


# R3DS: Reality-linked 3D Scenes for Panoramic Scene Understanding

Qirui Wu<sup>1</sup>, Sonia Raychaudhuri<sup>1</sup>, Daniel Ritchie<sup>2</sup>, Manolis Savva<sup>1</sup>, and  
Angel X. Chang<sup>1,3</sup>

<sup>1</sup> Simon Fraser University

<sup>2</sup> Brown University

<sup>3</sup> Alberta Machine Intelligence Institute (Amii)



**Fig. 1:** **Left:** the Reality-linked 3D Scenes dataset (R3DS) fills a gap between synthetic 3D scenes and reconstructions of real-world environments by providing 3D scene proxies linked to real-world panoramas from Matterport3D (three example panoramas and 3D scenes shown). **Right:** our dataset contains scenes with higher density and completeness compared to prior datasets, and provides additional annotations such as object support (what objects or architectural elements support other objects), and matching object sets (e.g., pairs of the same nightstand). We use our dataset for the panoramic scene understanding task and demonstrate its value for research on room layout estimation, as well as 2D and 3D object detection.

**Abstract.** We introduce the Reality-linked 3D Scenes (R3DS) dataset of synthetic 3D scenes mirroring the real-world scene arrangements from Matterport3D panoramas. Compared to prior work, R3DS has more complete and densely populated scenes with objects linked to real-world observations in panoramas. R3DS also provides an object support hierarchy, and matching object sets (e.g., same chairs around a dining table) for each scene. Overall, R3DS contains 19K objects represented by 3,784 distinct CAD models from over 100 object categories. We demonstrate the effectiveness of R3DS on the Panoramic Scene Understanding task. We find that: 1) training on R3DS enables better generalization; 2) support relation prediction trained with R3DS improves performance compared to heuristically calculated support; and 3) R3DS offers a challenging benchmark for future work on panoramic scene understanding.

## 1 Introduction

Datasets of 3D indoor environments are increasingly used for research on scene understanding [1, 3, 28], embodied AI [2, 19, 23], and scene generation [15, 25]. There are two strategies for constructing 3D scene datasets: reconstruction of real-world spaces [3], or authoring scenes using synthetic 3D objects [8] (CAD models). Reconstruction captures real spaces but is hard to scale, and the resulting scenes exhibit imperfections and artifacts. On the other hand, synthetic 3D scenes are complete and easy to manipulate but often do not match the statistics of real-world spaces and are artificially “clean”. Moreover, both strategies are time-consuming and require expertise. There have been some attempts to create “synthetic” replicas of real environments by matching CAD models to objects in scans [19, 23]. These efforts have been limited in scale and often result in partial and sparsely populated synthetic counterparts of the real environments.


We design a framework that allows users to create 3D scenes from RGB panoramas and use it to create **R3DS**: a dataset of ‘Reality-linked’ **3D Scenes**. Each 3D scene in our dataset is a complete proxy of an environment from the Matterport3D [3] dataset, representing both the 3D architecture and the objects. Thus, each scene is linked to a real space, with correspondences established between panorama observations of each object and the synthetic object. These reality-linked scenes reflect denser real-world arrangements of objects.

The use of panoramas for reference is advantageous compared to either perspective images or 3D reconstructions. Panoramas are not limited by the field of view unlike perspective images, enabling more complete 3D synthetic proxies. Panoramas also better capture relatively small objects and objects with challenging materials or illumination conditions compared to reconstructions. Additionally, there is a scarcity of synthetic 3D scenes coupled with real-world panoramas, with only one relatively small algorithmically constructed dataset provided by Zhang *et al.* [28] being available to the community.

Compared to prior efforts such as Scan2CAD [1] and CAD-Estate [14], our dataset provides more complete scenes, with salient observed objects being captured in the layout. Moreover, we provide a support hierarchy defining what objects are placed on other objects and specify sets of identical objects such as dining chairs around a table, allowing for creating realistic variations of the scene by swapping the entire set to a different chair design.

We demonstrate the value of our dataset by using it for the Panoramic Scene Understanding task. We show that leveraging the denser layouts and support hierarchy information in our scenes leads to improved object detection performance and better generalization compared to training using other datasets previously used for this task. In summary, we make the following contributions:

- We design a framework for efficient construction of synthetic scenes from real panoramas and use it to create R3DS: a dataset of reality-linked 3D scenes.
- R3DS provides more complete and realistic scenes with correspondences between real and synthetic objects, and object-object support relations.



**Fig. 2: Dataset comparison.** (Top) shows different views of a scene annotated in R3DS. Comparison with previous datasets (bottom) shows (1) R3DS has more complete scenes than the previous datasets; (2) Objects in R3DS are properly supported by either architecture or other objects unlike the others (e.g. floating objects with no proper support); (3) R3DS is annotated using the same 3D model for objects arranged together (chairs by the dining table, couches arranged together).

- We show that the more complete layouts and support relations in our dataset enable better performance and generalization in the Panoramic Scene Understanding task, and that our dataset offers a challenging benchmark for future work in scene understanding.

## 2 Related Work

**3D scene datasets.** A spectrum of scene datasets have been used for scene understanding tasks. One type provides annotated 3D reconstructions of real scenes based on RGB-D videos [3, 6, 10, 16, 20, 27]. These datasets are usually subject to the limitations of RGB-D reconstruction, typically containing noise, artifacts such as holes, and poor reconstructions of thin structures, shiny objects, or light sources. Another type of 3D datasets is authored by manually designing 3D object assets [5, 9] and inserting them into synthetic 3D scenes [8]. However, such datasets lack the realism of real-world reconstructions and demand expert knowledge, making them expensive to create. A third, hybrid approach which is closest to our work creates 3D scene datasets by aligning existing object CAD models to real world data.

**Datasets that align CAD models to real world.** There have been a number of recent efforts in aligning CAD models with real-world data. Prior work [13, 22, 26] has annotated object images with 3D models, typically using keypoint correspondences to perspective images. These perspective images usually do not

depict a complete scene; they typically focus on one or two objects and are limited in field of view, resulting in a sparse proxy of the real scene.

