

3D Vision and Language Pretraining with Large-Scale Synthetic Data

Dejie Yang¹, Zhu Xu¹, Wentao Mo¹, Qingchao Chen^{2,3}, Siyuan Huang⁴, Yang Liu^{1,3*}

¹ Wangxuan Institute of Computer Technology, Peking University

² National Institute of Health Data Science, Peking University

³ National Key Laboratory of General Artificial Intelligence, Peking University

⁴ State Key Laboratory of General Artificial Intelligence, BIGAI

{yjdj, xuzhu}@stu.pku.edu.cn, huangsiyuan@ucla.edu, {mowt, qingchao.chen, yangliu}@pku.edu.cn

Abstract

3D Vision-Language Pre-training (3D-VLP) aims to provide a pre-train model which can bridge 3D scenes with natural language, which is an important technique for embodied intelligence. However, current 3D-VLP datasets are hindered by limited scene-level diversity and insufficient fine-grained annotations (only 1.2K scenes and 280K textual annotations in ScanScribe), primarily due to the labor-intensive of collecting and annotating 3D scenes. To overcome these obstacles, we construct SynVL3D, a comprehensive synthetic scene-text corpus with 10K indoor scenes and 1M descriptions at object, view, and room levels, which has the advantages of diverse scene data, rich textual descriptions, multi-grained 3D-text associations, and low collection cost. Utilizing the rich annotations in SynVL3D, we pre-train a simple and unified Transformer for aligning 3D and language with multi-grained pretraining tasks. Moreover, we propose a synthetic-to-real domain adaptation in downstream task fine-tuning process to address the domain shift. Through extensive experiments, we verify the effectiveness of our model design by achieving state-of-the-art performance on downstream tasks including visual grounding, dense captioning, and question answering. Codes are available at: <https://github.com/idejie/3DSyn>.

1 Introduction

Bridging 3D scenes with natural language represents a pivotal advancement in the pursuit of embodied artificial intelligence [Huang *et al.*, 2023] including 3D vision-language (3D-VL) tasks such as 3D visual grounding [Chen *et al.*, 2020; Achlioptas *et al.*, 2020], dense captioning [Chen *et al.*, 2021], and question answering [Azuma *et al.*, 2022; Mo and Liu, 2024], etc.

Drawing inspiration from the impact of 2D vision-text pre-training models [Radford *et al.*, 2021; Alayrac *et al.*, 2022], recent researches [Zhu *et al.*, 2023; Jin *et al.*, 2023] demonstrate that 3D-VL pre-training can enhance the performance

*Corresponding Author

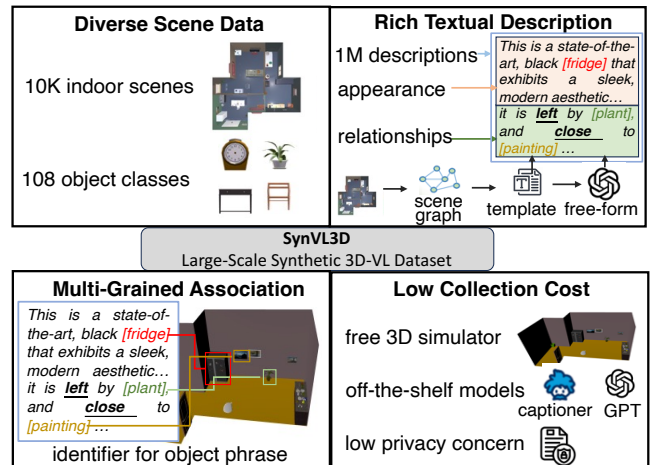


Figure 1: Advantages of our proposed dataset SynVL3D.

in subsequent tasks. However, the performance of these methods is limited by the size and diversity of 3D-text data, as 3D data collection largely depends on the scanning device, introducing inherent complexities and increased costs in contrast to the simpler process of gathering 2D images. There are limitations to the existing 3D-VL datasets: (1) *limited diversity of 3D vision*, the largest-scale 3D-VL dataset ScanScribe [Zhu *et al.*, 2023] combines multiple existing 3D datasets, but still only contains 1.2K indoor scenes, which causes limited vision diversity for pre-training; (2) *insufficient fine-grained 3D-text associations*, ScanScribe provides 278K descriptions of objects where only the sentence-level target objects is annotated and ignores the other fine-grained vision-language associations between the spatial region and other objects mentioned in the sentence. (3) *labor-intensive data collection procedure*, the 3D data collection of a real scene needs 30 minutes with advanced scanners [Yeshwanth *et al.*, 2023] and annotating a scene with the help of the automatic segmentation still requires 1.93 minutes, makes large-scale real-world 3D scene infeasible to collect.

To address the limitation of existing 3D-VL datasets, we propose the dataset SynVL3D, which contains the free-collected and easy-to-annotate synthetic 3D data generated by a 3D scene simulator [Deitke *et al.*, 2022], instead of using re-

alistic data. As shown in Figure 1, our dataset SynVL3D has four advantages: (1) **Diverse scene data.** Benefit from the capability of the 3D simulator [Deitke *et al.*, 2022], we collected 10,000 indoor 3D scenes with 108 object categories. And SynVL3D provides semantic (category), visual (high-quality mesh), and fine-grained visual annotation (error-free position, orientation and segmentation mask) of each object. To acquire fine-grained understanding for scenes, we formed a scene graph for each house (scene) with the relationships between objects based on their spatial position, size, and shape. (2) **Rich textual description.** We utilize templates, off-the-shelf image captioners [Dai *et al.*, 2023], and GPT3 [Brown *et al.*, 2020] to generate appearance and relationship descriptions of objects, and obtain more than 1M template-based or free-form descriptions. (3) **Multi-grained association.** Based on the above textual descriptions and scene graph, we retrain a distinct identifier for every noun phrase, creating a connection between each identifier and a spatial region. This establishes a robust phrase-region association, enhancing the grounded interpretation and facilitating a more meaningful evaluation. (4) **Low collection cost.** Based on simulators, we can obtain free 3D scene data, use off-the-shelf models to automatically describe scenes and obtain fine-grained annotations with few manual annotations. With virtual scenes, we also avoid privacy issues compared to real scenes.

Based on the diverse and fine-grained annotations of large scale of 3D scene provided by SynVL3D, we can construct more **fine-grained pre-training tasks.** The tasks include relationship prediction between objects, multi-level and view-aggregated region-word alignment, beyond scene-text matching and masked language/object modeling that previous pre-training methods [Zhu *et al.*, 2023; Jin *et al.*, 2023] use. The multi-grained 3D VLP tasks enable the model to capture rich 3D vision-language knowledge, which benefits downstream vision-language tasks. However, pre-training based on synthetic 3D data suffers from the domain shift between the synthetic data and the realistic data used in downstream tasks. Thus, we introduce **synthetic-to-real adaptation** with domain discriminator in the fine-tuning process of downstream tasks to reduce the representation distribution difference between the synthetic and realistic data of our pre-trained model. The efficacy of our model is evaluated across 3D-VL tasks, including visual grounding (e.g., ScanRefer [Chen *et al.*, 2020], Nr3D/Sr3D [Achlioptas *et al.*, 2020]), dense captioning (e.g., Scan2Cap [Chen *et al.*, 2021]), question answering (e.g., ScanQA [Azuma *et al.*, 2022]).

