VLMs are typically only trained with text and 2D images [35, 40]. Therefore, we also include a pseudo depth map as additional inputs to the large language model. For each image, we use the state-of-the-art depth estimation model DPT [59] to generate a pseudo depth map D . We use the same image encoder to encode the depth map D to produce depth features F_D , and concatenate it with image features. Our final model is thus

$$A, M = \text{AffordanceLLM}(F_I, F_D, F_T). \quad (3)$$

Training objectives. Following the same practice as [54], we train the affordance map using a binary focal loss [38], and call it affordance loss L_{aff} . We set the weight of positive examples to be 0.95 and that of negative ones to be 0.05 to balance positives and negatives, as there are more negatives than positives in ground truth affordance map. We follow the standard cross entropy loss for the text output of language models. Our final loss function is a linear combination of affordance loss and text loss, given by

$$L = L_{\text{aff}} + \lambda \cdot L_{\text{text}}. \quad (4)$$

In practice, we set $\lambda = 0.01$ to balance two losses, as the affordance loss can be quite small due to the imbalance of positive and negative values.

3.2. Network Architecture

Next, we discuss the network architecture, and the training details of our model.

Image encoder. The standard LLaVA uses CLIP image encoder and a linear projection layer [40, 67]. In practice, we find that the CLIP image encoder has low resolution (224x224) and does not capture sufficient information about grounding. Therefore, we use OWL-ViT [48] to replace the standard CLIP-ViT [57]. OWL-ViT has an input resolution of 768x768, which is significantly higher than CLIP. At the same time, OWL-ViT is pretrained to extract features that include precise location information of objects. As we will empirically show in experiments, using OWL-ViT is significantly better than CLIP. However, we note that our approach is general, and will benefit from any future improvements in pretrained VLM backbones.

Projection. Another problem of using OWL-ViT is about its embedding space. With a much higher input resolution, OWL-ViT produces 576 tokens with a hidden dimension of 768 for each image. In comparison, CLIP only produces 256 tokens for each image. Projecting each individual token into the hidden dimension of LLM (4096) consumes a lot of GPU memory. In practice, we project each token

of OWL-ViT to 1024, and concatenating four neighboring tokens into a single token.

Language model. We follow LLaVA [40] and LLama [63] to use the standard text tokenizer to encode our text query. We use LLama-7B [63] as the large language model.

Affordance decoder. We aim to keep a lightweight decoder, as it has been proved to produce good segmentation masks and affordance maps [6, 11, 29, 54]. However, we find the vanilla mask decoder is too lightweight in our case and does not produce high-resolution affordance map. Therefore, we add an additional transposed convolution layer to increase its output resolution.

Implementation. We implement our model using PyTorch and HuggingFace. We initialize our model with LLama-7B pretrained weights. Following LLaVA [40], we freeze the image encoder, pretrain the image projection layer to align OWL-ViT and LLama features, and then use GPT instructions to tune the language model. Finally, we add the mask encoder [29, 54] and tune the whole model on AGD20K [45], which has annotations of object affordance. We use eight NVIDIA A100 (40GB) to train our model, with Fully Sharded Data Parallel. We use a batch size of 4 and set the learning rate as $2e-5$.

4. Experiments

In experiments, we aim to systematically evaluate the performance of our approach. In particular, we are interested in answering these questions: (1) How well does it generalize, compared with state-of-the-art methods? (2) How does each design choice contribute to the final performance, including prompts, visual encoders, and depth?

4.1. Experimental Setup

Metrics. We evaluate primarily on AGD20K [45] and follow its metrics to evaluate our model, which is KLD, SIM and NSS [34, 45, 50, 54]. For KLD, the lower the better. And for SIM and NSS, the higher, the better. A full explanation is available in the supplemental.

Baselines. We compare our approach against state-of-the-art baselines. In general, affordance grounding methods belong to two categories: weakly supervised and fully supervised methods. We report performance of both categories. (*Weakly supervised methods*): They do not train on explicit labels of the affordance map. Instead, they are trained on a human demonstration of the same object. These approaches include InteractionHotspots [50], Cross-View-AG [45], Cross-View-AG+ [44], AffCorrs [21], and LOCATE [34]. Among them, LOCATE is the most recent model and has the best results on AGD20K. We use the reported number in LOCATE for the easy split and retrain them for the hard split. Among them, we cannot run AffCorrs, as it focuses on one-shot affordance learning. The




Figure 4. Qualitative results on the test set of the hard split. LOCATE-Sup fails to learn a reasonable affordance map due to limited training data. LOCATE [34] typically predicts an affordance map which covers the whole object. 3DOI [54] focuses on a small area of the object. Overall, our approach produces the best-quality affordance predictions.

reported model on the easy split is adapted by [34] and not publicly available. We also do not run InteractionHotspots because the pretrained model only supports 7 actions. The reported model is retrained by [45] but lacks sufficient implementation details to be reproduced. Therefore, we retrain Cross-View-AG [45], Cross-View-AG+ [44], and LOCATE [34] on the hard split. We maintain the object/action split, but allow them to use more images for weak supervision. Therefore, they have 11,889 images for training.

