

# Hypersim: A Photorealistic Synthetic Dataset for Holistic Indoor Scene Understanding

Mike Roberts Jason Ramapuram Anurag Ranjan Atulit Kumar  
Miguel Angel Bautista Nathan Paczan Russ Webb Joshua M. Susskind  
Apple

<http://github.com/apple/ml-hypersim>

## Abstract

For many fundamental scene understanding tasks, it is difficult or impossible to obtain per-pixel ground truth labels from real images. We address this challenge by introducing Hypersim, a photorealistic synthetic dataset for holistic indoor scene understanding. To create our dataset, we leverage a large repository of synthetic scenes created by professional artists, and we generate 77,400 images of 461 indoor scenes with detailed per-pixel labels and corresponding ground truth geometry. Our dataset: (1) relies exclusively on publicly available 3D assets; (2) includes complete scene geometry, material information, and lighting information for every scene; (3) includes dense per-pixel semantic instance segmentations and complete camera information for every image; and (4) factors every image into diffuse reflectance, diffuse illumination, and a non-diffuse residual term that captures view-dependent lighting effects.

We analyze our dataset at the level of scenes, objects, and pixels, and we analyze costs in terms of money, computation time, and annotation effort. Remarkably, we find that it is possible to generate our entire dataset from scratch, for roughly half the cost of training a popular open-source natural language processing model. We also evaluate *sim-to-real transfer performance* on two real-world scene understanding tasks – semantic segmentation and 3D shape prediction – where we find that pre-training on our dataset significantly improves performance on both tasks, and achieves state-of-the-art performance on the most challenging Pix3D test set. All of our rendered image data, as well as all the code we used to generate our dataset and perform our experiments, is available online.

## 1. Introduction

For many fundamental scene understanding tasks, it is difficult or impossible to obtain per-pixel ground truth labels from real images. In response to this challenge, the computer vision community has developed several photore-



Figure 1. Overview of the Hypersim dataset. For each color image (a), Hypersim includes the following ground truth layers: depth (b); surface normals (c); instance-level semantic segmentations (d,e); diffuse reflectance (f); diffuse illumination (g); and a non-diffuse residual image that captures view-dependent lighting effects like glossy surfaces and specular highlights (h). Our diffuse reflectance, diffuse illumination, and non-diffuse residual layers are stored as HDR images, and can be composited together to exactly reconstruct the color image.

alistic synthetic datasets and interactive simulation environments that have spurred rapid progress towards the goal of holistic indoor scene understanding [6, 7, 9, 10, 14, 15, 19, 22, 24, 31, 33, 36, 37, 39, 43, 44, 45, 46, 49, 52, 61, 62, 63, 65, 70, 72, 75, 76, 79, 83, 84].

Dataset/simulator	Images	3D	Seg.	Intrinsic
<i>Real (3D reconstruction)</i>				
SceneNN [33]		✓	S+I	
Stanford 2D-3D-S [9, 10, 79]	✓	✓	S+I	
Matterport3D [7, 15, 61, 62, 79]	✓	✓	S+I	
ScanNet [19]	✓	✓	S+I	
Gibson [62, 79]		✓		
Replica [62, 72]		✓	S+I	
<i>Synthetic (artist-created)</i>				
AI2-THOR [39]		✓	S+I	
ARAP [14]	✓	✓	I	D+R
SceneNet-RGBD [52]	✓	✓	S+I	
PBRs [70, 83]			S+I	
CGIntrinsics [45, 70]	✓			D
InteriorNet [43]	✓		S+I	D
Jiang et al. [36]			S+I	
RobotriX [24]	✓	✓	S+I	
CG-PBR [63, 70]			S	D+R
DeepFurniture [49]	✓		S	D
Structured3D [84]	✓		S+I	D
TartanAir [65, 76]	✓		I	
Li et al. [44, 70]				D+R
3D-FUTURE [22]	✓	✓	S+I	
OpenRooms [46]	✓	✓		D+R
<b>Hypersim (ours)</b>	✓	✓	<b>S+I</b>	<b>D+R</b>