Our main contributions can be summarized as follows: (1) We construct a large-scale synthetic dataset SynVL3D, featuring diverse 3D scenes, textual descriptions, and fine-grained 3D region-text associations at a low manual cost, which is a significant improvement in terms of data diversity and scale compared to prior datasets. (2) We introduce a pre-training framework featuring new auxiliary tasks such as pairwise object relationship prediction and multi-level 3D region-phrase alignment to enhance comprehensive understanding. To address the domain shift between synthetic and real 3D data, we introduce a domain adaptation technique, improving the fine-tuning process for downstream tasks. (3) We demonstrate that with the data scale-up and model design, our model

SynFormer3D achieves state-of-the-art performance on a variety of 3D-VL tasks, including visual grounding, dense captioning, and question answering.

2 Related Work

2.1 3D Vision-Language Dataset

There has been a growing interest in 3D-VL datasets. [Chen *et al.*, 2020] and [Achlioptas *et al.*, 2020] introduced the ScanRefer and ReferIt3D datasets, which serve as benchmarks for anchoring sentence-level natural language to ground the 3D object. [Abdelreheem *et al.*, 2024] and [Yuan *et al.*, 2022a] extend to provide phrase-level grounding annotations. Scan2Cap [Chen *et al.*, 2021] provides the textual descriptions about the scene in ScanRefer for captioning. ScanQA [Azuma *et al.*, 2022] developed a 3D question-answering dataset, which annotates the objects mentioned within the questions and answers. ScanScribe [Zhu *et al.*, 2023] endeavors to build a large-scale dataset by combining the aforementioned datasets for pre-training purposes. However, these 3D-VL datasets are constrained by a lack of scene-level diversity, fine-grained annotations (the largest existing dataset ScanScribe with only 1.2K scenes and 280K textual descriptions, and the laborious process involved in collecting and annotating 3D scenes. To overcome these obstacles, we constructed the dataset SynVL3D, which features freely collected and easily annotated 3D data generated by a 3D scene simulator [Deitke *et al.*, 2022], as opposed to utilizing realistic data. It encompasses diverse 3D scene data (10K scenes), provides 1M rich textual descriptions with low cost. Moreover, compared to the sentence-region association in ScanScribe, we provide the phrase-region finegrained association, which means that most noun phrase in description has an association with the spatial region in the 3D scene.

2.2 Large-scale 3D-VL Pre-training

In recent years, large-scale pre-training has emerged as a fundamental paradigm in natural language processing (NLP) [Radford *et al.*, 2018; Radford *et al.*, 2019; Brown *et al.*, 2020] and 2D vision-and-language (2D-VL) domains [Lu *et al.*, 2019; Radford *et al.*, 2021]. While the pre-training technique has become indispensable in these domains, its application in the 3D-VL field remains relatively uncharted. 3D-VISTA [Zhu *et al.*, 2023] seeks to pioneer a pre-training model for 3D-VL tasks, employing strategies such as masked language modeling, masked object modeling, and scene-sentence matching. Concurrently, 3D-VLP [Jin *et al.*, 2023] introduces a context-aware spatial-semantic alignment to forge connections between point clouds and the textual descriptions of ground objects. However, the efficacy of these pre-training tasks is often compromised by the lack of detailed annotations in existing datasets. In this paper, we propose three fine-grained pre-training tasks, including object relations prediction, multi-level region word alignment and view-aggregated region word alignment, to enhance the model’s perception capability of object relationships, scenes at different levels and from various views.

Dataset	Type	Text	Scene	Annotation
Nr3D [Achlioptas <i>et al.</i> , 2020]	Realistic	31K	707	Sentence-Region
ScanRefer [Chen <i>et al.</i> , 2020]	Realistic	37K	800	Sentence-Region
ScanQA [Azuma <i>et al.</i> , 2022]	Realistic	41K	800	Sentence-Region
ScanEnts3D [Abdelreheem <i>et al.</i> , 2024]	Realistic	84K	705	Phrase-Region
PhraseRefer [Yuan <i>et al.</i> , 2022a]	Realistic	88K	707	Phrase-Region
3R-Scan [Wald <i>et al.</i> , 2019]	Realistic	90K	478	Sentence-Region
Sr3D [Achlioptas <i>et al.</i> , 2020]	Realistic	91K	707	Sentence-Region
ScanScribe [Zhu <i>et al.</i> , 2023]	Realistic	278K	1.2K	Sentence-Region
Syn3D(Ours)	Synthetic	1M	10K	Phrase-Region

Table 1: The comparison between SynVL3D and other 3D-VL datasets. ‘‘Type’’ stands for the source of 3D scenes. ‘‘Text’’ denotes the textual description size.

2.3 Simulation-to-Real Domain Adaptation

Simulation-to-real domain adaptation refers to the process of transferring knowledge acquired in a simulated environment to the real world, and is widely applied in 2D vision [Pei *et al.*, 2024; Zhao *et al.*, 2023; Pei *et al.*, 2023; Xu *et al.*, 2022; Chen and Liu, 2020; Chen *et al.*, 2018] and 3D fields like autonomous vehicles [Zhao *et al.*, 2021b], indoor scenes layout [Ding *et al.*, 2022], and robotics [DeBortoli *et al.*, 2021]. Existing methods primarily aim to design a simulation-to-real domain adaptation framework for uni-modal applications, such as the 3D semantic segmentation [Zhao *et al.*, 2021b; Ding *et al.*, 2022], point cloud object detection [DeBortoli *et al.*, 2021] and point cloud registration [Chen *et al.*, 2023]. However, there are few works exploring the simulation-to-real domain adaptation in 3D cross-modality applications, where the domain shift exists in both 3D vision and language. To address this dual and joint domain shift, we propose a task-specific fine-tuning strategy with domain adaptation in 3D vision, language, and joint 3D-language domains.