(Fully supervised methods): Affordance map can also be learned from explicit labels, and we call it supervised methods. This includes 3DOI [54] and ours. We also adapt LOCATE to a fully supervised version for fair comparison.

- *3DOI* [54]: 3DOI is a SAM-based model [29], which takes a single image and a query point and predicts the segmentation mask and affordance map. Therefore, We randomly sample a pixel with score > 0.9 as the query point from the affordance map. We use the 3DOI pretrained model, which has never seen any images in AGD20K, in-

cluding the training set.

- *LOCATE-Sup* [34]: To ensure fair comparison, we also adopt LOCATE and train it using the same binary focal loss as our model. We call it LOCATE-Sup. LOCATE uses a Dino-ViT as its visual encoder [7]. To eliminate the effect of different pretrained visual encoders, we also report the performance of LOCATE-Sup-OWL, which uses the exact same pretrained visual encoder as ours.

4.2. Dataset

We follow LOCATE [34] to evaluate primarily on AGD20K [45], as it is the only large-scale dataset for affordance with action and object labels. Since our approach is not weakly supervised and requires dense annotations, we only use AGD20K images with dense annotations.

In this paper, we primarily evaluate the ability of an affordance model to generalize to unseen object categories, and thus evaluate on the *Unseen* split of the AGD20K benchmark. This split ensures that there is no overlap be-

Table 1. Difficulty score of different splits. The lower the score, the more similar are the object categories in the train and test set.

Splits	Same	Easy	Hard	Random
Difficulty Score \uparrow	0.000	0.356	0.412	0.491

Table 2. **Quantitative results on the *Easy* split of AGD20K [45].** InteractionHotspots, Cross-View-AG(+), AffCorrs and LOCATE are trained on AGD20K images with weak supervision (13,323 images). LOCATE-Sup and LOCATE-Sup-OWL, and AffordanceLLM are trained on AGD20K images with dense annotation (1,135 images). 3DOI is trained on their own dataset with dense annotation (10,000 images) [54]. AffordanceLLM is comparable to LOCATE [34] on the easy split, where test objects have similar counterparts in the training set. The **best** and second-best results are highlighted in bold and underlined, respectively.

Methods	KLD \downarrow	SIM \uparrow	NSS \uparrow
InteractionHotspots [50]	1.994	0.237	0.577
Cross-View-AG [45]	1.787	0.285	0.829
Cross-View-AG+ [44]	1.765	0.279	0.882
AffCorrs [21]	1.618	0.348	1.021
LOCATE [34]	1.405	<u>0.372</u>	1.157
LOCATE-Sup [34]	1.907	0.236	0.641
LOCATE-Sup-OWL [34, 48]	1.927	0.234	0.624
3DOI [54]	3.565	0.227	0.657
AffordanceLLM (Ours)	<u>1.463</u>	0.377	<u>1.070</u>

tween the object categories in the train and test set.

However, the *Unseen* split has one major issue: there are still a lot of similarities between the objects in the train and test set. Objects in the test set typically have similar counterparts in the training set, leaving models room for memorizing. For example, “skis” in the test set maps to “snowboards” and “skateboards” in training set, “basketball” maps to “baseball”, “knife” maps to “fork”, etc. To make the benchmark more reflective of the generalization ability of a model, we provide a more challenging split. We thus have the following two splits.

Easy split. This is the original *Unseen* split of AGD20K. We have 1135/540 images for train and test with dense annotations for the fully supervised setting, or 13,323/540 images for the weakly supervised setting. The test set remains the same for both settings.

Hard split. We *randomly* put 50% AGD20K object classes into the training set and the remaining classes into the test set to simulate in-the-wild generalization (details in the supplemental). The training and test object are not overlapping, and most objects in the test set do not have a similar counterpart in the training set, and is thus harder to generalize than the *Easy* split. We have 868/807 images for train and test with dense annotations for the fully supervised setting, and 11,889/807 images for weakly supervised setting. The test set is the same for both settings.

