

Semantic Abstraction: Open-World 3D Scene Understanding from 2D Vision-Language Models

Huy Ha Shuran Song
Columbia University
semantic-abstraction.cs.columbia.edu

Abstract: We study open-world 3D scene understanding, a family of tasks that require agents to reason about their 3D environment with an open-set vocabulary and out-of-domain visual inputs – a critical skill for robots to operate in the unstructured 3D world. Towards this end, we propose Semantic Abstraction (SemAbs), a framework that equips 2D Vision-Language Models (VLMs) with new 3D spatial capabilities, while maintaining their zero-shot robustness. We achieve this abstraction using relevancy maps extracted from CLIP, and learn 3D spatial and geometric reasoning skills on top of those abstractions in a semantic-agnostic manner. We demonstrate the usefulness of SemAbs on two open-world 3D scene understanding tasks: 1) completing partially observed objects and 2) localizing hidden objects from language descriptions. SemAbs can generalize to novel vocabulary for object attributes and nouns, materials/lighting, classes, and domains (i.e., real-world scans) from training on limited 3D synthetic data.

Keywords: 3D scene understanding, out-of-domain generalization, language




Figure 1: **Open-World 3D Scene Understanding.** Our approach, Semantic Abstraction, unlocks 2D VLM’s capabilities to 3D scene understanding. Trained with a small simulated data, our model generalizes to unseen classes in a novel domain (i.e., real world scans), from small objects like “rubiks cube”, to long-tail concepts like “harry potter”, to hidden objects like the “used N95s in the garbage bin”.

1 Introduction

To assist people in their daily life, robots must recognize a large, context-dependent, and dynamic set of semantic categories from their visual sensors. However, if their 3D training environment contains only *common* categories like desks, office chairs, and bookshelves, how could they represent and localize *long-tail* objects, such as the “used N95s in the garbage bin” (Fig 1)?

This is an example of an open-world 3D scene understanding task (Fig. 2), a family of 3D vision-language tasks which encompasses open-set classification (evaluation on novel classes) with extra generalization requirements to novel vocabulary (i.e., object attributes, synonyms of object nouns), visual properties (e.g. lighting, textures), and domains (e.g. sim v.s. real). The core challenge in such tasks is the limited data: all existing 3D datasets [1–6] are limited in diversity and scale compared to their internet-scale 2D counterparts [7–9], so training on them does not prepare robots for the open 3D world.

On the other hand, large-scale 2D vision language models (VLM) [7–11] have demonstrated impressive performance on open-set image classification. Through exposure to many image-caption pairs from the internet, they have acquired an unprecedented level of robustness to visual distributional shifts and a large repertoire of visual-semantic concepts and vocabulary. However, fine-tuning these pretrained models reduces their open-set classification robustness [7, 10, 12, 13]. Compared to internet-scale image-caption data, the finetuning datasets, including the existing 3D datasets [1–6], show a significant reduction in scale

and diversity, which causes finetuned models to become task-specialized and lose their generality. Hence, we investigate the following question:

*How can we equip 2D VLMs with new 3D capabilities,
while maintaining their zero-shot robustness?*

We propose **Semantic Abstraction (SemAbs)**, a framework for tackling visual-semantic reasoning in open-world 3D scene understanding tasks using 2D VLMs. We hypothesize that, while open-world visual-semantic reasoning requires exposure to internet-scale datasets, 3D spatial and geometric reasoning is tractable even with a limited synthetic dataset, and could generalize better if learned in a semantic-agnostic manner. For instance, instead of learning the concept of “behind the Harry Potter book”, the 3D localization model only needs to learn the concept of “behind *that object*”.

To achieve this abstraction, we leverage the relevancy maps extracted from 2D VLMs and used them as the “abstracted object” representation that is agnostic to their semantic labels. The SemAbs module factorizes into two submodules: 1) A semantic-aware wrapper that takes an input RGB-D image and object category label and outputs the relevancy map of a pre-trained 2D VLM model (i.e., CLIP [7]) for that label, and 2) a semantic-abstracted 3D module that uses the relevancy map to predict 3D occupancy. This 3D occupancy can either represent the 3D geometry of a partially observed object or the possible 3D locations of a hidden object (e.g. mask in the trash) conditioned on the language input.

While we only train the 3D network on a limited synthetic 3D dataset, it generalizes to any novel semantic labels, vocabulary, visual properties, and domains that the 2D VLM can generalize to. As a result, the SemAbs module inherits the VLM’s visual robustness and open-set classification abilities while adding the spatial reasoning abilities it lacks [14–16]. In summary, our contributions are three-fold¹:

- **Semantic Abstraction**, a framework for augmenting 2D VLMs with 3D reasoning capabilities for open-world 3D scene understanding tasks. By abstracting semantics away using relevancy maps, our SemAbs module generalizes to novel semantic labels, vocabulary, visual properties, and domains (i.e., sim2real) despite being trained on only a limited synthetic 3D dataset.
- **Efficient Multiscale Relevancy Extraction**. To support Semantic Abstraction, we propose a multi-scale relevancy extractor for vision transformers [17] (ViTs), which robustly detects small, long-tail objects and achieves over $\times 60$ speed up from prior work [18].
- Two novel **Open-world 3D Scene Understanding** tasks (open-vocabulary semantic scene completion and visually obscured object localization), a data generation pipeline with AI2-THOR [1] simulator for the tasks, and a systematic open-world evaluation procedure. We believe these two tasks are important primitives for bridging existing robotics tasks to the open-world.

2 Related Works

2D Visual Language Models. The recent advancements in scaling up contrasting learning have enabled the training of large 2D vision language models (VLM). Through their exposure to millions of internet image-caption pairs, these VLMs [7, 9–11] acquire a remarkable level of zero-shot robustness – they can recognize a diverse set of long-tail semantic concepts robustly under distributional shifts. However, the learned image-level representations are not directly applicable to spatial-reasoning tasks which require pixel-level information. To address this issue, approaches were proposed to extract dense features [19–21] or relevancy maps [18, 22]. These advancements complement our contribution, which leverages such techniques for interfacing 2D VLMs to open-world 3D scene understanding tasks.