3 Dataset Generation: SynVL3D

In this paper, we introduce SynVL3D, the first large-scale dataset of 3D synthetic scene-text pairs designed for 3D-VL pre-training. IAs indicated in Table 1, SynVL3D not only surpasses existing 3D-VL datasets in size but also offers a richer variety of 3D scenes, text, and fine-grained 3D-text annotations. Our dataset SynVL3D showcases key advantages in four aspects:

Diverse 3D Scene Data. We collected mesh data for 10,000 distinct scenes from an off-the-shelf 3D scene simulator [Deitke *et al.*, 2022]. It encompasses 108 object categories, totaling over 1 million distinct instances. Besides the variety surpasses the existing datasets as shown in Table 1, we provide following annotation: (1) For each *Single Instance*, we automatically acquire its semantic category and precise bounding boxes and segmentation masks directly from the object and scene information provided by the simulator [Deitke *et al.*, 2022] at no extra cost. (2) For *Multiple Instance* relationship, we construct a scene graph for each scene based on densely annotated objects and their attributes, where each node represents an object instance, and edges represent the relationships between an instance and its neighbors. Inspired by [Wald *et al.*, 2020], the scene graph encompasses a rich array of relationships, which can be classified into three types: a) support relationships (e.g., standing, lying); b) spa-

tial relationships (e.g., next to, in front of) that consider the relative positions of neighboring instances; and c) comparative relationships (e.g., bigger than, same shape as) that are based on the size and shape of an instance.

Rich Textual Description. We have generated more than 1 million textual descriptions, both templated and free-form, encapsulating detailed appearances and relationships between objects. (1) *Template-based*: For object appearances, we utilize an off-the-shelf image captioner [Dai *et al.*, 2023] to generate captions for the front view of each object. We manually reviewed these descriptions and make corrections where necessary. For relationship triplets of the form $\langle object1, relation, object2 \rangle$, we craft descriptions using the template ‘‘the *object1* is *relation* to *object2*’’. The description of an instance is structured as a paragraph by merging its appearance description with the relationship descriptions of its neighbors. (2) *Free-form*: We leverage GPT-3 [Brown *et al.*, 2020] to rephrase the templated description to enhance the naturalness. These descriptions not only enrich the dataset with diverse visual perspectives but also enhance the semantic understanding of spatial relations and object attributes.

Phrase-Region Association. The grounding annotations in previous datasets [Achlioptas *et al.*, 2020; Chen *et al.*, 2020; Zhu *et al.*, 2023] mainly consist sentence-region association for the single referred target object. In this paper, for 10K scene, we aim to curate grounding annotations of most objects in each description by providing phrase-region explicit association with an identifier of a phrase that refers to an object in a 3D scene. In practice, GPT-3 is tasked with rephrasing the descriptions while maintaining a critical requirement: the preservation of the identifier for each object, formatted as *category(id of instance)*, throughout the rephrasing process. This ensures that the region-word association remains intact, maintaining the accuracy of object referencing in the dataset.

Low Collection Cost. Unlike other datasets in Table 1 that are manually collected and annotated from real-world scenes, our SynVL3D leverages freely synthetic 3D scene data. We utilize off-the-shelf models for automatic scene description and require minimal manual annotation cost. By using virtual scenes, our approach also circumvents privacy issues that can arise from scanning real-world environments.

4 Methodology

As shown in Figure 2, our SynFormer3D first uses multi-modal encoder to extract and fuse 3D and text features. The pipeline mainly consists two steps: synthetic pre-training (Section 4.2) and task-specific fine-tuning (Section 4.3).

4.1 Multi-Modal Encoder

Following [Zhu *et al.*, 2023], firstly we employ text encoder to encode words to text tokens $\{\mathbf{w}_{CLS}, \mathbf{w}_{1:L}\}$, where \mathbf{w}_{CLS} denotes the $[cls]$ token and L denotes the length of a description. The 3D-object encoder is to encode M object proposals and extract features from the input point cloud, represented by $\mathbf{o}_{1:M}$. To achieve multi-modal fusion, we concatenate the text tokens and the 3D object tokens $\{\mathbf{w}_{CLS}, \mathbf{w}_{1:L}, \mathbf{o}_{1:M}\}$ and

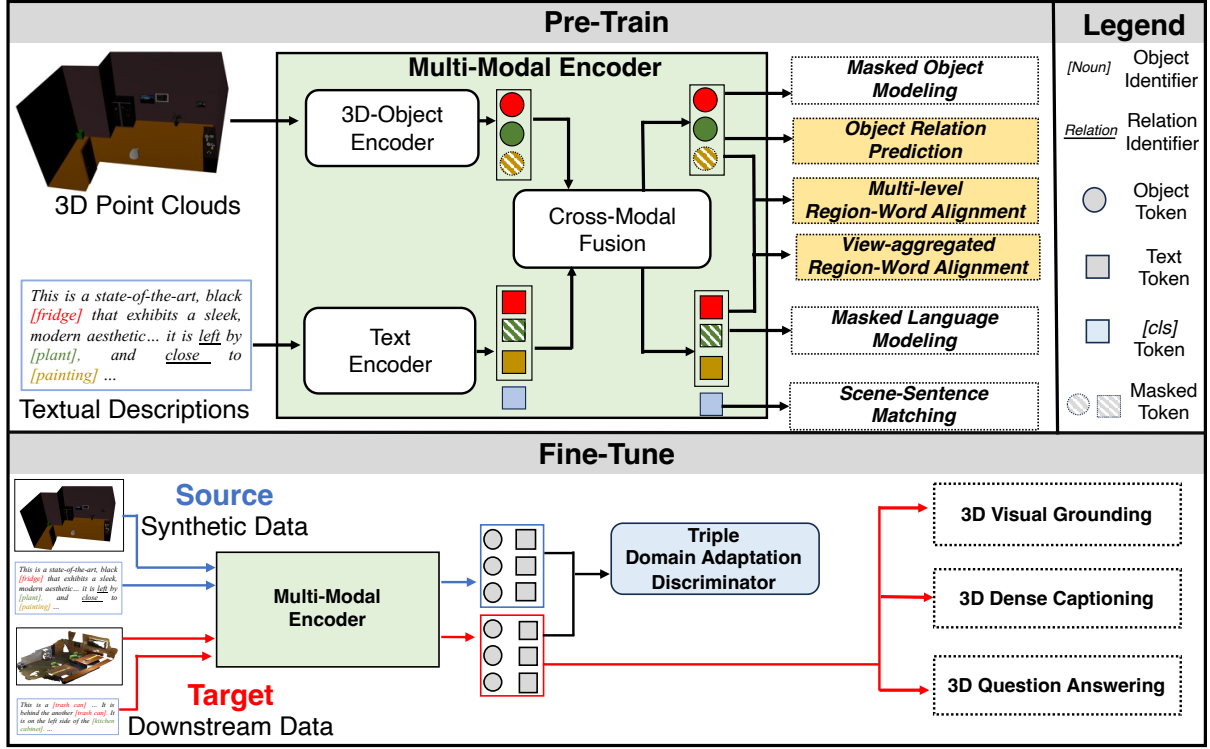


Figure 2: The model architecture of our SynFormer3D. The multi-modal encoder includes a 3D-object encoder, text encoder, and cross-modal fusion modules. Compared to previous pre-trained models, our SynFormer3D introduces more fine-grained auxiliary pre-training tasks, which include Object Relation Prediction, Multi-level and View-aggregated Region-Word Alignment.

forward them to a Transformer. The outputs of the cross-modal fusion module are denoted as $\{\mathbf{W}_{CLS}, \mathbf{W}_{1:L}, \mathbf{O}_{1:M}\}$, corresponding to $[cls]$ token, word tokens and object features.