Table 1. Comparison to previous datasets and simulators for indoor scene understanding. We broadly categorize these datasets and simulators as being either *real* (i.e., based on 3D triangle mesh reconstructions from real sensors) or *synthetic* (i.e., artist-created), and we sort chronologically within each category. We limit our comparisons to synthetic datasets and simulators that aim to be photorealistic. The *Images* and *3D* columns indicate whether or not images and 3D assets (e.g., triangle meshes) are publicly available. The *Seg.* column indicates what type of segmentation information is available: S indicates semantic; I indicates instance. The *Intrinsic* column indicates how images are factored into disentangled lighting and shading components: D indicates that each image is factored into diffuse reflectance and diffuse illumination; D+R indicates that each factorization additionally includes a non-diffuse residual term that captures view-dependent lighting effects. Our dataset is the first to include images, 3D assets, semantic instance segmentations, and a disentangled image representation.

However, existing synthetic datasets and simulators have important limitations (see Table 1). First, most synthetic datasets are derived from 3D assets that are not publicly available. These datasets typically include rendered images, but do not include the underlying 3D assets used during rendering (e.g., triangle meshes), and are therefore not suitable for geometric learning problems that require direct 3D supervision (e.g., [27]). Second, not all synthetic datasets and simulators include semantic segmentations. Although it is common for synthetic datasets to include some kind of segmentation information, these segmentations may not include semantic labels, and may group pixels together at the granularity of low-level object parts, rather than semantically meaningful objects. Third, most datasets and simulators do not factor images into disentangled lighting

and shading components, and are therefore not suitable for inverse rendering problems (e.g., [13, 40]). No existing synthetic dataset or simulator addresses all of these limitations, including those that target outdoor scene understanding [8, 20, 23, 34, 38, 41, 55, 56, 58, 60, 65, 76, 77].

In this work, we introduce Hypersim, a photorealistic synthetic dataset for holistic indoor scene understanding that addresses all of the limitations described above (see Figure 1). To create our dataset, we leverage a large repository of synthetic scenes created by professional artists, and we generate 77,400 images of 461 indoor scenes with detailed per-pixel labels and corresponding ground truth geometry. Our dataset: (1) relies exclusively on publicly available 3D assets; (2) includes complete scene geometry, material information, and lighting information for every scene; (3) includes dense per-pixel semantic instance segmentations and complete camera information for every image; and (4) factors every image into diffuse reflectance, diffuse illumination, and a non-diffuse residual term that captures view-dependent lighting effects. Together, these features make our dataset well-suited for geometric learning problems that require direct 3D supervision (e.g., [27]), multi-task learning problems that require reasoning jointly over multiple input and output modalities (e.g., [71]), and inverse rendering problems (e.g., [13, 40]).

To generate our dataset, we introduce a novel computational pipeline that takes as input a collection of scenes downloaded from an online marketplace, and produces as output a collection of images with ground truth labels and corresponding geometry (see Figure 2). Our pipeline has three main steps. First, we generate camera views of each input scene using a novel view sampling heuristic that does not require the scene to be semantically labeled. Second, we generate images using a cloud rendering system that we built on top of publicly available cloud computing services. Third, we obtain semantic segmentations from a human annotator using an interactive mesh annotation tool we built ourselves.

We analyze our dataset at the level of scenes, objects, and pixels, and we analyze costs in terms of money, computation time, and annotation effort. Remarkably, we find that it is possible to generate our entire dataset from scratch, for roughly half the cost of training a popular open-source natural language processing model. We also evaluate sim-to-real transfer performance on two scene understanding tasks – semantic segmentation on NYUv2 [68] and 3D shape prediction on Pix3D [73] – where we find that pre-training on our dataset significantly improves performance on both tasks, and achieves state-of-the-art performance on the most challenging Pix3D test set. All of our rendered image data, as well as all the code we used to generate our dataset and perform our experiments, is available online.<sup>1</sup>

<sup>1</sup><http://github.com/apple/ml-hypersim>




Figure 2. Overview of our computational pipeline. In this simplified diagram, our pipeline takes as input a triangle mesh, an artist-defined camera pose, and a V-Ray scene description file, and produces as output a collection of images with ground truth labels and corresponding geometry. The main steps of our pipeline are as follows. We estimate the free space in our scene, use this estimate to generate a collision-free camera trajectory, modify our V-Ray scene to include the trajectory, and invoke our cloud rendering system to render images. In parallel with the rest of our pipeline, we annotate the scene’s triangle mesh using our interactive tool. In a post-processing step, we propagate mesh annotations to our rendered images (not shown). This pipeline design enables us to render images before mesh annotation is complete, and also enables us to re-annotate our scenes (e.g., with a different set of labels) without needing to re-render images.