Table 3. **Quantitative results on the *Hard* split of AGD20K [45].**

Cross-View-AG(+) and LOCATE are trained on AGD20K images with weak supervision (11,889 images). LOCATE-Sup and LOCATE-Sup-OWL, and AffordanceLLM are trained on AGD20K images with dense annotation (868 images). 3DOI is trained on their own dataset with dense annotation (10,000 images) [54]. On the hard split, AffordanceLLM outperforms all baselines by a large margin, which demonstrates the superior generalization ability of our model. We do not run InteractionHotspots [50] and AffCorrs [21], as the reported model has ambiguous implementation details, or is not publicly available. The **best** and second-best results are highlighted in bold and underlined, respectively.

Methods	KLD \downarrow	SIM \uparrow	NSS \uparrow
Cross-View-AG [45]	2.092	0.209	0.138
Cross-View-AG+ [44]	2.034	0.218	0.342
LOCATE [34]	<u>1.829</u>	<u>0.282</u>	0.276
LOCATE-Sup [34]	2.003	0.224	0.435
LOCATE-Sup-OWL [34, 48]	2.127	0.206	0.314
3DOI [54]	4.017	0.200	<u>0.549</u>
AffordanceLLM (Ours)	1.661	0.361	0.947

Measuring split difficulty. We propose a metric to quantify the generalization difficulty of a split. Intuitively, the difficulty to generalize to an object class in the test set is defined by how different it is from the classes in the training set, which could be measured by its semantic distance to the most similar class in the training set [55]. The greater the distance, the harder it is to generalize to this test class. Therefore, for each semantic class c in the test set, we compute its distance d to the most similar class in the training set. We use the CLIP [58] text encoder to obtain an embedding to represent each object class. Assume train classes are C_{train} and test classes are C_{test} , the difficulty of this split is

$$D(C_{\text{train}}, C_{\text{test}}) = 1 - \frac{1}{|C_{\text{test}}|} \sum_{c \in C_{\text{test}}} \max_{c' \in C_{\text{train}}} d(c, c'). \quad (5)$$

We compare the difficulty metric among four settings: (1) *Same*: train and test share the same classes; (2) *Easy* split; (3) *Hard* split; (4) *Random*: constructed by randomly even-splitting 50 object classes from LVIS [20], which serves as a lower bound.

We show the difficulty metrics in Tab 1. The *Same* split has a similarity metric of 0.0, as all object classes in the test are present during training. The *Easy* split has a similarity metric of 0.356. The *Random* split has a score of 0.491. The *Hard* split has a higher score than *Easy*, meaning that the difference between test and train is more significance in *Hard* than in *Easy*, and is thus harder to generalize.

4.3. Results

Figure 4 shows qualitative results on the test set of the hard split. In this split, the objects in the test set bear little to none

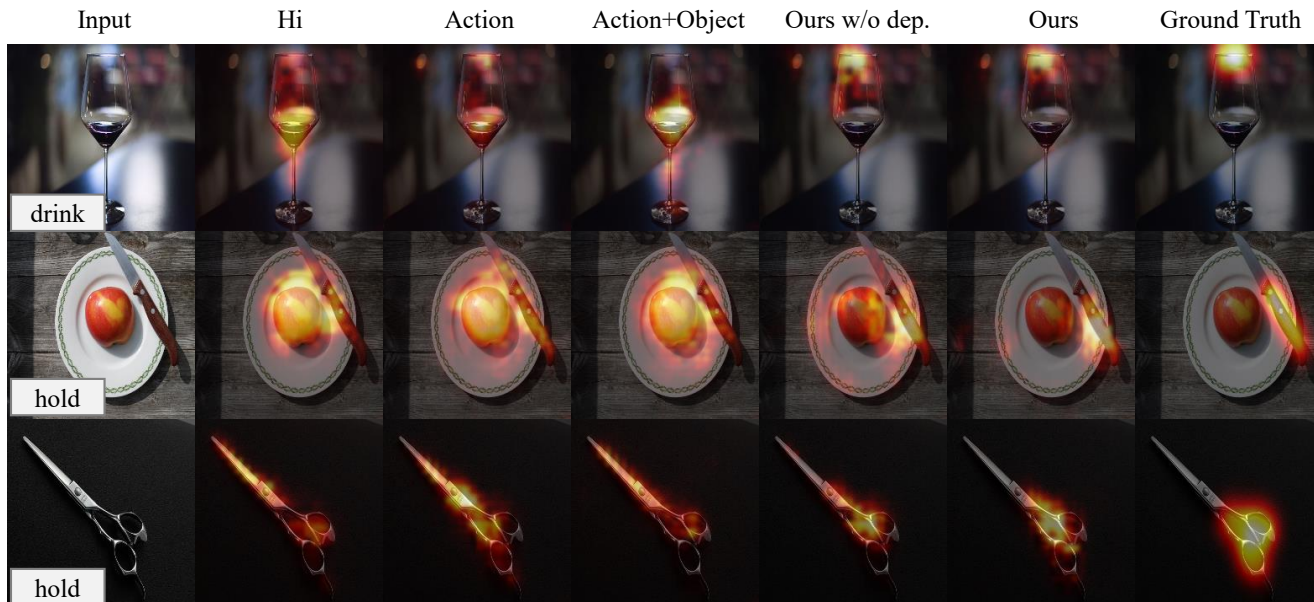


Figure 5. Ablation of different text prompts and depth. Ours w/o depth is our approach without pseudodepth as additional inputs. Ours is our full approach. We find constructing the correct text prompt typically helps our model to focus on the correct area. We believe it is because the correct text prompt would activate the world knowledge related to affordance embedded in the VLM.