4.2 Synthetic Pre-Training

To align 3D scene and text in self-supervised manner, pre-training tasks Masked Language Modeling (MLM), Masked Object Modeling (MOM), and Scene-Sentence Matching (SSM) are adopted. Further, to achieve better alignment, we propose three fine-grained pre-training tasks based on the fine-grained annotation in SynVL3D: Object Relationship Prediction (ORP), Multi-level Region-Word Alignment (MRWA) and View-aggregated Region-Word Alignment (VRWA), details are as follows:

ORP. The previous methods [Zhu *et al.*, 2023; Jin *et al.*, 2023] neglect the importance of relationship understanding between object pairs, resulting in limitations for downstream tasks. And we introduce object relationship prediction task, which predicting pairwise relationship with prediction head $r(\cdot, \cdot)$, and the relationship annotations in SynVL3D are utilized to form label.

$$\mathcal{L}_{ORP} = -\frac{1}{M^2} \sum_{i=1}^M \sum_{j=1}^M R_{i,j} \cdot \log(r(\mathbf{O}_i, \mathbf{O}_j)), \quad (1)$$

where i and j denote the indices of object features within the set of object proposal features \mathbf{O} , M is the total number

of object proposal and $R_{i,j}$ is the label of relations between object i and j .

MRWA. 3D tasks and operations involve multi-level targets with varying granularity, ranging from object, room, to entire scene. To achieve multi-level alignment, we align regions and words at object, room and scene levels respectively:

$$\mathcal{L}(\mathbf{x}, \mathbf{z}) = -\frac{1}{L \cdot M} \sum_{i=1}^L \sum_{j=1}^M y_{i,j} \cdot \log(p(\mathbf{x}_i, \mathbf{z}_j)), \quad (2)$$

$$\mathcal{L}_{MRWA} = \sum_l \mathcal{L}(\mathbf{W}^l, \mathbf{O}^l), \quad (3)$$

where L and M are the length of text(word) token sequence and the number of object features respectively. $y_{i,j}$ means the binary label of whether the i -th word \mathbf{x}_i and j -th region(object) \mathbf{z}_j is matching. $l \in \{object, room, scene\}$ indicate the levels of a single object, room, and the entire scene, respectively. \mathbf{W}^l represents the word feature sequence from descriptions at different levels, and \mathbf{O}^l is the object feature set derived from the point cloud of the corresponding level. $p(\cdot, \cdot)$ is the region-word matching prediction head. The set of objects involved in different levels of textual description is different. At the object level, only objects near the description are mentioned, while room and house levels provide a larger range of objects. Therefore, region-word alignment at different levels considers the relationships between objects

from the object set at different levels (the self-attention is conducted on different set of objects/words).

VRWA. The perspective changes when executing operations or object searching in 3D scenes, which makes view-aggregated object-text alignment critical. Following [Guo *et al.*, 2023], we rotate the point cloud of a 3D scene into V to different views using the annotated camera parameters from SynVL3D: $\mathbf{P}^v = R(\theta^v)\mathbf{P}$, where $R(\theta)$ is the rotation matrix [Guo *et al.*, 2023], \mathbf{P} is the scene point cloud and θ^v is the rotation angle adopted for v -th view: $\theta^v = \frac{2\pi}{V}(v-1)$, $v \in \{1, \dots, V\}$ for a total of V views. $R(\theta^v) \times \mathbf{P}$ represents rotating the point cloud of whole scene clockwise by θ^v degrees along the Z-axis. Then we encode the v -th view features \mathbf{O}^v , and apply multi-modal encoder and aggregate multi-view information for the i -th object: $\mathbf{G}_i = \frac{1}{V} \sum_{v=1}^V \mathbf{O}_i^v$. To achieve precise alignment between regions and words (as nouns in descriptions correspond to object identifiers in Figure 2), we perform region-phrase matching besides aligning scene and text descriptions:

$$\mathcal{L}_{VRWA} = -\frac{1}{L \cdot M} \sum_{i=1}^L \sum_{j=1}^M z_{i,j} \log(s(\mathbf{W}_i, \mathbf{G}_j)). \quad (4)$$

where $z_{i,j}$ is the matching label of i -th word \mathbf{W}_i and j -th region(object) \mathbf{G}_j . $s(\cdot, \cdot)$ is the region-word matching prediction head.

Final Pre-training Loss. The final pre-training objective combines the losses from the proxy tasks mentioned earlier:

$$\mathcal{L}_{\text{pre-train}} = \alpha(\mathcal{L}_{MLM} + \mathcal{L}_{MOM} + \mathcal{L}_{SSM}) + (1 - \alpha)(\mathcal{L}_{ORP} + \mathcal{L}_{MRWA} + \mathcal{L}_{VRWA}), \quad (5)$$

where α is a hyperparameter used to balance the basic auxiliary tasks(the losses \mathcal{L}_{MLM} , \mathcal{L}_{MOM} , \mathcal{L}_{SSM} is following [Zhu *et al.*, 2023]) with our proposed fine-grained pre-training tasks. Note that the pre-training approach we propose is self-supervised and downstream task-agnostic.

4.3 Task-Specific Fine-Tuning

By incorporating lightweight task-specific decoders, the pre-trained model can be easily adapted to a variety of 3D-VL tasks. Due to the possible domain shift in our synthetic dataset and downstream task data (real scene scan data), we propose a task-specific fine-tuning strategy for synthetic-to-real domain adaptation to be better generalized to downstream tasks. In particular, we fine-tune the model on the following tasks: 3D Visual Grounding, 3D Dense Captioning, 3D Question Answering. We equip each task a specific lightweight decoder and corresponding fine-tuning losses.

Synthetic-to-Real Adaptation. Domain shift from multiple perspectives between synthetic data for pre-training and realistic data used in downstream tasks may hinder the performance on various 3D VL tasks. Vision differences can be observed between the synthesized data the scanned real data as shown in Figure 2. Besides, the descriptions in our dataset are generated by templates or GPT, which also has a gap with artificially annotated natural language descriptions in downstream datasets. Moreover, downstream tasks also involve the use of joint vision language, such as 3D question

answering, where this type of domain shift requires a joint adaptation. Therefore, we introduce synthetic-to-real adaptation with a *triple domain discriminator* (shown in Figure2) in the fine-tuning process of downstream tasks to reduce the representation distribution difference between the synthetic and realistic data of our pre-trained model. The triple domain discriminator consists of three separate discriminator for vision, language and vision-language joint domains, addressing unimodal and cross-modal domain shifts.