Transferring 2D VLMs to 2D Applications. Encouraged by their impressive capabilities, many downstream tasks build around these pretrained VLMs [14, 23–28], often finetuning them (either end-to-end or train learnable weights on top of the pretrained encoder) on a small task-specific dataset. However, many recent studies show that finetuning these VLMs significantly weakens their zero-shot robustness [7, 10, 12, 13, 19], which motivates a new paradigm for using pretrained models. By combining a *fixed* pretrained VLM with prompting [29] or attention extraction [15, 19–21], the new generation of zero-shot transfer techniques can inherit the VLM’s robustness in their downstream tasks without overfitting to any visual domain. However, all existing zero-shot transfer techniques have been confined to the realms of 2D vision.

Closed-World 3D Scene Understanding. Understanding the semantics and 3D structure of an unstructured environment is a fundamental capability for robots. However, prior works in 3D object detection [30–32], semantic scene completion [33–35], and language-informed object localization [6, 36] are typically concerned with limited visual diversity and object categories, limitations imposed by their 3D training data. Despite great efforts from the community [1–6], existing 3D datasets’ scale, diversity,

¹All code, data, and models is publicly available at <https://github.com/columbia-ai-robotics/semantic-abstraction>.

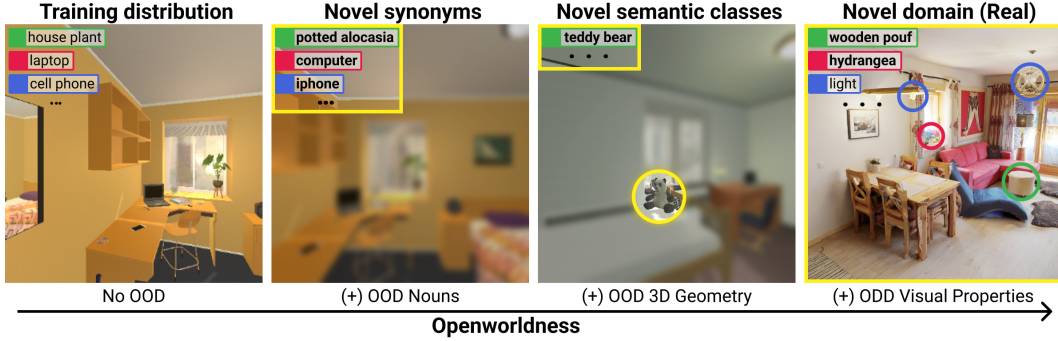


Figure 2: **Open-world generalization requirements** can build on top of each other, forming tiers of open-worldness (distinct property of each tier outlined yellow).

and coverage pale in comparison to internet-scale image-text pairs [7–9]. However, our key insight is that visual-semantic reasoning can be learned separately from 3D reasoning. That is, with the right abstraction, the complex visual-semantic reasoning can be offloaded to 2D VLMs, while the 3D model specializes in semantic-agnostic spatial and geometry reasoning.

3 Method: Semantic Abstraction

Evidence from linguistics and cognitive science suggests that semantic concepts are much more diverse than spatial and geometric concepts. While the average adult knows roughly 10,000 nouns for object categories, there are only 90 words in English for describing spatial relations [37]. In fact, there are so few spatial prepositions that they are typically considered a closed-class vocabulary [38]. Similarly for 3D geometry, the idea that our 3D world can be decomposed into sets of primitives [39] has been applied in many works, such as generalized cylinders [40] and block worlds [41]. In contrast, the space of nouns for semantic concepts is large and often cannot be further decomposed into common primitives. This suggests that while we may need a large and diverse dataset for semantic and visual concept learning, a relatively small 3D dataset could cover spatial and geometric reasoning. This observation motivates our approach, Semantic Abstraction.

3.1 Abstraction via Relevancy

The Semantic Abstraction (SemAbs) module (Fig 3c) takes as input an RGB-D image $\mathcal{I} \in \mathbb{R}^{H \times W}$ and an object class text label \mathcal{T} (e.g. “biede armchair”) and outputs the 3D occupancy \mathcal{O} for objects of class \mathcal{T} . It factorizes into two submodules:

The semantic-aware wrapper (Fig 3c, green background) abstracts \mathcal{I} and \mathcal{T} into a relevancy map $\mathcal{R} \in \mathbb{R}^{H \times W}$, where each pixel’s value denotes that pixel’s contribution to the VLM’s classification score for \mathcal{T} . Introduced for model explainability [18, 22], relevancy maps can be treated as a coarse localization of the text label. Using the depth image and camera matrices, \mathcal{R} is projected into a 3D point cloud, $\mathcal{R}^{\text{proj}} = \{r_i\}_{i=1}^{H \times W}$ where $r_i \in \mathbb{R}^4$ (a 3D location with a scalar relevancy value). Only $\mathcal{R}^{\text{proj}}$, but neither the text \mathcal{T} nor the image \mathcal{I} , is passed to the second submodule.

The semantic-abstracted 3D module (Fig 3c, grey background) treats the relevancy point cloud as the localization of a partially observed object and completes it into that object’s 3D occupancy. To do this, we first scatter $\mathcal{R}^{\text{proj}}$ into a 3D voxel grid $\mathcal{R}^{\text{vox}} \in \mathbb{R}^{D \times 128 \times 128 \times 128}$. Then, we encode \mathcal{R}^{vox} as a 3D feature volume: $f_{\text{encode}}(\mathcal{R}^{\text{vox}}) \mapsto Z \in \mathbb{R}^{D \times 128 \times 128 \times 128}$, where f_{encode} is a 3D UNet [42] with 6 levels. Z can be sampled using trilinear interpolation to produce local features $\phi_q^Z \in \mathbb{R}^D$ for any 3D query point $q \in \mathbb{R}^3$. Finally, decoding ϕ_q^Z with a learned MLP f_{decode} gives us an occupancy probability for each point q , $f_{\text{decode}}(\phi_q^Z) \mapsto o(q) \in [0, 1]$. In this submodule, only f_{encode} and f_{decode} (Fig. 3c, yellow boxes) are trained with the 3D dataset.