resemblance to the ones in the training set. We compare our approach, AffordanceLLM, with a set of state-of-the-art baselines. LOCATE [34] tends to predict an affordance map that covers the entire object, indicating poor generalization performance. 3DOI [54] typically focuses on a small area of the object, and sometimes fails to ground the correct region. LOCATE-Sup fails to predict reasonable affordance map, probably due to the small amount of training data. Despite being trained on the same training set as LOCATE-Sup, our approach is able to produce the best affordance map among all methods, showcasing superior generalization capability.

We further compare our model with baselines quantitatively and the results are summarized in Tab. 2, 3. On the *Hard* split, where the test set objects differ semantically from the training set, our method outperforms all baselines significantly. This improvement can be attributed to the extensive world knowledge and understanding embedded within the large language model. On the *Easy* split, our model is comparable to LOCATE [34] and outperform all other baselines. We hypothesize that the advantage of our approach is less pronounced when the test and train objects exhibit similarity, as the generalization capability becomes less critical. It is also worth noting that unlike LOCATE which is weakly supervised on 10k+ images, our model was fully supervised on some 1k images with dense annotations, which renders a more meaningful comparison with LOCATE-Sup that is trained on the same data. Our method significantly outperforms LOCATE-Sup on both splits, indicating the effectiveness of our approach.

Table 4. Ablation on the hard split. We validate the importance of text prompts, image encoders and pseudo depth to performance.

Depth	Text Prompt	Img Encoder	KLD ↓	SIM ↑	NSS ↑
Yes	Full	OWL-ViT	1.661	0.361	0.947
-	Full	OWL-ViT	1.713	0.352	0.881
-	Full	CLIP-ViT	1.759	0.286	0.776
-	Object, Action	OWL-ViT	1.769	0.329	0.827
-	Action	OWL-ViT	1.843	0.336	0.815
-	Hi	OWL-ViT	1.836	0.325	0.793

4.4. Ablation

We conduct a few ablation studies to understand how different components of the model contribute to the final performance. We test different text prompts, different image encoders, and the effect of pseudo depth as inputs. The results are summarized in Tab 4 and Fig 5.

Text prompts. Prompt tuning is known to have major effects on large VLMs. We test four different text prompts to understand the effect of text content on model performance:

- *Hi*: We use “Hi” as our text prompt.
- *Action*: We use the action (e.g. “hold”) as the prompt.
- *Object + Action*: We use the object name and action label as our text prompt, for example “hold, knife”.
- *Full*: We use a complete question as the text prompt — “What part of the motorcycle should we interact with in order to push it?”.

We notice that the *Full* prompt yields a higher performance compared with other simple text prompts. It demonstrates that specific question prompt is helpful for extracting the knowledge from pretrained large language models.

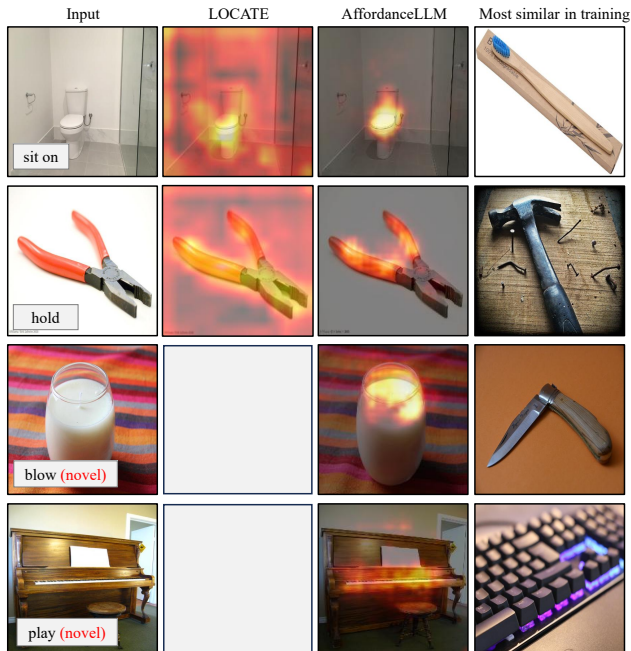


Figure 6. Generalization results on random Internet images. We show the most similar objects in the training set to demonstrate how different the objects are from the ones in the training set. **(Row 1, 2):** AffordanceLLM generalizes to novel objects from random Internet images, while LOCATE [34] fails. **(Row 3, 4):** AffordanceLLM generalizes to novel actions plus novel objects. LOCATE cannot infer novel actions thus we left it blank.