A vision domain discriminator D_v is placed after multi-modal encoder E to predict the domain label of the object features \mathbf{O} :

$$\mathcal{L}_{\text{vision}} = \max_E \min_{D_v} \mathcal{L}_{BCE}(\mathbf{O}) \quad (6)$$

where \mathcal{L}_{BCE} denotes the binary cross entropy loss. Gradient Reverse Layers(GRL) is adopted for min-max optimization. Similarly, a language domain discriminator D_l is adopted to predict the domain label of sentence tokens $\hat{\mathbf{W}}$:

$$\mathcal{L}_{\text{lang}} = \max_E \min_{D_l} \mathcal{L}_{BCE}(\hat{\mathbf{W}}) \quad (7)$$

where the sentence token $\hat{\mathbf{W}}$ is the mean of the text token $\mathbf{W}_{1:L}$ in a sentence. In order to reduce the gap of 3D-text alignment in different domains, we also introduce a joint discriminator D_{joint} . And the adaptation loss for shared-domain is:

$$\mathcal{L}_{\text{joint}} = \max_E \min_{D_{\text{joint}}} \mathcal{L}_{BCE}([\hat{\mathbf{W}}_i; \mathbf{O}_i]) \quad (8)$$

where $[\hat{\mathbf{W}}_i; \mathbf{O}_i]$ denote i -th data sample (setence and corresponding 3D object) which is the concatenation of the $\hat{\mathbf{W}}$ and \mathbf{O} .

The final synthetic-to-real adaptation loss can be combined as: $\mathcal{L}_{\text{align}} = \mathcal{L}_{\text{vision}} + \mathcal{L}_{\text{lang}} + \mathcal{L}_{\text{joint}}$.

Final Fine-tune Loss. The final fine-tuning objective combines the losses from the specific downstream task and synthetic-to-real domain adaptation:

$$L_f = \beta \mathcal{L}_{\text{task}} + (1 - \beta) \mathcal{L}_{\text{align}}, \quad (9)$$

where $\mathcal{L}_{\text{task}}$ is the specific loss of the downstream task mentioned above and where β is a hyperparameter used to balance the downstream task loss and domain-adaptation loss.

5 Experiments

5.1 Implementation Details

The pre-training runs for 100 epochs with a batch size of 64 with NVIDIA A100 GPUs. We set the balance hyper-parameters α, β as 0.5 and 0.8. We use the AdamW [Loshchilov and Hutter, 2019] optimizer and learning rate is set to $1e^{-4}$.

5.2 Downstream Task

3D Visual Grounding. In Table 2, we evaluate our model on Nr3D, Sr3D [Achlioptas *et al.*, 2020] and ScanRefer [Chen *et al.*, 2020]. For Nr3D/Sr3D, we report the results as the grounding accuracy. For ScanRefer, we follow [Chen *et al.*, 2020] to use detector-generated object proposals and report the results as $\text{Acc@IoU} > k$. As shown

Method	Nr3D					Sr3D					ScanRefer	
	Overall	Easy	Hard	View Dep	View Indep	Overall	Easy	Hard	View Dep	View Indep	acc 0.25	acc 0.5
3DVG-Trans [Zhao <i>et al.</i> , 2021a]	40.8	48.5	34.8	34.8	43.7	51.4	54.2	44.9	44.6	51.7	47.6	34.7
TransRefer3D [He <i>et al.</i> , 2021]	48.0	56.7	39.6	42.5	50.7	57.4	60.5	50.2	49.9	57.7	-	-
LAR [Bakr <i>et al.</i> , 2022]	48.9	58.4	42.3	47.4	52.1	59.4	63.0	51.2	50.0	59.1	-	-
SAT [Yang <i>et al.</i> , 2021]	56.5	64.9	48.4	54.4	57.6	57.9	61.2	50.0	49.2	58.3	44.5	30.1
3D-SPS [Luo <i>et al.</i> , 2022]	51.5	58.1	45.1	48.0	53.2	62.6	56.2	65.4	49.2	63.2	48.8	37.0
MVT [Huang <i>et al.</i> , 2022]	59.5	67.4	52.7	59.1	60.3	64.5	66.9	58.8	58.4	64.7	40.8	33.3
ViL3DRel [Chen <i>et al.</i> , 2022]	<u>64.4</u>	70.2	57.4	<u>62.0</u>	64.5	72.8	74.9	67.9	<u>63.8</u>	73.2	47.9	37.7
ScanRefer [Chen <i>et al.</i> , 2020]	-	-	-	-	-	-	-	-	-	-	41.2	27.4
3DJCG [Cai <i>et al.</i> , 2022]	-	-	-	-	-	-	-	-	-	-	49.6	37.3
3D-VLP[Jin <i>et al.</i> , 2023]	-	-	-	-	-	-	-	-	-	-	<u>51.4</u>	39.5
3D-VISTA[Zhu <i>et al.</i> , 2023]	64.2	<u>72.1</u>	<u>56.7</u>	61.5	<u>65.1</u>	<u>76.4</u>	<u>78.8</u>	<u>71.3</u>	58.9	<u>77.3</u>	<u>50.6</u>	<u>45.8</u>
Ours	65.5	73.2	<u>56.3</u>	63.2	<u>66.1</u>	77.9	79.2	73.3	64.1	79.2	52.3	46.2

 Table 2: Grounding accuracy (%) on Nr3D, Sr3D and ScanRefer. The best and second-best results are in **bold** and underlined.

Method	C	@0.25				@0.5			
		B-4	M	R	C	B-4	M	R	
No Pre-train									
Scan2Cap [Chen <i>et al.</i> , 2021]	53.7	34.3	26.1	55.0	35.2	22.4	21.4	43.5	
X-Trans2Cap[Yuan <i>et al.</i> , 2022b]	58.8	34.2	25.8	54.1	41.5	23.8	21.9	45.0	
MORE[Jiao <i>et al.</i> , 2022]	58.9	35.4	26.4	55.4	39.0	23.0	21.7	44.3	
3DJCG [Cai <i>et al.</i> , 2022]	60.9	<u>39.7</u>	27.5	59.0	47.7	31.5	24.3	51.8	
Pre-train									
3D-VLP [Jin <i>et al.</i> , 2023]	64.1	39.8	27.7	<u>58.8</u>	50.0	31.9	24.5	51.5	
3D-VISTA[Zhu <i>et al.</i> , 2023]	<u>71.0</u>	36.5	<u>28.4</u>	57.6	66.9	<u>34.0</u>	<u>27.1</u>	<u>54.3</u>	
Ours	72.1	37.3	28.9	58.2	<u>64.2</u>	34.6	28.0	54.8	

 Table 3: Dense Captioning results on Scan2Cap dataset. ‘‘C’’ stands for ‘‘CIDEr’’, ‘‘B-4’’ for ‘‘BLEU-4’’, ‘‘M’’ for ‘‘METEOR’’, and ‘‘R’’ for ‘‘ROUGE’’, respectively. ‘‘@0.25’’ and ‘‘@0.5’’ represent the 3D IoU between the predicted and annotated box. The best and second-best results are in **bold** and underlined.