Although the semantic-abstracted 3D module only observes relevancy abstractions of the semantic label \mathcal{T} but not \mathcal{T} itself, the semantic-abstracted 3D module’s output can be interpreted as the 3D occupancy for \mathcal{T} . This means it generalizes to any semantic label that can be recognized by the 2D VLMs’ relevancy maps even if it was trained on a limited 3D dataset. In our implementation, we use CLIP [7] as our VLM. However, our framework is *VLM-agnostic*, as long as relevancy maps can be generated. It is an interesting future direction to investigate how different VLMs [7–11], their training procedure [43], and different relevancy approaches [18, 22, 44, 45] affect the performance of different downstream 3D scene understanding tasks. In the next sections, we explain how to ensure the visual-semantic concepts are reliably recognized by the relevancy maps (Sec. 3.2), and how to apply SemAbs to 3D scene understanding tasks (Sec. 3.3, Sec. 3.4).

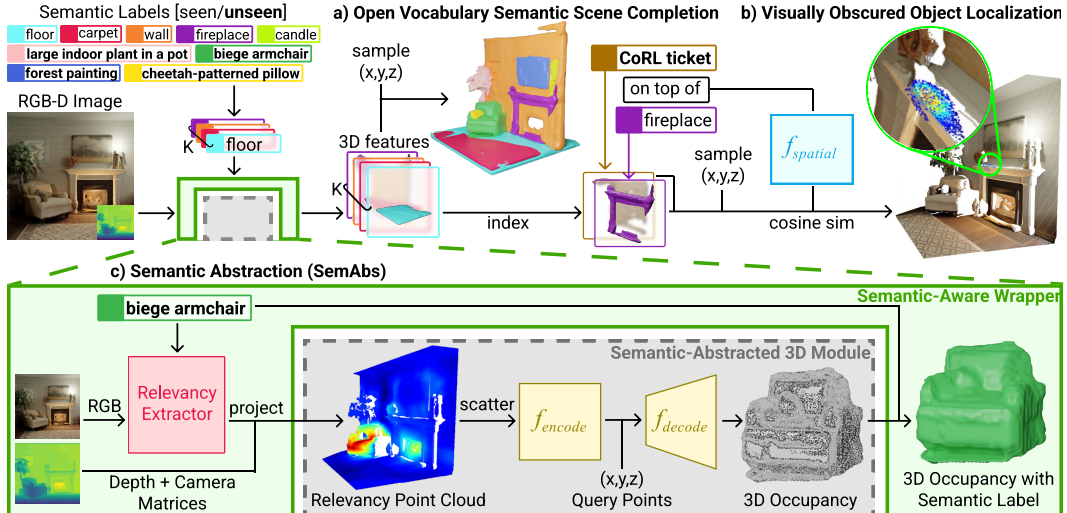


Figure 3: **Semantic Abstraction Overview.** Our framework can be applied to open-world 3D scene understanding tasks (a-b) using the SemAbs module (c). It consists of a **semantic-aware wrapper** (green background) that abstracts the input image and semantic label into a relevancy map, and a **semantic-abstracted 3D module** (grey background) that completes the projected relevancy map into a 3D occupancy. This abstraction allows our approach to generalize to long-tail semantic labels unseen (bolded) during 3D training, such as the “CoRL ticket on top of the fireplace”.

3.2 A Multi-Scale Relevancy Extractor

We choose to use ViT-based CLIP [7] due to their superior performance over the ResNet variants. However, existing ViT relevancy techniques [18] often produce noisy, low-resolution maps that highlight irrelevant regions or miss small objects [27, 46].

To reduce noise from irrelevant regions, we use horizontal flips and RGB augmentations [27] of \mathcal{I} and obtain a relevancy map for each augmentation. To reliably detect small objects, we propose a multi-scale relevancy extractor that *densely* computes relevancies at different scales and locations in a sliding window fashion (an analogous relevancy-based approach to Li et al.’s local attention-pooling). The final relevancy map is averaged across all augmentations and scales.




Figure 4: Our relevancy extractor robustly detects even small, long-tail objects, like the “nintendo switch”.

Qualitatively, while augmentations help in reducing irrelevant regions (e.g. Fig. 4, from localizing the entire fireplace in the 1st column to only the vase in the 2nd column), they still often miss small objects. However, with our multi-scale relevancy extractor, our relevancy maps can detect small objects such as the tiny “nintendo switch” (Fig. 4, bottom right). To support efficient multi-scale relevancy extraction, we implemented a batch-parallelized ViT relevancy extractor. At 0.4 seconds per text label on average², we achieve over $\times 60$ speed up compared to the non-parallelized implementation. With the rise in popularity of ViTs, we believe this efficient ViT relevancy extractor implementation will be a useful primitive for the research community and have released it on [Github](#) and a [Hugging Face Spaces](#).

3.3 Application to Open Vocabulary Semantic Scene Completion (OVSSC)

Task. Given as input a single RGB-D image \mathcal{I} and a set of K object class labels, represented as open-vocabulary text $\{\mathcal{T}_k\}_{k=1}^K$, the task is to complete the 3D geometry of all the partially observed objects referred by $\{\mathcal{T}_k\}_{k=1}^K$. In contrast to its closed-world variant [33], where the object list is fixed for all scenes, OVSSC formulation accounts for context-dependency, which allows the class labels to change from scene to scene and from training to testing. From this 3D semantic scene completion, robots can use the detailed 3D geometry for low level planning such as grasping and motion planning.

²measured over 10 repetitions with 100 text labels on a GTX 1080 Ti machine with a light load. We provide extra details in the appendix.