Vision encoders. Although LLaVA [40] uses CLIP-ViT, it may not be the optimal vision encoder for our affordance grounding task — CLIP-ViT is trained with an objective to align text-image pairs and is not explicitly optimized to perform localization, and therefore has limited visual grounding capability. We therefore switch to OWL-ViT [48], which is trained on detection datasets with 2M images, and achieves state-of-the-art open vocabulary detection performance. As shown in Tab 4, using OWL-ViT as vision backbone far excels using CLIP-ViT. It indicates the importance of grounding capability of visual backbone.

4.5. Pseudodepth as Inputs

Our model is trained with pseudo depth map produced by DPT [59]. To verify whether the additional depth inputs are effective, we compare the model trained with and without estimated depth (Tab 4 and Fig 5). With depth, our model can predict better affordance map, demonstrating the importance of 3D information in affordance reasoning.

4.6. Generalization to Internet Images

We further test the generalization of our model on random Internet images in Fig. 6. All test objects are novel. To showcase how different these objects are from the ones in train set, for each test object, we retrieve the most similar




Figure 7. Failure examples. **(Row 1:)** AffordanceLLM sometimes fails due to multiple objects present in the scene and it fails to refer to the correct object. **(Row 2:)** AffordanceLLM thinks humans should hold the handle to cut something using the knife, while AGD20K annotators think “cut with” should refer to the blade.

object in the train set using the metric defined in Eq 5. As we can see, these objects are not similar to any objects in the train set. We go even further to test if our approach is able to generalize to novel actions in addition to novel objects, such as “blow” and “play”. Generalization to novel actions is even more challenging, requiring open vocabulary understanding of actions, which is beyond the capability of LOCATE [34]. Despite these challenges, our approach not only produces very reasonable affordance maps for novel objects, but is also able to handle novel actions plus novel objects, once again demonstrating the extraordinary capability to generalize.

4.7. Failure Examples

Finally, we show our failure examples in Fig 7. First, we find AffordanceLLM fails on some ambiguous questions. For example, in AGD20K, “cut with” refers to the blade of a knife. However, AffordanceLLM thinks humans should hold the handle of the knife to cut anything. Second, when there are multiple objects in the image, it sometimes cannot refer to the correct object.

5. Conclusion

We have presented AffordanceLLM, a novel approach which can ground affordance for in-the-wild objects unseen during training. By tapping into the world knowledge embedded in a Vision Language Model, our proposed approach generalizes much better to in-the-wild objects, compared with state-of-the-art affordance grounding models.

Our approach can have positive impacts by helping build intelligence robots which can manipulate in-the-wild objects. On the other hand, it can be misused to cause physical damage or harm if applied in an adversarial manner.