Method	EM@1	EM@10	BLEU-4	ROUGE	METEOR	CIDEr
No Pre-train						
ScanQA [Azuma <i>et al.</i> , 2022]	23.5 / 20.9	56.5 / <u>54.1</u>	12.0 / 10.8	34.3 / 31.1	13.6 / 12.6	67.3 / 60.2
CLIP-Guided [Parelli <i>et al.</i> , 2023]	23.9 / 21.4	- / -	14.6 / 11.7	35.2 / 32.4	13.9 / <u>13.3</u>	69.5 / 62.8
Multi-CLIP [Delitzas <i>et al.</i> , 2023]	24.0 / 21.6	- / -	12.7 / 12.9	35.5 / 32.6	14.0 / 13.4	68.7 / 63.2
Pre-train						
3D-VLP[Jin <i>et al.</i> , 2023]	25.2 / 20.4	55.2 / 51.5	10.5 / 8.7	35.5 / 29.6	13.8 / 11.6	68.6 / 55.7
3D-VISTA[Zhu <i>et al.</i> , 2023]	<u>27.0</u> / 23.0	<u>57.9</u> / 53.5	16.0 / 11.9	38.6 / 32.8	15.2 / 12.9	76.6 / 62.6
Ours	27.6 / 24.1	58.3 / 54.5	16.3 / <u>12.3</u>	39.2 / 33.3	<u>14.9</u> / 13.1	<u>76.2</u> / 62.7

 Table 4: Answer accuracy on ScanQA using object proposals from Mask3D. Each entry denotes ‘‘test w/ object’’ / ‘‘test w/o object’’. The best and second-best results are in **bold** and underlined.

in Table 2, SynFormer3D obtains an overall accuracy of 67.5% and 78.6% on Nr3D and Sr3D, which has gains of +1.1% and +1.5% compared to the previous leading methods. It indicates that SynFormer3D is trained on our large-scale SynVL3D with diverse 3D vision and text can understand complex 3D scenes better. Notably, our approach can be adaptive to different views better(+1.7% and +5.2% on ‘‘View-Dep’’ compared with 3D-VISTA), benefiting from the well-explored view-aggregated region-word alignment during the pre-training stage. Our SynFormer3D obtains an overall accuracy of 52.3% @0.25 and 46.2% @0.5, which outperforms all other methods on ScanRefer. Compared to the pre-training methods [Jin *et al.*, 2023; Zhu *et al.*, 2023], we introduce more fine-grained pre-training tasks benefiting from

the scale-up data and rich and accurate annotations of our SynVL3D.

3D Dense Captioning. In Table 3, we evaluate our model on the Scan2cap. Our method outperforms the second best scores on CIDEr@0.25 (+1.1%), METEOR@0.25 (+0.5%), BLEU-4@0.5 (+2.6%), METEOR@0.5 (+0.9%) and ROUGE@0.5 (+0.5%). While our method achieves state-of-the-art in most metrics, it is only comparable with previous methods on BLEU-4@0.25, ROUGE@0.25 and CIDEr@0.5. The cause could be attributed to the fact that the textual descriptions utilized in pretraining are automatically generated with minimal manual annotations. This may result in gaps and biases about fluency and free-form expression.

ORP	VRWA	MRWA	VG	Cap	QA
×	×	×	60.5	63.0	25.1
✓	×	×	62.9	65.8	25.9
×	✓	×	63.3	67.1	26.2
×	×	✓	63.1	66.8	26.5
✓	✓	×	66.7	69.8	27.1
✓	✓	✓	67.5	72.1	27.6

Table 5: Ablation on Fine-grained Pre-training Tasks.

vision	lang	joint	VG	Cap	QA
×	×	×	53.1	58.7	23.7
✓	×	×	62.3	65.4	24.5
✓	✓	×	65.1	68.1	25.5
✓	✓	✓	67.5	72.1	27.6

Table 6: Ablation of Synthetic-to-Realistic Adaptation.

3D Question Answering. We evaluate our model on the ScanQA dataset [Azuma *et al.*, 2022]. The quantitative comparison (Table 4) shows our method surpasses existing models on both EM@1 (+0.6%/+1.1%), EM@10(+0.4%/+0.4%) and captioning metric ROUGE(+0.6%/+0.5%). While the textual input of 3D-QA is a question rather than referring expression, our method still benefits the 3D QA task with the multi-modal features from multi-grained alignment learning.

5.3 Ablation Studies

We conduct the ablation studies on three tasks in Tables 5 to 7. ‘VG’ stands for overall accuracy of 3D visual grounding on Nr3D, ‘Cap’ for the CIDEr@0.25 result on Scan2Cap and ‘QA’ for the EM@1 results(w/ objects) on ScanQA.

Effectiveness of the fine-grained pre-training tasks. We evaluate the effectiveness of fine-grained alignment tasks in enhancing 3D visual grounding, dense captioning, and question-answering tasks (Table 5): (1) *ORP* (Row 2): Compared to the baseline (Row 1), *ORP* achieves improvements of +2.4%, +2.8%, and +0.8% for 3D visual grounding, dense captioning, and question-answering tasks, respectively. This suggests *ORP*’s effectiveness in understanding object-environment relations in 3D scenes. (2) *VRWA* (Row 3): *VRWA* leads to gains of +2.8%, +4.1%, and +1.1%. Its focus on viewpoint-based region-word alignment seems to enhance adaptability from various perspectives. (3) *MRWA* (Row 4): Incorporating *MRWA* results in +2.6%, +3.8%, and +1.4% improvements. *MRWA*’s multi-granularity alignment (object/room/scene) appears to enhance model perception in 3D environments. (4) *VRWA and MRWA v.s. ORP*: Both *VRWA* and *MRWA* show more significant improvements in visual grounding and question-answering than *ORP*, likely due to the enhanced perception across multiple perspectives and granularities from *VRWA* and *MRWA*, that is required by *VG* and *QA* tasks. (5) *Integrating All Tasks*: The combination of *ORP*, *VRWA*, and *MRWA* yields the highest gains: +7.0%, +9.1%, and +2.5%, demonstrating their capabilities.