Another line of work aligns 3D CAD models to RGB-D scans either through annotation as in Scan2CAD [1], or automated heuristics as in iGibson [19]. OpenRooms [12] extends Scan2CAD [1] with photorealistic material annotations and focuses on inverse rendering tasks. Conceptually, these allow for more complete synthetic scene proxies. However, statistics from these datasets show that they are still relatively sparse (see Tab. 1). In addition, the poor quality of reconstruction makes aligning CAD models challenging without referring to the original RGB images. A prominent exception is Replica [20] which has fairly high-quality reconstructions and the artist-created Replica-CAD [23]. However, creating such high quality “replicas” is labor intensive and costly. Szot *et al.* [23] report 900+ work hours required to model approximately 90 objects, resulting in a dataset of limited scale with 105 different layouts of what is effectively a single room.

More recently, Maninis *et al.* [14] introduced CAD-Estate, which aligns CAD models to RGB videos for over 19K spaces. Because the data is based on monocular video, the coverage of the spaces is incomplete. In addition, the annotation is relatively sparse, with an average of only 6 objects per scene.

**Datasets for panoramic scene understanding.** There have been relatively few datasets introduced for Panoramic Scene Understanding [7, 28, 29]. In the initial PanoContext dataset [29], the data did not have aligned CAD models and only included object cuboids. The ground truth data was collected on 2D panorama images by annotating visible cuboid vertices; 3D cuboids were obtained by minimizing the re-projection error from the annotated 2D vertices. Moreover, these 3D cuboids and the room layout are obtained with the assumption that the room layout is a cuboid and that the objects are vertically aligned. Thus, the resulting object layout may deviate from the real arrangement of objects. More recently, datasets for Panoramic Scene Understanding have been built by taking 3D scans, aligning CAD objects to them, and then generating panoramas [7, 28]. Compared to these datasets, our R3DS is manually curated for a larger number of distinct regions and provides support hierarchy and matching object set annotations.

### 3 The R3DS Dataset

We describe the construction of the R3DS dataset and present a statistical analysis of the scenes it contains. Compared to previous datasets [1, 14, 28] (Fig. 2), our scenes are more densely populated, and objects are annotated with a hierarchy of support relations. Moreover, our dataset specifies matching object instances in furniture arrangements. Figure 1 shows example annotations from our dataset.

#### 3.1 Dataset construction

We developed a 3D annotation interface (Fig. 3) showing a panorama of a room from Matterport3D and allowing users to insert 3D CAD objects into a 3D scene



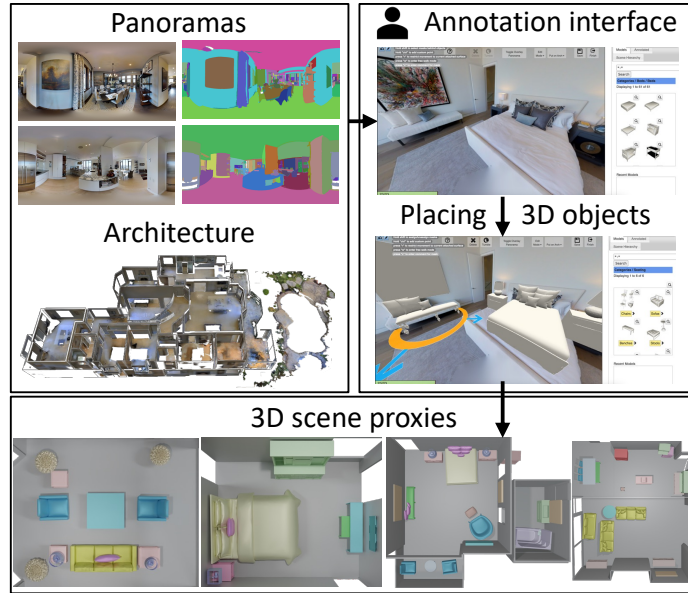
**Table 1: Comparison with 3D indoor scene datasets aligned with real-world images, videos, or scans.** Our R3DS dataset contains more densely populated annotations compared to other datasets, with objects from 110 different categories. We report the unique models (#CAD), object categories (#Cat), object instances (#Objects) as well as average number of objects and object categories per annotation. For Scan2Cad [1] and ReplicaCAD [23] the average is per scan. Note that ReplicaCAD consists of 105 different layouts (arrangements) of effectively one room. CAD-Estate has partial views into 19K spaces, many of which are 1-2 rooms. Of the datasets used for panoramic scene understanding, our R3DS dataset covers more rooms with 842 panoramas over 22 room types. Replica-Pano (in gray) was not released, so we report statistics from the paper. Our annotations per panorama are more complete and our dataset has both support relations (Sup) and matching object instance sets (Match).

Dataset	Source	CAD Alignment	Type	Houses/Rooms	Panos	#CAD	#Objects	#Cat	Ave Obj	Ave Cat	Sup	Match
Scan2CAD [1]	ScanNet [6]	Annotator	scan	- / 1506	✗	3,049	14,225	35	9.4	4.1	✗	✗
OpenRooms [12]	ScanNet [6]	Scan2CAD [1]	scan	- / 1288	✗	2,651	16,014	38	12.4	6.3	✗	✗
ReplicaCAD [23]	Replica [20]	Artist recreation	scan	- / 105*	✗	92	2,293	44	21.8	14.4	✗	✗
CAD-Estate [14]	RealEstate10K [31]	Annotator	video	19,512	✗	12,024	100,882	49	6.3	3.4	✗	✗
Replica-Pano [7]	Replica [20]	Heuristic	pano	- / 27	2700	-	-	25	-	-	✗	✗
iGibson-DPC [28]	iGibson [19]	Heuristic	pano	15 / 100	1500	500	26,998	57	17.9	10.2	✗	✗
<b>R3DS (Ours)</b>	Matterport3D [3]	Annotator	pano	20 / 370	842	3,784	19,050	110	22.9	10.4	✓	✓

which is *visually overlaid* on the panorama. The 3D scene is initially empty, consisting only of 3D architectural geometry which specifies the walls, floor, ceiling as well as the placement of openings (e.g. doors, windows, and other openings) on the walls. We create this 3D architecture by taking 20 houses from Matterport3D, constructing an initial architecture based on the region and object annotations for the windows and doors, and manually refining the placement of walls and openings. By combining panoramas and 3D architectures, users can see through openings and annotate objects located in other rooms.