Applying SemAbs to OVSSC is a matter of applying the SemAbs module K times, once per object label. First, the semantic-aware wrapper abstracts \mathcal{I} and $\{\mathcal{T}_k\}_{k=1}^K$ into K relevancy point clouds $\{\mathcal{R}_k^{\text{proj}}\}_{k=1}^K$. Then, the semantic-abstracted 3D module extracts K feature volumes $\{Z_k\}_{k=1}^K$, used to decode K 3D occupancies $\{\mathcal{O}_k\}_{k=1}^K$. The final OVSSC output is a point-wise argmax of occupancy probability over the K classes. Query points with the maximum occupancy probability less than a threshold c are labeled as empty. We provide data generation and training details in the appendix.

3.4 Application to Visually Obscured Object Localization (VOOL)

Motivation. Even after completing partially visible objects in the scene, OVSSC’s output may not be complete – the scene may contain visually obscured or hidden objects (e.g., the ticket on the fireplace in Fig. 3b). In such cases, the human user can provide additional hints to the robot using natural language to help localize the object, which motivates our VOOL task.

Task. Given as input a single RGB-D image \mathcal{I} of a 3D scene and a description \mathcal{D} of an object’s location in the scene, the task is to predict the possible occupancies of the referred object. Similar to prior works [47], we assume that the description is in the standard linguistic representation of an object’s place [37]. Specifically, $\mathcal{D} = \langle \mathcal{T}_{\text{target}}, \mathcal{S}, \mathcal{T}_{\text{ref}} \rangle$, where the open-vocabulary target object label $\mathcal{T}_{\text{target}}$ (e.g. “rapid test”) is the visually obscured object to be located, the open-vocabulary reference object label \mathcal{T}_{ref} (e.g. “harry potter book”) is presumably a visible object, and the closed-vocabulary spatial preposition \mathcal{S} describes $\mathcal{T}_{\text{target}}$ ’s location with respect to \mathcal{T}_{ref} (e.g. “behind”). Unlike the referred expression 3D object localization task [6] which assumes full visibility (i.e., fused pointcloud), VOOL takes as input only the current RGB-D image, and thus is more suited for a dynamic, partially observable environment. Further, in contrast to bounding boxes, VOOL’s occupancy can be disjoint regions in space, which accounts the multi-modal uncertainty when objects are hidden. Downstream robotic applications, such as object searching, benefit more from the uncertainty information encoded in a VOOL’s output than bounding boxes.

Applying SemAbs to VOOL. We propose to learn an embedding $f_{\text{spatial}}(\mathcal{S}) \in \mathbb{R}^{2D}$ for each spatial relation (Fig. 3b, blue box), since the set of spatial prepositions is small and finite [37], and use them as follows. As in OVSSC (§3.3), we can get occupancy feature volumes Z_{target} and Z_{ref} for the target and reference objects respectively. Given a set of fixed query points $Q = \{q_i\}_{i=1}^N$, we extract their local feature point clouds $\phi_Q^{Z_{\text{target}}}, \phi_Q^{Z_{\text{ref}}} \in \mathbb{R}^{N \times D}$, then concatenate them to get a new feature pointcloud $\phi_Q^{Z_{\text{target}} \parallel Z_{\text{ref}}} \in \mathbb{R}^{N \times 2D}$. Finally, to get the 3D localization occupancy \mathcal{O} , we perform a point-wise cosine-similarity between $\phi_Q^{Z_{\text{target}} \parallel Z_{\text{ref}}}$ and $f_{\text{spatial}}(\mathcal{S})$. For this task, we only need to learn the spatial embeddings f_{spatial} conditioned on the occupancy information encoded in Z_{target} and Z_{ref} , since semantic reasoning of $\mathcal{T}_{\text{target}}$ and \mathcal{T}_{ref} has been offloaded to the SemAbs (and, therefore, to CLIP). We provide data generation and training details in the appendix.

4 Experiments

We design experiments to systemically investigate Semantic Abstraction’s open-world generalization abilities to novel³ concepts in 3D scene understanding. Our benchmark include the following categories: 1) **Novel Rooms**: we follow AI2-THOR [1]’s split, which holds out 20 rooms for as novel rooms, while training on 100 rooms. 2) **Novel Visual**: we use AI2-THOR’s randomize material and lighting functionality to generate more varied 2D visual conditions of novel rooms. 3) **Novel Synonyms**: we manually picked out 17 classes which had natural synonyms (e.g. “pillow” to “cushion”, complete list included in appendix) and replace their occurrences in text inputs (i.e., object label in OVSSC, description in VOOL) to these synonyms. 4) **Novel Class**: we hold out 6 classes out of the 202 classes in AI2-THOR. For OVSSC, any view which contains one of these novel classes is held out for testing. For VOOL, any scene with a description that contains a novel class (as either the target or reference object) is held out. 5) **Novel Domain**: We also provide qualitative evaluations (Fig. 5) using real world scans from HM3D [4].

Metrics & Baselines. For both tasks, we measure the learner’s performance with voxel IoU of dimension $32 \times 32 \times 32$, where the ground truth of each task is generated as described in §3.3 for OVSSC and §3.4 for VOOL. We compare with following categories of baselines:

- **Semantic-Aware (SemAware).** To evaluate the effectiveness of semantic abstraction, we design baselines that perform visual-semantic reasoning themselves by taking RGB point clouds as input

³Here, we define “novel” as to whether a concept was observed during the 3D training process.