## 2. Related Work

**Synthetic Data in Computer Vision** Synthetic data plays a critical role in a wide variety of computer vision applications. See the excellent recent survey by Nikolenko [54].

**Synthetic Data for Indoor Scene Understanding** We discuss photorealistic datasets and simulation environments for indoor scene understanding in Section 1. Non-photorealistic datasets and environments [29, 30, 70, 78] also play an important role in scene understanding research because they can be rendered very efficiently. However, these datasets and environments introduce a large domain gap between real and synthetic images that must be handled carefully [54]. In contrast, our dataset aims to be as photo-realistic as possible, thereby significantly reducing this domain gap.

Several datasets provide 3D CAD models that have been aligned to individual objects in real images [19, 47, 73, 80, 81]. In these datasets, the CAD models may not be exactly aligned to each image, and many objects that are visible in an image do not have a corresponding CAD model. In contrast, our dataset provides segmented 3D models that are exactly aligned to each image, and every pixel of every image is associated with a 3D model.

Shi et al. [66] provide photorealistic renderings of individual objects in the ShapeNet dataset [16], where each image is factored into diffuse reflectance, diffuse illumination, and a non-diffuse residual term. Our images are factored in the same way, but we render entire scenes, rather than individual objects.

**View Sampling Methods for Synthetic Scenes** Existing view sampling methods for synthetic scenes include: sampling views uniformly at random [76, 84]; using data-driven methods to generate realistic camera jitter [43]; and using semantic segmentation information to match the distribution of semantic classes to an existing dataset [26], prefer foreground semantic classes [29, 30, 70], and maintain a minimum number of semantic instances in each view

[52, 83]. However, existing methods are not directly applicable in our setting. First, many of our scenes are highly *staged*, i.e., they contain realistic clutter in some parts of the scene, but are unrealistically empty in other parts. As a result, for our scenes, sampling views uniformly at random produces many uninformative views with no foreground objects. Second, our pipeline supports rendering images and annotating scenes in parallel, and therefore we do not have access to semantic segmentation information when we are sampling views. Our view sampling method addresses these challenges by choosing salient views without relying on segmentation information, and is therefore directly applicable in our setting.

**Interactive Annotation Tools for 3D Scenes** Our interactive tool is similar in spirit to existing tools for annotating reconstructed triangle meshes [19, 53, 72] and synthetic images from video games [8, 55]. All of these tools, including ours, leverage some kind of pre-segmentation strategy to reduce annotation effort. Tools for annotating reconstructed meshes [19, 53, 72] obtain pre-segmentations by applying an unsupervised grouping method (e.g., [21]) to the input mesh. As a result, the quality of 2D semantic labels obtained from these tools is limited by the quality of the reconstructed input mesh, as well as the unsupervised grouping. On the other hand, tools for annotating video game images [8, 55] obtain clean pre-segmentations by analyzing per-pixel rendering metadata, but do not allow freeform 3D navigation through the scene. In contrast, our tool leverages clean artist-defined mesh pre-segmentations, leading to 2D semantic labels that are exactly aligned with our rendered images, and our tool allows freeform 3D navigation.

## 3. Data Acquisition Methodology

To assemble an appropriate collection of scenes for our dataset, we browsed through online marketplaces looking for ready-made indoor scenes that satisfied three main desiderata. First, we wanted as many scenes as possible.



Figure 3. Randomly selected images from our dataset. From these images, we see that the scenes in our dataset are visually diverse, and our view sampling heuristic generates informative views without requiring our scenes to be semantically labeled.