We ask annotators to select and place 3D object models to best match the panoramic image. We use CAD models from Wayfair [17] and ShapeNet [4] models collected from 3D Trimble warehouse. Wayfair provides a large collection of furniture CAD models that match real-world products and are sized based on real-world dimensions. However, it does not include bathroom fittings, electronic equipment and kitchen appliances, for which we manually scale and align CAD models from ShapeNet. Compared with ShapeNetCore, the CAD models we use are already sized to real-world sizes (instead of normalized to a unit cube).

To assist the annotators, we provide segmented masks of objects visible in the panorama. Since Matterport3D has annotated 3D object masks on the scans we use those annotations, but it is also possible to run an instance segmentation on the panorama. When the user clicks one of these masks, a search panel automatically opens and shows objects matching the clicked mask category label. For each mask, the annotator selects a matching object and positions and aligns it to match the mask. Annotators are instructed to choose objects which match the *shape* of the corresponding object in the panorama (rather than its color or texture). To help annotators focus on shape, we render all 3D objects in a neutral gray color. Annotators are also explicitly asked to select the same




**Fig. 3: R3DS annotation pipeline.** Annotators see an empty scene (architecture only). They then insert and manipulate 3D object models from a panorama viewpoint to create a populated 3D scene proxy corresponding to the panorama.

3D asset for objects that should be the same; our interface provides a list of recently selected assets to make this process easier. In addition, annotators are instructed to add annotations for any objects that are not segmented (due to errors in Matterport3D) through simple clicks. These additional objects provide a more complete annotation that covers poorly reconstructed objects such as glass tables, lamps, and other small objects. The interface enforces that each object is placed on a support surface (either an architecture element or another object). The annotator can review their work by toggling off the panorama overlay or by switching to a perspective view of the 3D scene. For more annotated scene examples and details on the annotation process please refer to the supplement.

### 3.2 Dataset analysis and statistics

We collect annotations for 20 Matterport3D houses with 808 panoramas in total. We discard panoramas taken on stairs or outside a house, since they have a limited number of objects that can be placed. After filtering we have 769 panoramas for our analysis and experiments. For 73 panoramas, we collect two sets of annotations for each, to obtain a total of 842 annotated object arrangements across 22 different Matterport3D region types. The panoramas with two annotations serve as a test of annotator consistency and add diversity. In total, R3DS contains 19,050 object instances from 3,784 unique 3D CAD models spanning



**Fig. 4: Architecture comparison.** Compared to Scan2CAD (no architecture) and CAD-Estate (partial architecture), R3DS provides complete architecture with door/window portals.

over 110 fine-grained object categories. Table 1 shows a comparison of overall statistics with previous 3D indoor scene datasets. See the supplement for more statistics about annotated objects.

Compared to prior datasets that align CAD models to real-world scenes, R3DS is more complete, providing annotated object support hierarchies and matching object instances. CAD-Estate [14] annotates RGB videos with 3D objects and architectural room layouts. However, the architecture is partial as the videos have limited view (Fig. 4), and not all objects in the scenes are annotated (Fig. 2). This results in annotations of objects floating mid-air and not properly supported (e.g. lamp Fig. 2). Scan2CAD [1] also lacks support structure (e.g., lamp not supported by cabinet in Fig. 2). In addition, because Scan2CAD does not provide clean 3D architecture on which objects are placed (Fig. 4), objects on the floor are not always placed such that their bottom face is parallel with a horizontal plane. In contrast, our R3DS scenes have an accurate support hierarchy by construction. OpenRooms [12] augments Scan2CAD with room layouts representing the architecture. However, the architecture in R3DS is more complex and realistic, especially due to inclusion of more doors (1.92 doors per room in R3DS vs 0.67 in OpenRooms).

Evaluation of CAD object annotation quality is non-trivial as the ‘ground truth’ from the semantically annotated 3D reconstructions is itself imperfect. We measured how closely our annotated CAD objects conform to the real objects using the average 2D IoU between CAD object mask and ground truth 2D mask. R3DS is at 42.6% vs 38.5% for Scan2CAD, across 8 common object categories (bed, sofa, chair, cabinet, tv/monitor, table, shelving, bathtub).

Of the datasets previously used for Panoramic Scene Understanding, Replica-Pano [7] has not been released, and iGibson-DPC [28] is the only dataset with synthetic panoramic images annotated with 3D objects and room layout. iGibson-DPC is built on scenes from iGibson [18] by randomly replacing objects with different models from the same category and rendering using the iGibson simulator to render panoramas. The selection and placement of objects in iGibson-DPC is based on heuristic algorithms, while our R3DS is manually annotated and placed. 3D models are verified in terms of match and alignment to the object masks.

Moreover, iGibson-DPC contains unrealistic object arrangements (e.g., floating TV in Fig. 2).

## 4 Experiments

We showcase the value of R3DS on the Panoramic Scene Understanding (PanoSun) task [7, 28, 29]. Given an input RGB panorama, the goal is to estimate the room layout, detect objects in 2D, estimate their 3D oriented bounding boxes and also reconstruct 3D object meshes. Our experiments show that methods trained on R3DS data benefit from its realism and generalize better when evaluated on photorealistic images. We also investigate the role of object support hierarchy information in improving performance.