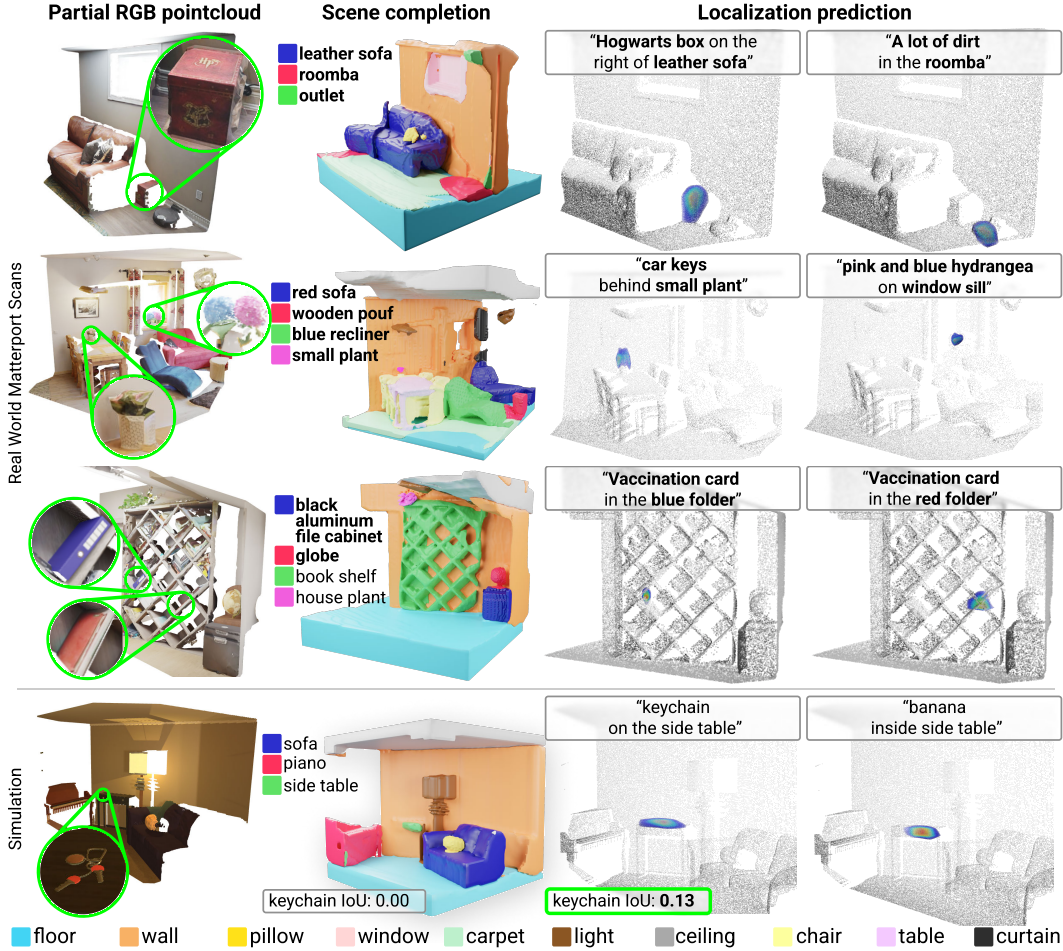


Figure 5: **Semantic Abstraction inherits CLIP’s visual-semantic reasoning skills** From distinguishing colors (e.g. “red folder” v.s. “blue folder”) to recognizing cultural (e.g. “hogwarts box”) and long-tail semantic concepts (e.g. “roomba”, “hydrangea”), our approach offloads such visual-semantic reasoning challenges to CLIP. In doing so, its learned 3D spatial and geometric reasoning skills transfers sim2real in a zero-shot manner.

- (instead of relevancy point clouds). For OVSSC, the baseline’s feature point cloud is supervised with BCE to be aligned (*i.e.*, using cosine similarity) with the correct object class text embedding. Similarly, for VOOL, the feature point cloud’s cosine similarity with an embedding of the entire localization description (e.g. “N95s inside the trash”) is supervised using BCE to the target object occupancy.
- **Semantic & Spatial Abstraction (ClipSpatial).** What if we off-load both 3D and visual-semantic reasoning to 2D VLMs (instead of just the latter, as in our approach)? To answer this, we designed CLIPSpatial, a VOOL baseline that uses relevancy maps for the entire description (e.g. “N95s inside the trash”) as input. We expect this approach to perform poorly since it has been demonstrated that current 2D VLMs [7, 8] struggles with spatial relations [14–16, 48].
 - **Naive Relevancy Extraction. (SemAbs + [18])** To investigate the effects of relevancy quality, we replace our multi-scale relevancy extractor with [Chefer et al. \[18\]](#)’s approach.

The first two baselines can be seen as two extremes on a spectrum of how much we rely on pretrained 2D VLMs. While SemAware learns both semantics and 3D spatial reasoning, CLIPSpatial delegates both to pretrained 2D VLMs. Our approach is in between these two extremes, designed such that it builds on 2D VLMs’ visual robustness and open-classification strengths while addressing its spatial reasoning weaknesses.

4.1 Open-world Evaluation Results

Tab. 1 and 2 summarize the quantitative results, and Fig. 5 shows real world qualitative results tested on Matterport scenes [4]. More results and comparisons can be found on the [project website](#).

Semantic Abstraction simplifies open-world generalization. SemAware baselines especially struggle to generalize to novel semantic classes (Tab. 1, Tab. 2).




Figure 6: **Qualitative result on NYUCAD [49, 50] dataset.** Our model can distinguish between “puppy toy”, “ipod”, “certificate”, and “hand sculpture”, while all prior works would’ve output “obj”.

We hypothesize that observing text embeddings during training causes these baselines to specialize to this distribution of embeddings, such that some synonyms and most novel class labels are out-of-distribution inputs. Our approach significantly outperforms the SemAware baseline across the board, suggesting that Semantic Abstraction not only simplifies open-world generalization but is also a strong inductive bias for learning the task.

Semantic & Spatial Abstraction generalizes poorly. The CLIPSpatial baseline achieves 8.0-9.0 IoU worse than our approach in all generalization scenarios (Tab. 2). This indicates that our design of learning spatial reasoning instead of relying on CLIP was crucial to semantic abstraction’s performance.

Semantic Abstraction requires high quality relevancy maps. Intuitively, the model can only output high quality completions if its relevancy maps are also high quality (*i.e.*, highlights the target object, especially when the object is small). Our average IoU is $\times 1.5$ (novel room, visual, synonyms, Table 1), over $\times 2.0$ (novel classes, Table 1, novel room, visual, synonyms, Table 2) or $\times 5.0$ (novel class, Table 2) compared to SemAbs + [18], which demonstrates the importance of our relevancy extractor in a successful application of Semantic Abstraction.