Effectiveness of synthetic-to-realistic domain adaptation. We evaluated the effectiveness of utilising three specific domain adaptation discriminators that address the domain gap

Object	Room	House	VG	Cap	QA
✓	×	×	66.7	69.8	27.1
✓	✓	×	67.1	71.0	27.5
✓	×	✓	66.9	70.6	27.3
✓	✓	✓	67.5	72.1	27.6

Table 7: Ablation of Multi-level Region-Word Alignment .

between synthetic and real-world down-stream data (Table 6). (1) *3D Vision Domain Adaptation* (Row 2): By applying adaptation to the 3D vision domain, we observe performance increases of 9.2%, 6.7%, and 0.8% for 3D visual grounding, dense captioning, and question-answering tasks, respectively. This highlights the significance of 3D domain adaptation. (2) *Language Domain Adaptation* (Row 3): Extending adaptation to the language domain further enhances results, with gains of 2.8%, 2.7%, and 1.0%. This step demonstrates the value of refining the language domain. (3) *Joint Vision-Language Domain Adaptation* (Row 4): Finally, incorporating vision-language joint domain adaptation leads to additional improvements of 2.4%, 4.0%, and 2.1%. This confirms the effectiveness of our comprehensive approach in reducing domain discrepancies and improving overall performance.

Ablation of MRWA levels. We analyze the impact of different *MRWA* levels (object, room and house) across three tasks (Table 7), by comparing with baseline object-only alignment. (1) *Room-Level Alignment* (Row 2): Adding room-level alignment to the baseline object-only alignment (Row 1) yields improvements of 0.4%, 1.2%, and 0.4%. This suggests that room-level associations between words and objects enhance the model’s understanding of object relationships within rooms. (2) *House-Level Alignment* (Row 3): Extending alignment to the house level (whole scene) leads to gains of 0.2%, 0.8%, and 0.2%. This level of alignment assists the model in comprehending the entire scene by providing associations across the house. (3) *Combining All Levels* (Row 4): Combining alignments at both room and house levels results in the most substantial improvements of 0.8%, 2.3%, and 0.5%. It indicates the effectiveness of multi-level alignment in improving the model’s 3D scene perception.

6 Conclusion

To tackle the scarcity of data and concerns regarding data collection cost for the 3D vision-language pretraining, we employ synthetic data generation using 3D simulators and off-the-shelf models and create a substantial and diverse synthetic dataset SynVL3D. With the rich annotations in SynVL3D, we pre-train a model SynFormer3D aligning 3D and language with diverse auxiliary tasks, including predicting object relationships, multi-level and view-aggregated region-word alignment. A domain adaptation strategy in fine-tuning is adapted to downstream tasks. Extensive comparisons on various benchmarks show the effectiveness of our method.

Acknowledgements

This work was supported by the grants from the National Natural Science Foundation of China 62372014.

References

- [Abdelreheem *et al.*, 2024] Ahmed Abdelreheem, Kyle Olszewski, Hsin-Ying Lee, Peter Wonka, and Panos Achlioptas. Scanents3d: Exploiting phrase-to-3d-object correspondences for improved visio-linguistic models in 3d scenes. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 3524–3534, 2024.
- [Achlioptas *et al.*, 2020] Panos Achlioptas, Ahmed Abdelreheem, Fei Xia, Mohamed Elhoseiny, and Leonidas Guibas. Referit3d: Neural listeners for fine-grained 3d object identification in real-world scenes. In *ECCV*, pages 422–440. Springer, 2020.
- [Alayrac *et al.*, 2022] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katie Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. *arXiv preprint arXiv:2204.14198*, 2022.
- [Azuma *et al.*, 2022] Daichi Azuma, Taiki Miyanishi, Shuhei Kurita, and Motoaki Kawanabe. Scanqa: 3d question answering for spatial scene understanding. In *CVPR*, pages 19129–19139, 2022.
- [Bakr *et al.*, 2022] Eslam Mohamed Bakr, Yasmeen Alsaedy, and Mohamed Elhoseiny. Look around and refer: 2d synthetic semantics knowledge distillation for 3d visual grounding. *NeurIPS*, 2022.
- [Brown *et al.*, 2020] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *NeurIPS*, 33:1877–1901, 2020.
- [Cai *et al.*, 2022] Daigang Cai, Lichen Zhao, Jing Zhang, Lu Sheng, and Dong Xu. 3djcg: A unified framework for joint dense captioning and visual grounding on 3d point clouds. In *CVPR*, pages 16464–16473, 2022.
- [Chen and Liu, 2020] Qingchao Chen and Yang Liu. Structure-aware feature fusion for unsupervised domain adaptation. In *AAAI Conference on Artificial Intelligence (AAAI)*, pages 10567–10574, 2020.
- [Chen *et al.*, 2018] Qingchao Chen, Yang Liu, Zhaowen Wang, Ian Wassell, and Kevin Chetty. Re-weighted adversarial adaptation network for unsupervised domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7976–7985, 2018.
- [Chen *et al.*, 2020] Dave Zhenyu Chen, Angel X Chang, and Matthias Nießner. Scanrefer: 3d object localization in rgb-d scans using natural language. In *ECCV*, pages 202–221. Springer, 2020.
- [Chen *et al.*, 2021] Zhenyu Chen, Ali Gholami, Matthias Nießner, and Angel X Chang. Scan2cap: Context-aware dense captioning in rgb-d scans. In *CVPR*, pages 3193–3203, 2021.
- [Chen *et al.*, 2022] Shizhe Chen, Pierre-Louis Guhur, Makarand Tapaswi, Cordelia Schmid, and Ivan Laptev. Language conditioned spatial relation reasoning for 3d object grounding. *NeurIPS*, 2022.
- [Chen *et al.*, 2023] Suyi Chen, Hao Xu, Ru Li, Guanghui Liu, Chi-Wing Fu, and Shuaicheng Liu. Sira-pcr: Sim-to-real adaptation for 3d point cloud registration. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14394–14405, 2023.
- [Dai *et al.*, 2023] Wenliang Dai, Junnan Li, Dongxu Li, Anthony Meng Huat Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale Fung, and Steven Hoi. Instructclip: Towards general-purpose vision-language models with instruction tuning. *arXiv preprint arXiv:2305.06500*, 2023.
- [DeBortoli *et al.*, 2021] Robert DeBortoli, Li Fuxin, Ashish Kapoor, and Geoffrey A Hollinger. Adversarial training on point clouds for sim-to-real 3d object detection. *IEEE Robotics and Automation Letters*, 6(4):6662–6669, 2021.
- [Deitke *et al.*, 2022] Matt Deitke, Eli VanderBilt, Alvaro Herrasti, Luca Weihs, Kiana Ehsani, Jordi Salvador, Winsong Han, Eric Kolve, Aniruddha Kembhavi, and Roozbeh Mottaghi. Proctor: Large-scale embodied ai using procedural generation. *Advances in Neural Information Processing Systems*, 35:5982–5994, 2022.
- [Delitzas *et al.*, 2023] Alexandros Delitzas, Maria Parelli, Nikolas Hars, Georgios Vlassis, Sotirios Anagnostidis, Gregor Bachmann, and Thomas Hofmann. Multi-clip: Contrastive vision-language pre-training for question answering tasks in 3d scenes. *arXiv preprint arXiv:2306.02329*, 2023.
- [Ding *et al.*, 2022] Runyu Ding, Jihan Yang, Li Jiang, and Xiaojuan Qi. Doda: Data-oriented sim-to-real domain adaptation for 3d semantic segmentation. In *European Conference on Computer Vision*, pages 284–303. Springer, 2022.
- [Guo *et al.*, 2023] Ziyu Guo, Yiwen Tang, Renrui Zhang, Dong Wang, Zhigang Wang, Bin Zhao, and Xuelong Li. Viewrefer: Grasp the multi-view knowledge for 3d visual grounding with gpt and prototype guidance. *arXiv preprint arXiv:2303.16894*, 2023.
- [He *et al.*, 2021] Dailan He, Yusheng Zhao, Junyu Luo, Tianrui Hui, Shaofei Huang, Aixi Zhang, and Si Liu. Transrefer3d: Entity-and-relation aware transformer for fine-grained 3d visual grounding. In *Proceedings of the 29th ACM International Conference on Multimedia*, pages 2344–2352, 2021.
- [Huang *et al.*, 2022] Shijia Huang, Yilun Chen, Jiaya Jia, and Liwei Wang. Multi-view transformer for 3d visual grounding. In *CVPR*, pages 15524–15533, 2022.
- [Huang *et al.*, 2023] 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.
- [Jiao *et al.*, 2022] Yang Jiao, Shaoxiang Chen, Zequn Jie, Jingjing Chen, Lin Ma, and Yu-Gang Jiang. More: Multi-order relation mining for dense captioning in 3d scenes.