Second, we wanted scenes that are as photorealistic and visually diverse as possible. Because our goal in this work is to construct a static dataset, rather than an interactive environment, we are willing to sacrifice rendering speed to achieve larger scale and greater photorealism. Third, we wanted scenes that are as consistent as possible in terms of their file formats and internal data representations. This last criterion arises because it is easier to implement our automated computational pipeline if we can parse, modify, and render our scenes in a consistent way.

Motivated by these desiderata, we chose the Evermotion Archinteriors Collection [3] as our starting point<sup>2</sup>. This collection consists of over 500 photorealistic indoor scenes, and has several features that are especially helpful for our purposes. First, each scene is represented as a standalone asset file that is compatible with V-Ray [2]. This representation is helpful because V-Ray has a powerful Python API for programmatically manipulating scenes. Second, each scene is scaled appropriately in metric units, and has a consistent up direction. Third, each scene is grouped into object parts, and includes a small set of artist-defined camera poses (i.e., typically between 1 and 5) that frame the scene in an aesthetically pleasing way. In our pipeline, we use the grouping into object parts as a pre-segmentation that reduces annotation effort, and we use the artist-defined camera poses as a form of weak supervision in several of our processing steps. Fourth, almost every scene is distributed under a permissive license that allows publicly releasing rendered images, e.g., in academic publications and public benchmarks.

We excluded scenes from our dataset according to the following criteria. First, we excluded any scenes that depict isolated objects, rather than complete environments. Second, we excluded any scenes that are not distributed under a royalty-free license. Third, for each scene, we manually rendered test images, attempted to export mesh data from the standalone asset file, and attempted to generate camera trajectories through the scene using our view sampling heuristic. We excluded any scenes with noticeable render-

<sup>2</sup>We purchased the Evermotion Archinteriors Collection from TurboSquid [4].

ing artifacts, any scenes where we were unable to successfully export mesh data, and any scenes where our view sampling heuristic failed to generate views inside the scene’s intended viewing region. After applying these exclusion criteria, we were left with 461 scenes (568 scenes in the Evermotion Archinteriors collection, 107 scenes excluded, 21 of these were because our view sampling heuristic failed). In our public code release, we provide a complete list of every scene in our dataset.

## 4. Computational Pipeline

After acquiring our scenes, we apply our computational pipeline to generate images with ground truth labels and corresponding geometry (see Figure 2). In this section, we describe our computational pipeline, and we assume for simplicity that we are processing a single scene. To generate our full dataset, we apply the same pipeline to all of our scenes. For any pipeline steps that require manual data filtering, we provide a complete record of our filtering decisions in our public code release, so our dataset can be reproduced exactly. In the supplementary material, we describe our procedure for estimating free space, our procedure for modifying V-Ray scenes, and our cloud rendering system.

**Pre-processing** We assume that we are given as input a standalone asset file that describes our scene. We begin our pipeline by programmatically exporting a triangle mesh, all artist-defined camera poses, and a V-Ray scene description file from the original asset file. We manually remove from our dataset any artist-defined camera poses that are outside the intended viewing region for the scene.

*Output.* In total, we exported 784 artist-defined camera poses (809 exported initially, 25 removed manually).

**Generating Camera Trajectories** We generate camera trajectories using a simple view sampling heuristic that does not require the input scene to be semantically labeled, works well for our scenes, and has not previously appeared in the literature. When designing our heuristic, we make the observation that salient objects (e.g., chairs, couches, lamps) tend to be more finely tessellated than non-salient objects (e.g., walls, floors, ceilings) (see Figure 5). Motivated by




Figure 4. Our interactive mesh annotation tool. Our tool has a *semantic instance view* (a,b,c) and a *semantic label view* (d,e), as well as a set of *selection filters* that can be used to limit the extent of editing operations based on the current state of the mesh. To see how these filters can be useful, consider the following scenario. The table in this scene is composed of multiple object parts, but initially, these object parts have not been grouped into a semantic instance (a). Our filters enable the user to paint the entire table by drawing a single rectangle, without disturbing the walls, floor, or other objects (b,c). Once the table has been grouped into an instance, the user can then apply a semantic label with a single button click (d,e). Parts of the mesh that have not been painted in either view are colored white (e.g., the leftmost chair). Parts of the mesh that have not been painted in the current view, but have been painted in the other view, are colored dark grey, (e.g., the table in (d)). Our tool enables the user to accurately annotate an input mesh with very rough painting gestures.