Inheriting VLM zero-shot robustness enables Sim2Real transfer. Despite training in simulation, our approach can perform completion and localization from real-world matterport scans [4] (Fig. 1 and 5). From novel vocabulary for describing materials (*e.g.* “leather sofa”, “wooden pouf”), colors (*e.g.* “red folder”, “blue folder”), conjunctions of them (*e.g.* “black aluminum file cabinet”) to novel semantic classes (*e.g.* “roomba”, “Hogwarts box”), our approach unlocks CLIP’s visual robustness and large repertoire of visual-semantic concepts to 3D scene understanding. When OVSSC’s completion misses objects (*e.g.* “hydrangea”, “keychain”), VOOL can still localize them (*e.g.* +0.13 IoU for keychain).

4.2 Zero-shot Evaluation Results

The NYUv2 CAD dataset [49, 50] is a popular closed-vocabulary SSC benchmark. Although it contains 894 object classes, prior works [33, 51, 52] collapse it down to 11 classes to avoid the challenge in learning the “long-tail” categories (*e.g.* “puppy toy”, “ipod”, “hand sculpture”). We evaluate our model zero-shot on NYUv2 directly on **all 894 classes** as well as the 11 classes for comparison (Fig. 6, Tab. 3, Ours). We also train a SemAbs model from scratch in NYU to investigate the generalization gap between our dataset and NYUv2 (Tab. 3, Ours NYU). We use the same training/testing split [33]. Similar to prior works [33], we report volumetric IoU at a voxel dimension of 60. As in our THOR dataset, we only train f_{encode} and f_{decode} in the semantic abstracted 3D module, not the 2D VLM. To our knowledge, our approach is the first to output the full list of 894 object categories in this benchmark.

Approach	Novel			
	Room	Visual	Synonyms	Class
SemAware	32.2	31.9	20.2	0.0
SemAbs+[18]	26.6	24.3	17.8	12.2
Ours	40.1	36.4	33.4	37.9

Table 1: **Open-Vocab Semantic Scene Completion**

Approach	Novel			
	Room	Visual	Synonyms	Class
SemAware	12.1	11.9	11.4	3.0
SemAbs+[18]	9.1	8.8	10.7	4.0
CLIPSpatial	11.7	10.1	14.3	11.2
Ours	20.9	19.2	23.4	19.7

Table 2: **Visually Obscured Object Localization**

Approach	ceil.	floor	wall	win.	chair	bed	sofa	table	tv	furn.	objs.	avg.
SSCNet [33]	32.5	92.6	40.2	8.9	33.9	57	59.5	28.3	8.1	44.8	25.1	36.6
SISNet [51]	63.4	94.4	67.2	52.4	59.2	77.9	71.1	51.8	46.2	65.8	48.8	63.5
Ours Supervised	22.6	46.1	33.9	35.9	23.9	55.9	37.9	19.7	30.8	39.8	27.7	34.0
Ours Zeroshot	13.7	17.3	13.5	25.2	15.2	33.3	31.5	12.0	23.7	25.6	19.9	21.0

Table 3: NYUv2 CAD Semantic Scene Completion.

Open-world versus Closed-world formulation. Our work investigates open-world 3D scene understanding, which means specializing to one data distribution is not our focus. However, we hope these results on a standard closed-world benchmark along with our discussion below are helpful to the research community in understanding practical trade-offs in algorithm and evaluation design between the open-world versus closed-world formulation. Below summarizes our key findings:

- + **Flexibility to object labels.** Our framework directly outputs the full 894 object categories, of which 108 do not appear in the training split, and its predictions can map back to the collapsed 11 classes. This has never been attempted by prior works. This flexibility is achieved by building on top of 2D VLMs.
- + **Robust detection of small, long-tail object categories.** “Due to the heavy-tailed distribution of object instances in natural images, a large number of objects will occur too infrequently to build usable recognition models” [41]. However, our approach can robustly detect not just long-tail object classes, but also small ones – challenging scenarios for perception systems. For instance, from Fig. 6, our multi-scale relevancy extractor can output highly localized relevancy maps around “puppy toy”, “ipod”, and “hand sculpture”.
- **Inability to finetune large pretrained models.** As shown in prior works [7, 10, 12, 13, 19], directly finetuning 2D VLMs on small datasets hurts their robustness. We cannot finetune the visual encoders (and thus, its relevancy activations) to specialize on the visual and class distribution on a small dataset (e.g., NYUv2) without sacrificing its robustness and generality. This is an open and important research direction and recent work [53] has demonstrated promising results in image classification.
- **Bias from image caption data.** The typical training paradigm of current 2D VLMs [7, 10, 11] is contrastive learning over image-caption pairs on the internet. While their dataset covers a large number of semantic concepts, they still have their limits. For instance, we observed that for “wall” and “ceiling” – two classes which closed-world approaches typically achieve high performance for – our approach achieved a much lower IoU of only 33.9 and 22.6. Since these classes are often not the subject of internet captions, our approach inherits this bias from 2D VLMs. However, our framework is **VLM-agnostic**, which means it can continue to benefit from future and more powerful versions of 2D VLMs.

4.3 Limitations and assumptions

Our key assumption is that geometry and spatial reasoning can be done in a semantic agnostic manner. This assumption does not hold when we care about completing detailed object geometry, which is often semantic-dependent. However, we believe that coarse 3D reasoning with paired semantic labels may already be useful for open-world robotic applications such as object retrieval. Second, while we have assumed a simple description syntax and a small set of spatial prepositions for object localization, incorporating natural language understanding (*i.e.*, to support open-set phrases for richer spatial descriptions) into open-world 3D scene understanding is an important extension of our work. Lastly, we assumed that relevancy maps highlight only relevant regions. Limited by current VLMs, our relevancy extractor can sometimes fail to distinguish highly related concepts (*e.g.* book v.s. bookshelf). However, since Semantic Abstraction is VLM-agnostic, this limitation could be addressed by using other VLMs.