### 4.1 Task setup

**Method.** DeepPanoContext (DPC) [28] predicts the room layout, detects objects in 3D and recovers object meshes from a panorama image using a relation-based graph convolutional network and a differentiable relation optimization procedure. Since DPC has a publicly-available implementation, we use it to benchmark the R3DS data on the PanoSun task. We keep all hyperparameters unchanged except lowering the relation optimization loss weight of 3D bounding box back-projection from 10 to 1, since the ground truth 2D masks are noisy.

**Datasets.** We train and evaluate DPC on the iGibson-DPC (IG) [7, 18, 28], Structured3D (S3D) [30], and R3DS datasets. Zhang *et al.* [28] render 1,500 panoramas from 15 iGibson houses composed of 500+ objects spanning 57 object categories. We use the same data and splits for IG. Structured3D consists of 3500 houses and around 18K photo-realistic rendered panoramas in total. We use 14K for training and the remaining 4K for testing. Note that Structured3D does not provide ground truth object meshes.

To prepare R3DS for this task, we generate the ground truth room layout from the 3D architecture based on the camera viewpoint and obtain 3D oriented bounding boxes (OBBs) from all objects. We use 2D object masks from the Matterport3D mesh instance segmentation. We consider three variants of R3DS based on the input panorama: *R3DS-real* where we use the Matterport3D panoramas, *R3DS-syn* where we use rendered panoramas (at the same camera poses) from the annotated synthetic scenes, and *R3DS-mix* where we combine the two types of panoramas and double the available data. We follow the MP3D house split and merge the train and val sets to obtain a disjoint split of 15 train and 5 test houses. Based on the split, we have 696 annotated panoramas for train and 146 for test. To fairly evaluate methods trained on different datasets, we curate a list of 25 object classes common to all datasets.

**Metrics.** Following Zhang *et al.* [28] we use separate metrics for room layout estimation, 3D object detection, and scene relation prediction. For room layout estimation, we use 2D IoU for predicted 2D floorplan, 3D IoU for lifted 3D room geometry, and dRMSE for predicted depth with respect to the camera location.

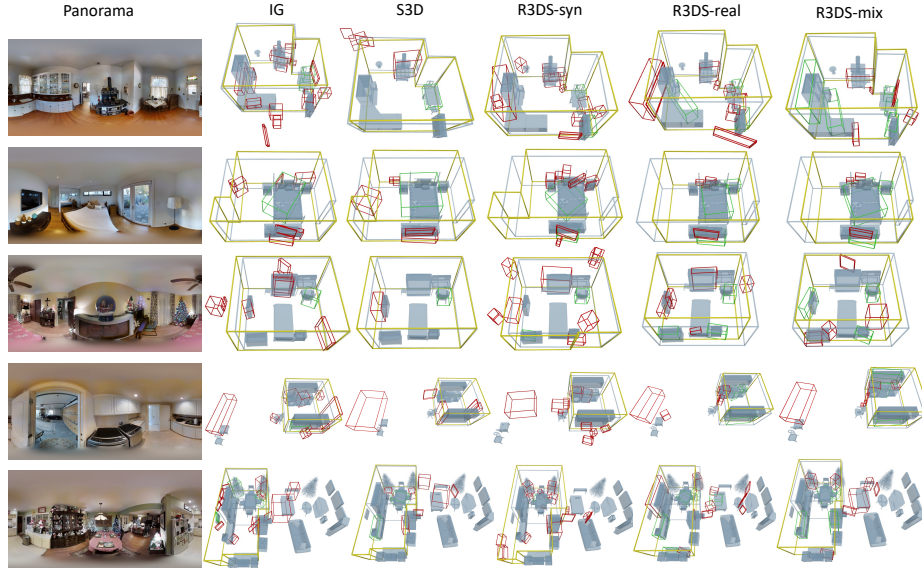
**Table 2: Room layout estimation on *R3DS-real* test set.** DPC [28] was pre-trained on IG and S3D. For the last three rows, we fine-tune the pretrained weights on variants of R3DS.

Train	2D IoU $\uparrow$	3D IoU $\uparrow$	dRMSE $\downarrow$
DPC [28]	53.4	50.3	0.682
R3DS-real	55.1	53.1	0.610
R3DS-syn	59.0	56.1	0.629
R3DS-mix	<b>59.6</b>	<b>57.0</b>	<b>0.572</b>

**Table 3: Cross-dataset evaluation for the Panoramic Scene Understanding task.** We evaluate 3D detections with class-agnostic IoU and mAP at IoU of 0.15, and report object collisions. The highlighted rows indicate the most challenging scenario.

Test	Train	3D detection $\uparrow$		Collision $\downarrow$		Attachment F1 $\uparrow$			
		IoU	mAP	mesh	arch	obj	wall	floor	ceil
IG	IG	<b>27.5</b>	<b>30.3</b>	1.662	2.594	53.1	<b>76.8</b>	<b>95.0</b>	<b>86.2</b>
	IG+R3DS	24.0	30.2	1.404	2.254	<b>59.7</b>	64.1	94.6	2.7
	R3DS-real	17.3	13.4	<b>0.242</b>	1.456	38.8	64.0	92.8	0.0
	R3DS-syn	23.2	14.2	0.480	1.938	48.5	46.7	93.8	28.6
	R3DS-mix	21.6	15.6	0.434	<b>1.248</b>	43.1	67.2	90.1	9.8
S3D	IG	19.5	3.5	1.016	2.651	<b>50.9</b>	<b>68.7</b>	90.8	<b>11.6</b>
	IG+R3DS	<b>19.7</b>	7.0	0.868	2.089	52.0	67.4	<b>91.2</b>	1.8
	R3DS-real	18.4	7.1	0.600	2.598	45.0	61.6	89.7	0.7
	R3DS-syn	19.0	4.8	0.644	2.561	49.3	49.7	<b>91.2</b>	2.4
	R3DS-mix	19.6	<b>7.5</b>	<b>0.463</b>	<b>1.673</b>	47.8	64.1	87.2	0.9
R3DS	IG	15.6	5.9	0.575	1.959	<b>53.8</b>	50.7	51.2	0.0
	IG+R3DS	17.5	14.1	0.281	1.267	49.5	<b>61.6</b>	58.6	0.0
	R3DS-real	16.4	15.0	0.226	1.562	44.0	57.3	58.9	0.0
	R3DS-syn	14.0	8.4	0.390	1.664	54.1	40.6	49.1	0.0
	R3DS-mix	<b>17.6</b>	<b>15.8</b>	<b>0.171</b>	<b>1.007</b>	48.5	58.3	<b>60.1</b>	0.0