References

- [1] Shikhar Bahl, Abhinav Gupta, and Deepak Pathak. Human-to-robot imitation in the wild. 2022. [1](#)
- [2] Shikhar Bahl, Russell Mendonca, Lili Chen, Unnat Jain, and Deepak Pathak. Affordances from human videos as a versatile representation for robotics. In *CVPR*, 2023. [3](#)
- [3] Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Xi Chen, Krzysztof Choromanski, Tianli Ding, Danny Driess, Avinava Dubey, Chelsea Finn, et al. Rt-2: Vision-language-action models transfer web knowledge to robotic control. *arXiv preprint arXiv:2307.15818*, 2023. [2](#)
- [4] Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrlke, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, et al. Sparks of artificial general intelligence: Early experiments with gpt-4. *arXiv preprint arXiv:2303.12712*, 2023. [2](#)
- [5] Ang Cao, Chris Rockwell, and Justin Johnson. Fwd: Real-time novel view synthesis with forward warping and depth. In *CVPR*, 2022. [3](#)
- [6] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *ECCV*, 2020. [4](#)
- [7] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *ICCV*, 2021. [5](#)
- [8] Yu-Wei Chao, Zhan Wang, Yugeng He, Jiakuan Wang, and Jia Deng. Hico: A benchmark for recognizing human-object interactions in images. In *ICCV*, 2015. [1](#)
- [9] Joya Chen, Difei Gao, Kevin Qinghong Lin, and Mike Zheng Shou. Affordance grounding from demonstration video to target image. In *CVPR*, 2023. [3](#)
- [10] Weifeng Chen, Shengyi Qian, and Jia Deng. Learning single-image depth from videos using quality assessment networks. In *CVPR*, 2019. [3](#)
- [11] Bowen Cheng, Alex Schwing, and Alexander Kirillov. Pixel classification is not all you need for semantic segmentation. *Advances in Neural Information Processing Systems*, 34:17864–17875, 2021. [4](#)
- [12] Yinpei Dai, Run Peng, Sikai Li, and Joyce Chai. Think, act, and ask: Open-world interactive personalized robot navigation. *arXiv preprint arXiv:2310.07968*, 2023. [2](#)
- [13] Karan Desai and Justin Johnson. Virtex: Learning visual representations from textual annotations. In *CVPR*, 2021. [2](#)
- [14] Danny Driess, Fei Xia, Mehdi SM Sajjadi, Corey Lynch, Aakanksha Chowdhery, Brian Ichter, Azyaan Wahid, Jonathan Tompson, Quan Vuong, Tianhe Yu, et al. Palme: An embodied multimodal language model. *arXiv preprint arXiv:2303.03378*, 2023. [2](#)
- [15] David Eigen and Rob Fergus. Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture. In *ICCV*, 2015. [3](#)
- [16] Kuan Fang, Te-Lin Wu, Daniel Yang, Silvio Savarese, and Joseph J Lim. Demo2vec: Reasoning object affordances from online videos. In *CVPR*, 2018. [3](#)
- [17] Georgia Gkioxari, Ross Girshick, Piotr Dollar, and Kaiming He. Detecting and recognizing human-object interactions. In *CVPR*, 2018. [1](#)
- [18] Georgia Gkioxari, Jitendra Malik, and Justin Johnson. Mesh r-cnn. In *ICCV*, 2019. [3](#)
- [19] Ankit Goyal, Jie Xu, Yijie Guo, Valts Blukis, Yu-Wei Chao, and Dieter Fox. Rvt: Robotic view transformer for 3d object manipulation. *arXiv preprint arXiv:2306.14896*, 2023. [3](#)
- [20] Agrim Gupta, Piotr Dollar, and Ross Girshick. Lvis: A dataset for large vocabulary instance segmentation. In *CVPR*, 2019. [6](#)
- [21] Denis Hadjivelichkov, Sichelukwanda Zwane, Lourdes Agapito, Marc Peter Deisenroth, and Dimitrios Kanoulas. One-shot transfer of affordance regions? affcorr! In *Conference on Robot Learning*, pages 550–560. PMLR, 2023. [3](#), [4](#), [6](#)
- [22] Yining Hong, Haoyu Zhen, Peihao Chen, Shuhong Zheng, Yilun Du, Zhenfang Chen, and Chuang Gan. 3d-llm: Injecting the 3d world into large language models. *arXiv preprint arXiv:2307.12981*, 2023. [2](#)
- [23] Cheng-Chun Hsu, Zhenyu Jiang, and Yuke Zhu. Ditto in the house: Building articulation models of indoor scenes through interactive perception. In *ICRA*, 2023. [1](#), [3](#)
- [24] Haifeng Huang, Zehan Wang, Rongjie Huang, Luping Liu, Xize Cheng, Yang Zhao, Tao Jin, and Zhou Zhao. Chat-3d v2: Bridging 3d scene and large language models with object identifiers. *arXiv preprint arXiv:2312.08168*, 2023. [2](#)
- [25] Jianguo Huang, Silong Yong, Xiaojian Ma, Xiongkun Linghu, Puhao Li, Yan Wang, Qing Li, Song-Chun Zhu, Baoxiong Jia, and Siyuan Huang. An embodied generalist agent in 3d world. *arXiv preprint arXiv:2311.12871*, 2023. [2](#)
- [26] Stephen James, Zicong Ma, David Rovick Arrojo, and Andrew J Davison. Rlbench: The robot learning benchmark & learning environment. *IEEE Robotics and Automation Letters*, 5(2):3019–3026, 2020. [3](#)
- [27] Juntao Jian, Xiuping Liu, Manyi Li, Ruizhen Hu, and Jian Liu. Affordpose: A large-scale dataset of hand-object interactions with affordance-driven hand pose. In *ICCV*, 2023. [3](#)
- [28] Justin Kerr, Letian Fu, Huang Huang, Yahav Avigal, Matthew Tancik, Jeffrey Ichnowski, Angjoo Kanazawa, and Ken Goldberg. Evo-nerf: Evolving nerf for sequential robot grasping of transparent objects. In *6th Annual Conference on Robot Learning*, 2022. [3](#)
- [29] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollár, and Ross Girshick. Segment anything. In *ICCV*, 2023. [3](#), [4](#), [5](#)
- [30] Sumith Kulal, Tim Brooks, Alex Aiken, Jiajun Wu, Jimei Yang, Jingwan Lu, Alexei A Efros, and Krishna Kumar Singh. Putting people in their place: Affordance-aware human insertion into scenes. In *CVPR*, 2023. [3](#)
- [31] Ashish Kumar, Saurabh Gupta, David F. Fouhey, Sergey Levine, and Jitendra Malik. Visual memory for robust path following. In *NIPS*, 2018. [1](#)