5 Conclusion

We proposed Semantic Abstraction, a framework that equips 2D VLMs with 3D spatial capabilities for open-world 3D scene understanding tasks. By abstracting RGB and semantic label inputs as relevancy maps, our learned 3D spatial and geometric reasoning skills generalize to novel vocabulary for object attributes and nouns, visual properties, semantic classes, and domains. Since scene completion and 3D object localization are fundamental primitives for many robotic tasks, we believe our approach can assist in bridging currently close-world robotics tasks, such as object retrieval, into the open-world formulation. More importantly, we hope our framework encourages roboticists to embrace building systems around large pretrained models, as a path towards open-world robotics.

Acknowledgments: We would like to thank Samir Yitzhak Gadre, Cheng Chi and Zhenjia Xu for their helpful feedback and fruitful discussions. This work was supported in part by NSF Award #2143601, #2132519 JP Morgan Faculty Research Award, and Google Research Award. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the sponsors.

References

- [1] E. Kolve, R. Mottaghi, W. Han, E. VanderBilt, L. Weihs, A. Herrasti, D. Gordon, Y. Zhu, A. Gupta, and A. Farhadi. AI2-THOR: An Interactive 3D Environment for Visual AI. *arXiv*, 2017.
- [2] S. Song, S. P. Lichtenberg, and J. Xiao. Sun rgb-d: A rgb-d scene understanding benchmark suite. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 567–576, 2015.
- [3] A. X. Chang, T. Funkhouser, L. Guibas, P. Hanrahan, Q. Huang, Z. Li, S. Savarese, M. Savva, S. Song, H. Su, et al. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*, 2015.
- [4] S. K. Ramakrishnan, A. Gokaslan, E. Wijmans, O. Maksymets, A. Clegg, J. M. Turner, E. Undersander, W. Galuba, A. Westbury, A. X. Chang, M. Savva, Y. Zhao, and D. Batra. Habitat-matterport 3d dataset (HM3d): 1000 large-scale 3d environments for embodied AI. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*, 2021. URL <https://openreview.net/forum?id=-v4OuqNs5P>.
- [5] A. Szot, A. Clegg, E. Undersander, E. Wijmans, Y. Zhao, J. Turner, N. Maestre, M. Mukadam, D. Chaplot, O. Maksymets, A. Gokaslan, V. Vondrus, S. Dharur, F. Meier, W. Galuba, A. Chang, Z. Kira, V. Koltun, J. Malik, M. Savva, and D. Batra. Habitat 2.0: Training home assistants to rearrange their habitat. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- [6] D. Z. Chen, A. X. Chang, and M. Nießner. Scanrefer: 3d object localization in rgb-d scans using natural language. In *European Conference on Computer Vision*, pages 202–221. Springer, 2020.
- [7] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021.
- [8] J. Li, R. Selvaraju, A. Gotmare, S. Joty, C. Xiong, and S. C. H. Hoi. Align before fuse: Vision and language representation learning with momentum distillation. *Advances in Neural Information Processing Systems*, 34, 2021.
- [9] J.-B. Alayrac, J. Donahue, P. Luc, A. Miech, I. Barr, Y. Hasson, K. Lenc, A. Mensch, K. Millican, M. Reynolds, et al. Flamingo: a visual language model for few-shot learning. *arXiv preprint arXiv:2204.14198*, 2022.
- [10] H. Pham, Z. Dai, G. Ghiasi, H. Liu, A. W. Yu, M.-T. Luong, M. Tan, and Q. V. Le. Combined scaling for zero-shot transfer learning. *arXiv preprint arXiv:2111.10050*, 2021.
- [11] C. Jia, Y. Yang, Y. Xia, Y.-T. Chen, Z. Parekh, H. Pham, Q. Le, Y.-H. Sung, Z. Li, and T. Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *International Conference on Machine Learning*, pages 4904–4916. PMLR, 2021.
- [12] M. Wortsman, G. Ilharco, M. Li, J. W. Kim, H. Hajishirzi, A. Farhadi, H. Namkoong, and L. Schmidt. Robust fine-tuning of zero-shot models. *arXiv preprint arXiv:2109.01903*, 2021.
- [13] A. Andreassen, Y. Bahri, B. Neyshabur, and R. Roelofs. The evolution of out-of-distribution robustness throughout fine-tuning. *arXiv preprint arXiv:2106.15831*, 2021.
- [14] S. Subramanian, W. Merrill, T. Darrell, M. Gardner, S. Singh, and A. Rohrbach. Reclip: A strong zero-shot baseline for referring expression comprehension. *arXiv preprint arXiv:2204.05991*, 2022.