For 3D object detection, we report bounding box-based class-agnostic 3D IoU as well as mean average precision (mAP) across the 25 object classes, where an IoU greater than 0.15 counts as a “true” result. For scene relation prediction, we report F1 scores for relation classification. We also report the average number of objects colliding with each other or with architectural structures (wall, floor, ceiling). Specifically, we follow Zhang *et al.* [28] and measure collisions using the Separating Axis Theorem (SAT) to test whether object bounding boxes overlap. Since bounding box-based collision is a poor proxy for real-world physical collision, we also compute mesh-based collision by checking if the meshes for object pairs have any interpenetrating triangles [11, 24].



**Fig. 5: Qualitative results for cross-dataset Panoramic Scene Understanding task.** Correct and incorrect object detections shown in green and red boxes. Ground truth room layout and meshes shown in gray color, while layout prediction is in yellow. Training on R3DS leads to fewer errors compared to training on other datasets, especially when mixed with real data.

## 4.2 Results

**1) Does R3DS help DPC generalize to real images?** Since the original DPC work only trained and evaluated on synthetic data, it is unclear how well it performs on realistic panoramic imagery. We hypothesize that training on R3DS will lead to better performance. We separately show results on room layout estimation (Tab. 2) and 3D object detection.

*Room layout estimation.* For room layout estimation, DPC uses HorizonNet [21] pretrained on iGibson (IG) and Structured3D (S3D) panoramas. This model achieves good performance on IG data (91.0 3D IoU). When directly testing the official pretrained model on *R3DS-real* panoramas, we notice a significant performance drop compared to results on the rendered panoramas from iGibson (Tab. 2 shows that the 3D IoU drops to 50.3). By finetuning the pretrained model with *R3DS-real*, we can predict more precise room layouts for real cluttered scenes. Even only trained on *R3DS-syn*, we outperform the original DPC model by 5.6% and 5.8% on 2D and 3D IoU, respectively. This is likely due to renderings from *R3DS-syn* reflecting more realistic object arrangements in a room instead of pushing all objects against walls. Best performance on 2D IoU (59.6), 3D IoU (57.0) and depth RMSE (0.572) is achieved by fine tuning on *R3DS-mix*.

**Table 4:** Performance of models trained on three synthetic datasets (IG, S3D, and IG+S3D) evaluated on the *R3DS-full* dataset, where “full” indicates all 840 panoramas are used for testing.

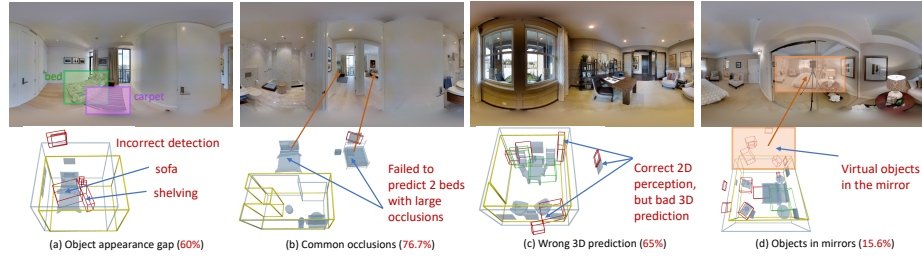
Train	3D detection		Collision ↓		Support F1 ↑		
	IoU↑	mAP↑	mesh	arch	obj	floor	ceiling
IG	14.2	5.2	0.703	1.639	<b>4.1</b>	85.3	0.0
S3D	16.4	<b>10.0</b>	<b>0.112</b>	1.226	3.1	<b>86.8</b>	0.0
IG+S3D	<b>17.1</b>	9.7	0.133	<b>1.162</b>	3.3	84.8	0.0

*Object detection.* For 3D object detection, we train DPC on different data settings and conduct a cross-dataset evaluation (see Tab. 3). To investigate how models perform on out-of-distribution scenes, we evaluate models on Structured3D, as its images are near-realistic. To explore whether DPC training benefits from R3DS given the same amount of data, we create a special data input *IG+R3DS* that combines iGibson and R3DS panoramas by randomly replacing half (500) of iGibson data with *R3DS-real* data. The results show that *IG+R3DS* performs almost the same as *IG* with fewer collisions on iGibson, but it remarkably outperforms *IG* on the test set of *R3DS* and *S3D* by 8.2 and 3.5 improvements on 3D mAP, respectively. It also averages 0.221 fewer mesh collisions. There are noticeable performance gaps on iGibson for models trained on R3DS data likely due to the data domain shift. Among the three variants of R3DS data, *R3DS-mix* outperforms the others on all three test sets regarding 3D IoU and 3D mAP with the fewest mesh and architecture collisions. Although *R3DS-syn* underperforms on *R3DS* and *S3D* test sets, it achieves better performance than *IG* with even less data.