- In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXV*, pages 528–545. Springer, 2022.
- [Jin *et al.*, 2023] Zhao Jin, Munawar Hayat, Yuwei Yang, Yulan Guo, and Yinjie Lei. Context-aware alignment and mutual masking for 3d-language pre-training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10984–10994, 2023.
- [Loshchilov and Hutter, 2019] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *ICLR*, 2019.
- [Lu *et al.*, 2019] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. *NeurIPS*, 32, 2019.
- [Luo *et al.*, 2022] Junyu Luo, Jiahui Fu, Xianghao Kong, Chen Gao, Haibing Ren, Hao Shen, Huaxia Xia, and Si Liu. 3d-sps: Single-stage 3d visual grounding via referred point progressive selection. In *CVPR*, pages 16454–16463, 2022.
- [Mo and Liu, 2024] Wentao Mo and Yang Liu. Bridging the gap between 2d and 3d visual question answering: A fusion approach for 3d vqa. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2024.
- [Parelli *et al.*, 2023] Maria Parelli, Alexandros Delitzas, Nikolas Hars, Georgios Vlassis, Sotirios Anagnostidis, Gregor Bachmann, and Thomas Hofmann. Clip-guided vision-language pre-training for question answering in 3d scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5606–5611, 2023.
- [Pei *et al.*, 2023] Jiangbo Pei, Zhuqing Jiang, Aidong Men, Liang Chen, Yang Liu, and Qingchao Chen. Uncertainty-induced transferability representation for source-free unsupervised domain adaptation. *IEEE Transactions on Image Processing (TIP)*, 2023.
- [Pei *et al.*, 2024] Jiangbo Pei, Aidong Men, Yang Liu, Xiahai Zhuang, and Qingchao Chen. Evidential multi-source-free unsupervised domain adaptation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024.
- [Radford *et al.*, 2018] Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. Improving language understanding by generative pre-training. *OpenAI Blog*, 2018.
- [Radford *et al.*, 2019] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- [Radford *et al.*, 2021] 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*, pages 8748–8763. PMLR, 2021.
- [Wald *et al.*, 2019] Johanna Wald, Armen Avetisyan, Nassir Navab, Federico Tombari, and Matthias Nießner. Rio: 3d object instance re-localization in changing indoor environments. In *ICCV*, pages 7658–7667, 2019.
- [Wald *et al.*, 2020] Johanna Wald, Helisa Dharmo, Nassir Navab, and Federico Tombari. Learning 3d semantic scene graphs from 3d indoor reconstructions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3961–3970, 2020.
- [Xu *et al.*, 2022] Yinsong Xu, Zhuqing Jiang, Aidong Men, Yang Liu, and Qingchao Chen. Delving into the continuous domain adaptation. In *Proceedings of the 30th ACM international conference on Multimedia*, 2022.
- [Yang *et al.*, 2021] Zhengyuan Yang, Songyang Zhang, Liwei Wang, and Jiebo Luo. Sat: 2d semantics assisted training for 3d visual grounding. In *ICCV*, pages 1856–1866, 2021.
- [Yeshwanth *et al.*, 2023] Chandan Yeshwanth, Yueh-Cheng Liu, Matthias Nießner, and Angela Dai. Scannet++: A high-fidelity dataset of 3d indoor scenes. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 12–22, October 2023.
- [Yuan *et al.*, 2022a] Zhihao Yuan, Xu Yan, Zhuo Li, Xuhao Li, Yao Guo, Shuguang Cui, and Zhen Li. Toward explainable and fine-grained 3d grounding through referring textual phrases. *arXiv preprint arXiv:2207.01821*, 2022.
- [Yuan *et al.*, 2022b] Zhihao Yuan, Xu Yan, Yinghong Liao, Yao Guo, Guanbin Li, Shuguang Cui, and Zhen Li. X-trans2cap: Cross-modal knowledge transfer using transformer for 3d dense captioning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8563–8573, June 2022.
- [Zhao *et al.*, 2021a] Lichen Zhao, Daigang Cai, Lu Sheng, and Dong Xu. 3dvg-transformer: Relation modeling for visual grounding on point clouds. In *ICCV*, pages 2928–2937, 2021.
- [Zhao *et al.*, 2021b] Sicheng Zhao, Yezhen Wang, Bo Li, Bichen Wu, Yang Gao, Pengfei Xu, Trevor Darrell, and Kurt Keutzer. epointda: An end-to-end simulation-to-real domain adaptation framework for lidar point cloud segmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 3500–3509, 2021.
- [Zhao *et al.*, 2023] Zijing Zhao, Sitong Wei, Qingchao Chen, Dehui Li, Yifan Yang, Yuxin Peng, and Yang Liu. Masked retraining teacher-student framework for domain adaptive object detection. In *ICCV*. IEEE, 2023.
- [Zhu *et al.*, 2023] Ziyu Zhu, Xiaojian Ma, Yixin Chen, Zhi-dong Deng, Siyuan Huang, and Qing Li. 3d-vista: Pre-trained transformer for 3d vision and text alignment. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2911–2921, 2023.