- [32] 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
- [33] Boyi Li, Kilian Q Weinberger, Serge Belongie, Vladlen Koltun, and Rene Ranftl. Language-driven semantic segmentation. In *ICLR*, 2022. 2
- [34] Gen Li, Varun Jampani, Deqing Sun, and Laura Sevilla-Lara. Locate: Localize and transfer object parts for weakly supervised affordance grounding. In *CVPR*, 2023. 1, 3, 4, 5, 6, 7, 8
- [35] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *ICML*, 2023. 2, 4
- [36] Yunzhu Li, Jiajun Wu, Russ Tedrake, Joshua B Tenenbaum, and Antonio Torralba. Learning particle dynamics for manipulating rigid bodies, deformable objects, and fluids. *arXiv preprint arXiv:1810.01566*, 2018. 3
- [37] Zeju Li, Chao Zhang, Xiaoyan Wang, Ruilong Ren, Yifan Xu, Ruifei Ma, and Xiangde Liu. 3dmit: 3d multi-modal instruction tuning for scene understanding. *arXiv preprint arXiv:2401.03201*, 2024. 2
- [38] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *ICCV*, 2017. 4
- [39] Chen Liu, Kihwan Kim, Jinwei Gu, Yasutaka Furukawa, and Jan Kautz. PlaneRCNN: 3D plane detection and reconstruction from a single image. In *CVPR*, 2019. 3
- [40] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. In *NeurIPS*, 2023. 1, 2, 3, 4, 8
- [41] Ruoshi Liu, Rundi Wu, Basile Van Hoorick, Pavel Tokmakov, Sergey Zakharov, and Carl Vondrick. Zero-1-to-3: Zero-shot one image to 3d object. In *ICCV*, 2023. 3
- [42] Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Chunyuan Li, Jianwei Yang, Hang Su, Jun Zhu, et al. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. *arXiv preprint arXiv:2303.05499*, 2023. 2
- [43] Timo Lüddecke and Alexander Ecker. Image segmentation using text and image prompts. In *CVPR*, 2022. 2
- [44] Hongchen Luo, Wei Zhai, Jing Zhang, Yang Cao, and Dacheng Tao. Grounded affordance from exocentric view. *arXiv preprint arXiv:2208.13196*, 2022. 1, 3, 4, 5, 6
- [45] Hongchen Luo, Wei Zhai, Jing Zhang, Yang Cao, and Dacheng Tao. Learning affordance grounding from exocentric images. In *CVPR*, 2022. 1, 2, 3, 4, 5, 6
- [46] Hongchen Luo, Wei Zhai, Jing Zhang, Yang Cao, and Dacheng Tao. Leverage interactive affinity for affordance learning. In *CVPR*, 2023. 3
- [47] Chuofan Ma, Yi Jiang, Xin Wen, Zehuan Yuan, and Xiaojuan Qi. Codet: Co-occurrence guided region-word alignment for open-vocabulary object detection. In *Advances in Neural Information Processing Systems*, 2023. 2
- [48] Matthias Minderer, Alexey Gritsenko, Austin Stone, Maxim Neumann, Dirk Weissenborn, Alexey Dosovitskiy, Aravindh Mahendran, Anurag Arnab, Mostafa Dehghani, Zhuoran Shen, et al. Simple open-vocabulary object detection. In *ECCV*, 2022. 2, 3, 4, 6, 8
- [49] Lorenzo Mur-Labadia, Jose J Guerrero, and Ruben Martinez-Cantin. Multi-label affordance mapping from ego-centric vision. In *ICCV*, 2023. 3
- [50] Tushar Nagarajan, Christoph Feichtenhofer, and Kristen Grauman. Grounded human-object interaction hotspots from video. In *ICCV*, 2019. 1, 3, 4, 6
- [51] Yinyu Nie, Xiaoguang Han, Shihui Guo, Yujian Zheng, Jian Chang, and Jian Jun Zhang. Total3dunderstanding: Joint layout, object pose and mesh reconstruction for indoor scenes from a single image. In *CVPR*, 2020. 3
- [52] OpenAI. Gpt-4 technical report, 2023. 2
- [53] 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. In *ICLR*, 2023. 2
- [54] Shengyi Qian and David F Fouhey. Understanding 3d object interaction from a single image. In *ICCV*, 2023. 3, 4, 5, 6, 7
- [55] Shengyi Qian, Linyi Jin, and David F. Fouhey. Associative3d: Volumetric reconstruction from sparse views. In *ECCV*, 2020. 6
- [56] Shengyi Qian, Alexander Kirillov, Nikhila Ravi, Devendra Singh Chaplot, Justin Johnson, David F Fouhey, and Georgia Gkioxari. Recognizing scenes from novel viewpoints. *arXiv preprint arXiv:2112.01520*, 2021. 3
- [57] 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 *ICML*, 2021. 2, 4
- [58] 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. 6
- [59] René Ranftl, Alexey Bochkovskiy, and Vladlen Koltun. Vision transformers for dense prediction. In *ICCV*, 2021. 3, 4, 8
- [60] Dandan Shan, Jiaqi Geng, Michelle Shu, and David Fouhey. Understanding human hands in contact at internet scale. In *CVPR*, 2020. 1
- [61] Cheng Shi and Sibeí Yang. Edadet: Open-vocabulary object detection using early dense alignment. In *ICCV*, 2023. 2
- [62] Mohit Shridhar, Lucas Manuelli, and Dieter Fox. Perceiver-actor: A multi-task transformer for robotic manipulation. In *Conference on Robot Learning*, pages 785–799. PMLR, 2023. 3
- [63] 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. 4
- [64] Peng Wang, An Yang, Rui Men, Junyang Lin, Shuai Bai, Zhikang Li, Jianxin Ma, Chang Zhou, Jingren Zhou, and Hongxia Yang. Ofa: Unifying architectures, tasks, and