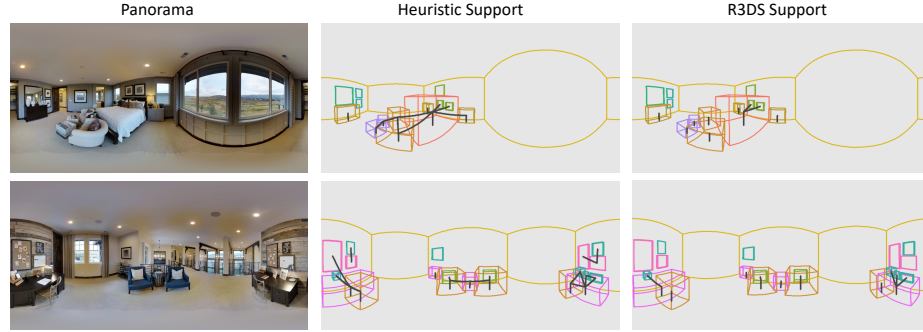
*Scene relation classification.* We report F1 scores for identifying attachment relationships of objects to other objects and architecture elements (see Tab. 3). We note that models trained with synthetic renderings perform better than those trained on real images. That is because synthetic renderings present cleaner and simpler scenes with fewer objects than real world and simpler illumination such that DPC finds it easier to learn object-object and object-architecture relations. Also, note that the predictions of object-ceiling attachments can be extremely low because few objects are attached to the ceiling in the ground truth data. We show qualitative examples in Figure 5.

**2) Is R3DS a challenging, high-quality test set?** How would a model trained on pure synthetic data perform on complex real data (R3DS)? Due to its modest scale, we propose using R3DS as a challenging, high-quality test set rather than a train set. Specifically, we evaluate the synthetic-to-real performance of DPC by training on iGibson and/or Structured3D and testing on all panoramas in R3DS-real. Table 4 shows that a model trained with Structured3D performs the best (10.0 3D mAP and 0.112 mesh collision) as it observes the most photo-realistic images. DPC benefits from the synthetic data for higher bounding box IoUs, since it possesses accurately aligned 3D bounding box and





**Fig. 6: Error analysis on R3DS panoramas using DPC pretrained on S3D.** Typical prediction errors are categorized into 4 groups: (a) Appearance gap between real and synthetic images. (b) Detection failure due to common occlusions. (c) Bad performance of monocular 3D prediction. (d) Perception error for objects in mirrors. Numbers in parentheses mean the error rates.



**Fig. 7: Comparison between heuristic and R3DS support relations.** Dark lines between objects show support relations. Heuristic support often mistakenly assigns relations to spatially close objects. In contrast, R3DS support relations are annotated.

more unoccluded objects. However, mAP performance is lower due to worse object recognition ability. All models struggle to predict correct object-wise support relations but do a better job of predicting object-floor support relations.

We conduct error analysis on 120 randomly sampled panoramas using the model pretrained on *S3D* to identify typical errors (see Fig. 6). Errors are categorized into 4 groups: (a) 60% panoramas have 2D perception errors due to the synthetic-to-real appearance gap; (b) 76.7% panoramas show detection failures due to occlusions; (c) 65% panoramas exhibit correct 2D detections but fail to correctly perform 3D predictions; and (d) 15.6% out of 45 panoramas with mirrors mistakenly predict virtual objects in mirrors.

**3) Are R3DS support relations helpful for PanoSun?** We investigate whether the support relationships between objects provided in our R3DS scene hierarchy help boost performance of holistic scene understanding. We augment DPC’s Relation Scene-GCN module with additional support relation prediction branches. Besides obtaining explicitly annotated scene support relations from



**Table 5:** Performance on *R3DS-real* of DPC models trained on variants of R3DS with different support relation settings. We compare the original model without support (*none*) against models supervised with support that is heuristically computed (*heur*) or annotated from R3DS (*anno*). Classification results are evaluated on annotated ground-truth scene hierarchy.

Train	Supp.	3D detection		Collision ↓		Support F1 ↑		
		IoU↑	mAP↑	mesh	arch	obj	floor	ceil
R3DS-real	none	16.4	15.0	0.226	1.562	-	-	-
	heur	16.2	14.1	<b>0.219</b>	<b>1.329</b>	3.2	69.9	0.0
	anno	<b>16.6</b>	<b>15.9</b>	0.349	1.404	<b>38.5</b>	<b>94.4</b>	0.0
R3DS-syn	none	14.0	<b>8.4</b>	0.390	1.664	-	-	-
	heur	<b>14.6</b>	7.8	<b>0.281</b>	1.301	4.6	82.9	<b>52.6</b>
	anno	14.3	8.2	0.349	<b>1.219</b>	<b>32.0</b>	<b>95.0</b>	0.0
R3DS-mix	none	17.6	15.8	0.171	<b>1.007</b>	-	-	-
	heur	<b>19.2</b>	17.7	<b>0.151</b>	1.267	3.6	83.7	0.0
	anno	18.6	<b>18.2</b>	0.158	1.308	<b>12.0</b>	<b>96.2</b>	<b>85.8</b>

**Table 6: Comparison of the average number of bounding box and mesh-based object collisions per scene in IG and R3DS.** R3DS exhibits more bounding box-based collisions, but almost none of these are actual physical collisions between object meshes. Measuring collisions between bounding boxes is a poor collision measure for fully-populated, real-world scenes.

Datasets	Mesh Collisions	Box Collisions			
		obj-obj	obj-wall	obj-floor	obj-ceil
IG	-	1.185	0.075	0.000	0.790
R3DS	0.0006	2.823	0.388	0.035	0.064

R3DS, it is also possible to compute heuristic support relations from object bounding boxes. Specifically, an object is supported by another if their bounding boxes intersect within tolerance distance of 0.1m and the centroid of the former object is higher than that of the latter. Support by wall/floor/ceiling is calculated in the same way without the height judgment. This definition is similar to how DPC defines object attachment. Figure 7 compares these two ways of computing support relations, showing that heuristic computation can mistakenly designate support relations to two nearby objects. Table 5 shows that incorporating support relation prediction indeed influences the performance of DPC. Heuristic support information may worsen 3D object detection (mAP in *R3DS-real* and *R3DS-syn*), but it eliminates mesh collisions the most. Learning support relations from R3DS annotations leads to a 2.4 improvement on mAP in *R3DS-mix*, although the classification F1 score is low.