- [15] S. Y. Gadre, M. Wortsman, G. Ilharco, L. Schmidt, and S. Song. Clip on wheels: Zero-shot object navigation as object localization and exploration. *arXiv preprint arXiv:2203.10421*, 2022.
- [16] A. Ramesh, P. Dhariwal, A. Nichol, C. Chu, and M. Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022.
- [17] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- [18] H. Chefer, S. Gur, and L. Wolf. Generic attention-model explainability for interpreting bi-modal and encoder-decoder transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 397–406, 2021.
- [19] C. Zhou, C. C. Loy, and B. Dai. Denseclip: Extract free dense labels from clip. *arXiv preprint arXiv:2112.01071*, 2021.
- [20] Y. Rao, W. Zhao, G. Chen, Y. Tang, Z. Zhu, G. Huang, J. Zhou, and J. Lu. Denseclip: Language-guided dense prediction with context-aware prompting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18082–18091, 2022.
- [21] J. Li, G. Shakhnarovich, and R. A. Yeh. Adapting clip for phrase localization without further training. *arXiv preprint arXiv:2204.03647*, 2022.
- [22] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, pages 618–626, 2017.
- [23] R. Mokady, A. Hertz, and A. H. Bermano. Clipcap: Clip prefix for image captioning. *arXiv preprint arXiv:2111.09734*, 2021.
- [24] H. Song, L. Dong, W.-N. Zhang, T. Liu, and F. Wei. Clip models are few-shot learners: Empirical studies on vqa and visual entailment. *arXiv preprint arXiv:2203.07190*, 2022.
- [25] Z. Gao, J. Liu, S. Chen, D. Chang, H. Zhang, and J. Yuan. Clip2tv: An empirical study on transformer-based methods for video-text retrieval. *arXiv preprint arXiv:2111.05610*, 2021.
- [26] M. Wang, J. Xing, and Y. Liu. Actionclip: A new paradigm for video action recognition. *arXiv preprint arXiv:2109.08472*, 2021.
- [27] N. Zabari and Y. Hoshen. Semantic segmentation in-the-wild without seeing any segmentation examples. *CoRR*, abs/2112.03185, 2021. URL <https://arxiv.org/abs/2112.03185>.
- [28] M. Shridhar, L. Manuelli, and D. Fox. Cliport: What and where pathways for robotic manipulation. In *Proceedings of the 5th Conference on Robot Learning (CoRL)*, 2021.
- [29] A. Zeng, A. Wong, S. Welker, K. Choromanski, F. Tombari, A. Purohit, M. Ryoo, V. Sindhwani, J. Lee, V. Vanhoucke, et al. Socratic models: Composing zero-shot multimodal reasoning with language. *arXiv preprint arXiv:2204.00598*, 2022.
- [30] Shuran Song and J. Xiao. Sliding Shapes for 3D object detection in depth images. 2014.
- [31] Shuran Song and J. Xiao. Deep sliding shapes for amodal 3D object detection in rgb-d images. 2016.
- [32] J. Hou, A. Dai, and M. Nießner. 3d-sis: 3d semantic instance segmentation of rgb-d scans. In *Proc. Computer Vision and Pattern Recognition (CVPR), IEEE*, 2019.
- [33] S. Song, F. Yu, A. Zeng, A. X. Chang, M. Savva, and T. Funkhouser. Semantic scene completion from a single depth image. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1746–1754, 2017.
- [34] A. Dai, D. Ritchie, M. Bokeloh, S. Reed, J. Sturm, and M. Nießner. Scancomplete: Large-scale scene completion and semantic segmentation for 3d scans. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4578–4587, 2018.

- [35] A. Avetisyan, M. Dahnert, A. Dai, M. Savva, A. X. Chang, and M. Nießner. Scan2cad: Learning cad model alignment in rgb-d scans. In *Proceedings of the IEEE/CVF Conference on computer vision and pattern recognition*, pages 2614–2623, 2019.
- [36] J. Roh, K. Desingh, A. Farhadi, and D. Fox. Languagerefer: Spatial-language model for 3d visual grounding. In *Conference on Robot Learning*, pages 1046–1056. PMLR, 2022.
- [37] L. Barbara and R. Jackendoff. “what” and “where” in spatial language and spatial cognition. *Behavioral and brain sciences*, 16:217–238, 1993.
- [38] W. G. Hayward and M. J. Tarr. Spatial language and spatial representation. *Cognition*, 55(1):39–84, 1995.
- [39] S. Tulsiani, H. Su, L. J. Guibas, A. A. Efros, and J. Malik. Learning shape abstractions by assembling volumetric primitives. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2635–2643, 2017.
- [40] I. Binford. Visual perception by computer. In *IEEE Conference of Systems and Control*, 1971.
- [41] A. Gupta, A. A. Efros, and M. Hebert. Blocks world revisited: Image understanding using qualitative geometry and mechanics. In *European Conference on Computer Vision*, pages 482–496. Springer, 2010.
- [42] Ö. Çiçek, A. Abdulkadir, S. S. Lienkamp, T. Brox, and O. Ronneberger. 3d u-net: Learning dense volumetric segmentation from sparse annotation. *CoRR*, abs/1606.06650, 2016. URL <http://arxiv.org/abs/1606.06650>.
- [43] H. Chefer, I. Schwartz, and L. Wolf. Optimizing relevance maps of vision transformers improves robustness. In *Thirty-Sixth Conference on Neural Information Processing Systems*, 2022. URL <https://openreview.net/forum?id=upuYKQiyxa..>
- [44] H. Chefer, S. Gur, and L. Wolf. Transformer interpretability beyond attention visualization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 782–791, 2021.
- [45] Y. Liu, H. Li, Y. Guo, C. Kong, J. Li, and S. Wang. Rethinking attention-model explainability through faithfulness violation test. *arXiv preprint arXiv:2201.12114*, 2022.
- [46] S. Shen, L. H. Li, H. Tan, M. Bansal, A. Rohrbach, K.-W. Chang, Z. Yao, and K. Keutzer. How much can clip benefit vision-and-language tasks? *arXiv preprint arXiv:2107.06383*, 2021.
- [47] N. Liu, S. Li, Y. Du, J. Tenenbaum, and A. Torralba. Learning to compose visual relations. *Advances in Neural Information Processing Systems*, 34, 2021.
- [48] M. Yuksekgonul, F. Bianchi, P. Kalluri, D. Jurafsky, and J. Zou. When and why vision-language models behave like bags-of-words, and what to do about it?, 2022.
- [49] N. Silberman, D. Hoiem, P. Kohli, and R. Fergus. Indoor segmentation and support inference from rgb-d images. In *European conference on computer vision*, pages 746–760. Springer, 2012.
- [50] R. Guo, C. Zou, and D. Hoiem. Predicting complete 3d models of indoor scenes. *arXiv preprint arXiv:1504.02437*, 2015.
- [51] Y. Cai, X. Chen, C. Zhang, K.-Y. Lin, X. Wang, and H. Li. Semantic scene completion via integrating instances and scene in-the-loop. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 324–333, 2021.
- [52] P. Zhang, W. Liu, Y. Lei, H. Lu, and X. Yang. Cascaded context pyramid for full-resolution 3d semantic scene completion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7801–7810, 2019.
- [53] M. Wortsman, G. Ilharco, S. Y. Gadre, R. Roelofs, R. Gontijo-Lopes, A. S. Morcos, H. Namkoong, A. Farhadi, Y. Carmon, S. Kornblith, and L. Schmidt. Model soups: averaging weights of multiple fine-tuned models improves accuracy without increasing inference time, 2022. URL <https://arxiv.org/abs/2203.05482>.