- modalities through a simple sequence-to-sequence learning framework. In *ICML*, 2022. 2
- [65] Xiaolong Wang, David Fouhey, and Abhinav Gupta. Designing deep networks for surface normal estimation. In *CVPR*, 2015. 3
- [66] Ruihai Wu, Chuanruo Ning, and Hao Dong. Learning foresightful dense visual affordance for deformable object manipulation. In *ICCV*, 2023. 3
- [67] Xin Xu, Tianyi Xiong, Zheng Ding, and Zhuowen Tu. Masq-clip for open-vocabulary universal image segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 887–898, 2023. 2, 4
- [68] Jianing Yang, Xuweiyi Chen, Shengyi Qian, Nikhil Madaan, Madhavan Iyengar, David F Fouhey, and Joyce Chai. Llm-grounder: Open-vocabulary 3d visual grounding with large language model as an agent. In *ICRA*, 2024. 2
- [69] Yuhang Yang, Wei Zhai, Hongchen Luo, Yang Cao, Jiebo Luo, and Zheng-Jun Zha. Grounding 3d object affordance from 2d interactions in images. 2023. 3
- [70] Yufei Ye, Xueting Li, Abhinav Gupta, Shalini De Mello, Stan Birchfield, Jiaming Song, Shubham Tulsiani, and Sifei Liu. Affordance diffusion: Synthesizing hand-object interactions. In *CVPR*, 2023. 3
- [71] Sriram Yenamandra, Arun Ramachandran, Karmesh Yadav, Austin Wang, Mukul Khanna, Theophile Gervet, Tsung-Yen Yang, Vidhi Jain, Alexander William Clegg, John Turner, et al. Homerobot: Open-vocabulary mobile manipulation. *arXiv preprint arXiv:2306.11565*, 2023. 3
- [72] Wei Yin, Yifan Liu, Chunhua Shen, and Youliang Yan. Enforcing geometric constraints of virtual normal for depth prediction. In *ICCV*, 2019. 3
- [73] Wei Yin, Jianming Zhang, Oliver Wang, Simon Niklaus, Long Mai, Simon Chen, and Chunhua Shen. Learning to recover 3d scene shape from a single image. In *CVPR*, 2021. 3
- [74] Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. pixelnerf: Neural radiance fields from one or few images. In *CVPR*, 2021. 3
- [75] Yichi Zhang, Ziqiao Ma, Xiaofeng Gao, Suhaila Shakiah, Qiaozi Gao, and Joyce Chai. Groundhog: Grounding large language models to holistic segmentation. In *CVPR*, 2024. 2
- [76] Brady Zhou, Philipp Krähenbühl, and Vladlen Koltun. Does computer vision matter for action? *Science Robotics*, 4(30): eaaw6661, 2019. 3
- [77] Gengze Zhou, Yicong Hong, and Qi Wu. Navgpt: Explicit reasoning in vision-and-language navigation with large language models. *arXiv preprint arXiv:2305.16986*, 2023. 2