**Table 7: Ablation of relation optimization (RO) on R3DS-real.** The 2nd and 3rd row remove optimization terms in RO. The last row replaces bounding-box collisions with mesh collisions.

Rel. Opt.	3D mAP $\uparrow$	Mesh Collisions $\downarrow$	Box Collisions $\downarrow$	
			obj-obj	obj-arch
DPC	18.2	0.158	0.062	1.308
w/o obj col	18.9	1.342	1.130	1.301
w/o obj col+tch	19.6	1.219	1.062	1.295
w/ mesh col	19.7	1.027	0.856	1.394

**4) Is relation optimization (RO) effective on R3DS?** Test-time relation optimization (RO) was introduced by Zhang *et al.* [28] to reduce physical violations, floating objects, and misalignment between objects and architecture. The original cost function based on bounding box collisions succeeds in optimizing object poses, since there are few such collisions in the IG data originally used for evaluation (see in Tab. 6). However, the same data assumption does not hold for R3DS, which has more bounding-box-based collisions but nearly zero mesh-based collisions. We ablate the design of RO on R3DS-real in Table 7. By removing two optimization terms (bounding-box-based object-wise collision and touching step-by-step), the model outperforms the original one in 3D mAP (+1.4) but degrades in mesh-based and box-based collisions. We show that using mesh-based collision optimization leads to the best performance. The increase in collisions is unsurprising as the R3DS data reflects more cluttered real interiors.

**Limitations.** Our dataset construction relied on 3D architectures for each Matterport3D scan which are simplifications of the geometry of the real environment. One issue is imperfect wall positions, resulting in objects attached to these virtual walls being offset from the true surface. In addition, objects in our 3D scenes were placed without regard to the materials, meaning that the detailed surface appearance does not match that of the observed object. Future work can investigate transfer of surface appearance to the synthetic objects by projecting textures from the RGB-D data and 3D reconstructed meshes.

## 5 Conclusion

We introduced the R3DS dataset. R3DS provides more complete, densely populated, and richly annotated synthetic 3D scene proxies of real-world environments with linked panoramic images. We showed the usefulness of R3DS on the Panoramic Scene Understanding task. Our experiments demonstrate the value of realistic synthetic recreations in this task, in particular through the use of object support information. While we focused on the PanoSun task, R3DS can also be useful for other tasks such as single-view shape retrieval, single-view object pose estimation, and panoramic scene graph prediction.

## Acknowledgements

This work was funded in part by a CIFAR AI Chair, a Canada Research Chair, NSERC Discovery Grant, NSF award #2016532, and enabled by support from WestGrid and Compute Canada. Daniel Ritchie is an advisor to Geopipe and owns equity in the company. Geopipe is a start-up that is developing 3D technology to build immersive virtual copies of the real world with applications in various fields, including games and architecture. We thank Madhawa Vidanapathirana, Weijie Lin, and David Han for help with development of the annotation tool, and Denys Iliash, Mrinal Goshalia, Brandon Robles, Paul Brown, Chloe Ye, Coco Kaleel, Elizabeth Wu and Hannah Julius for data annotation, and Ivan Tam, Austin Wang, and Ning Wang for feedback on the paper draft.

## References

1. Avetisyan, A., Dahnert, M., Dai, A., Savva, M., Chang, A.X., Nießner, M.: Scan2CAD: Learning CAD model alignment in RGB-D scans. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
2. Batra, D., Chang, A.X., Chernova, S., Davison, A.J., Deng, J., Koltun, V., Levine, S., Malik, J., Mordatch, I., Mottaghi, R., Savva, M., Su, H.: Rearrangement: A challenge for embodied AI. arXiv preprint arXiv:2011.01975 (2020)
3. Chang, A., Dai, A., Funkhouser, T., Halber, M., Niebner, M., Savva, M., Song, S., Zeng, A., Zhang, Y.: Matterport3D: Learning from RGB-D data in indoor environments. In: Proceedings of the International Conference on 3D Vision (3DV). pp. 667–676. IEEE (2017)
4. Chang, A.X., Funkhouser, T., Guibas, L., Hanrahan, P., Huang, Q., Li, Z., Savarese, S., Savva, M., Song, S., Su, H., et al.: Shapenet: An information-rich 3d model repository. arXiv preprint arXiv:1512.03012 (2015)
5. Collins, J., Goel, S., Deng, K., Luthra, A., Xu, L., Gundogdu, E., Zhang, X., Vicente, T.F.Y., Dideriksen, T., Arora, H., et al.: Abo: Dataset and benchmarks for real-world 3d object understanding. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 21126–21136 (2022)
6. Dai, A., Chang, A.X., Savva, M., Halber, M., Funkhouser, T., Nießner, M.: Scannet: Richly-annotated 3d reconstructions of indoor scenes. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 5828–5839 (2017)
7. Dong, Y., Fang, C., Dong, Z., Bo, L., Tan, P.: PanoContext-Former: Panoramic total scene understanding with a transformer. arXiv preprint arXiv:2305.12497 (2023)
8. Fu, H., Cai, B., Gao, L., Zhang, L., Li, C., Xun, Z., Sun, C., Fei, Y., Zheng, Y., Li, Y., et al.: 3D-FRONT: 3D Furnished Rooms with layOuts and se-maNtics. arXiv preprint arXiv:2011.09127 (2020)

9. Fu, H., Jia, R., Gao, L., Gong, M., Zhao, B., Maybank, S., Tao, D.: 3D-Future: 3D Furniture shape with TextURE. arXiv preprint arXiv:2009.09633 (2020)
10. Hua, B.S., Pham, Q.H., Nguyen, D.T., Tran, M.K., Yu, L.F., Yeung, S.K.: SceneNN: A scene meshes dataset with annotations. In: Proceedings of the International Conference on 3D Vision (3DV). pp. 92–101. IEEE (2016)
11. Karras, T.: Maximizing parallelism in the construction of bvhs, octrees, and k-d trees. In: Proceedings of the Fourth ACM SIGGRAPH / Eurographics Conference on High-Performance Graphics. pp. 33–37. Eurographics Association (2012). <https://doi.org/10.2312/EGGH/HPG12/033-037>, <https://doi.org/10.2312/EGGH/HPG12/033-037>
12. Li, Z., Yu, T.W., Sang, S., Wang, S., Song, M., Liu, Y., Yeh, Y.Y., Zhu, R., Gundavarapu, N., Shi, J., Bi, S., Xu, Z., Yu, H.X., Sunkavalli, K., Hasan, M., Ramamoorthi, R., Chandraker, M.: OpenRooms: An end-to-end open framework for photorealistic indoor scene datasets. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2021)
13. Lim, J.J., Pirsiavash, H., Torralba, A.: Parsing ikea objects: Fine pose estimation. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV). pp. 2992–2999 (2013)
14. Maninis, K.K., Popov, S., Nießner, M., Ferrari, V.: CAD-estate: Large-scale CAD model annotation in RGB videos. arXiv preprint arXiv:2306.09011 (2023)
15. Paschalidou, D., Kar, A., Shugrina, M., Kreis, K., Geiger, A., Fidler, S.: Atiss: Autoregressive transformers for indoor scene synthesis. *Advances in Neural Information Processing Systems* **34**, 12013–12026 (2021)
16. Ramakrishnan, S.K., Gokaslan, A., Wijmans, E., Maksymets, O., Clegg, A., Turner, J., Undersander, E., Galuba, W., Westbury, A., Chang, A.X., et al.: Habitat-Matterport 3D dataset (hm3d): 1000 large-scale 3D environments for embodied AI. arXiv preprint arXiv:2109.08238 (2021)
17. Sadalgi, S.: Wayfair’s 3D Model API. <https://www.aboutwayfair.com/tech-innovation/wayfairs-3d-model-api> (2016), [Online; accessed 15-Nov-2023]
18. Shen, B., Xia, F., Li, C., Martín-Martín, R., Fan, L., Wang, G., Buch, S., D’Arpino, C., Srivastava, S., Tchapmi, L.P., Vainio, K., Fei-Fei, L., Savarese, S.: iGibson, a simulation environment for interactive tasks in large realistic scenes. In: Proceedings of the International Conference on Intelligent Robots and Systems (IROS) (2021)
19. Shen, B., Xia, F., Li, C., Martín-Martín, R., Fan, L., Wang, G., Pérez-D’Arpino, C., Buch, S., Srivastava, S., Tchapmi, L., et al.: iGibson 1.0: A simulation environment for interactive tasks in large realistic scenes. In: Proceedings of the International Conference on Intelligent Robots and Systems (IROS). pp. 7520–7527. IEEE (2021)
20. Straub, J., Whelan, T., Ma, L., Chen, Y., Wijmans, E., Green, S., Engel, J.J., Mur-Artal, R., Ren, C., Verma, S., et al.: The replica dataset: A digital replica of indoor spaces. arXiv preprint arXiv:1906.05797 (2019)

21. Sun, C., Hsiao, C.W., Sun, M., Chen, H.T.: Horizonnet: Learning room layout with 1d representation and pano stretch data augmentation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 1047–1056 (2019)
22. Sun, X., Wu, J., Zhang, X., Zhang, Z., Zhang, C., Xue, T., Tenenbaum, J.B., Freeman, W.T.: Pix3D: Dataset and methods for single-image 3D shape modeling. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 2974–2983 (2018)
23. Szot, A., Clegg, A., Undersander, E., Wijmans, E., Zhao, Y., Turner, J., Maestre, N., Mukadam, M., Chaplot, D.S., Maksymets, O., et al.: Habitat 2.0: Training home assistants to rearrange their habitat. *Advances in Neural Information Processing Systems* **34**, 251–266 (2021)
24. Tzionas, D., Ballan, L., Srikantha, A., Aponte, P., Pollefeys, M., Gall, J.: Capturing hands in action using discriminative salient points and physics simulation. *International Journal of Computer Vision (IJCV)* **118**(2), 172–193 (Jun 2016), <https://doi.org/10.1007/s11263-016-0895-4>
25. Wang, K., Lin, Y.A., Weissmann, B., Savva, M., Chang, A.X., Ritchie, D.: Planit: Planning and instantiating indoor scenes with relation graph and spatial prior networks. *ACM Transactions on Graphics (TOG)* **38**(4), 1–15 (2019)
26. Xiang, Y., Kim, W., Chen, W., Ji, J., Choy, C., Su, H., Mottaghi, R., Guibas, L., Savarese, S.: Objectnet3D: A large scale database for 3D object recognition. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. pp. 160–176. Springer (2016)
27. Yadav, K., Ramrakhya, R., Ramakrishnan, S.K., Gervet, T., Turner, J., Gokaslan, A., Maestre, N., Chang, A.X., Batra, D., Savva, M., et al.: Habitat-Matterport 3D semantics dataset. *arXiv preprint arXiv:2210.05633* (2022)
28. Zhang, C., Cui, Z., Chen, C., Liu, S., Zeng, B., Bao, H., Zhang, Y.: Deep-PanoContext: Panoramic 3D scene understanding with holistic scene context graph and relation-based optimization. In: *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*. pp. 12632–12641 (2021)
29. Zhang, Y., Song, S., Tan, P., Xiao, J.: PanoContext: A whole-room 3D context model for panoramic scene understanding. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. pp. 668–686. Springer (2014)
30. Zheng, J., Zhang, J., Li, J., Tang, R., Gao, S., Zhou, Z.: Structured3D: A large photo-realistic dataset for structured 3D modeling. *arXiv preprint arXiv:1908.00222* (2019)
31. Zhou, T., Tucker, R., Flynn, J., Fyffe, G., Snavely, N.: Stereo magnification: learning view synthesis using multiplane images. *ACM Transactions on Graphics (TOG)* **37**(4), 1–12 (2018)