Figure 5. Color and wireframe renderings of a typical scene. In the wireframe rendering, we observe that the salient objects (e.g., chair, couch, lamp) are more finely tessellated than the non-salient objects (e.g., walls, floor, ceiling). This observation motivates our view sampling heuristic that takes triangle density into account, and does not require an input scene to be semantically labeled.




Figure 6. We include a tight 9-DOF bounding box for each semantic instance, so that our dataset can be applied directly to 3D object detection problems (e.g., [69]).

this observation, we define a *view saliency model* that takes triangle density into account, and we sample views based on this model. In our model, we also include a term that penalizes views for observing empty pixels, i.e., pixels that do not contain any scene geometry.

Stating our approach formally, we define the view saliency  $v(\mathbf{c})$  of the camera pose  $\mathbf{c}$  as follows,

$$v(\mathbf{c}) = t^\alpha p^\beta \quad (1)$$

where  $t$  is the number of unique triangles observed by  $\mathbf{c}$ ;  $p$  is the fraction of non-empty pixels observed by  $\mathbf{c}$ ; and  $\alpha, \beta > 0$  are parameters that control the sensitivity of our model to triangle counts and empty pixels, respectively.

Using our view saliency model, we generate camera trajectories by constructing random walks through free space that begin at each artist-defined camera pose, and are biased towards upright salient views. In Figure 3, we show randomly selected images from our dataset that were generated according to this sampling procedure. We include our exact random walk formulation in the supplementary material. Our trajectories can occasionally drift outside the scene’s intended viewing region, e.g., through an open window. To address this issue, we manually remove any such trajectories from our dataset.

*Output.* In total, we generated 774 camera trajectories using our random walk sampling approach (784 generated initially, 10 removed manually). In order to achieve reasonable visual coverage of our scenes while maintaining tolerable rendering costs, we defined each trajectory to consist of 100 camera poses, ultimately leading to a total of 77,400 distinct views that must be rendered.

**Interactive Mesh Annotation** In parallel with the rest of our pipeline, we obtain semantic segmentations using an interactive mesh annotation tool we built ourselves (see Figure 4). In addition to providing a mesh painting interface that operates at the granularity of object parts [8, 19, 53, 55, 72], our tool has two unique features that are especially helpful for annotating our scenes.

First, our tool provides a set of *selection filters* that can be used to limit the extent of editing operations based on the current state of the mesh. These filters enable the user to accurately annotate the mesh using very rough painting gestures (see Figure 4).

Second, many of our meshes are large (i.e., millions of triangles) and incoherently laid out in memory. These meshes cannot be rendered at interactive rates, even on modern GPUs, without the use of advanced acceleration structures [5, 35, 50]. In our tool, we address this issue by including an option to *decimate* the input mesh, which simply removes triangles at random until the mesh fits within a user-specified triangle budget. We bias our decimation procedure to remove very small triangles, and optionally trian-

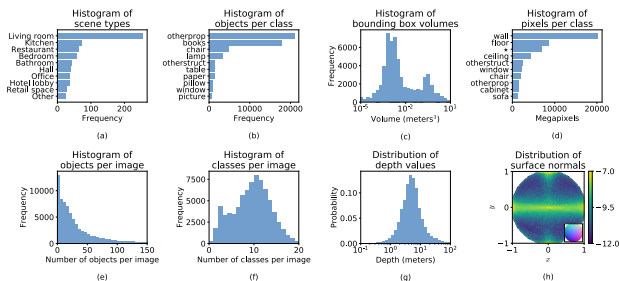


Figure 7. Dataset statistics at the granularity of scenes (a), objects (b,c), and pixels (d,e,f,g,h). We truncate the histograms in (a,b,d) to show the 10 most common scene types and object classes. In (a), we assign a scene type to each camera trajectory, and we count the number of trajectories belonging to each scene type. In (b), we count the number of unique objects in each scene that are visible in at least one image. In (d),  $\star$  indicates pixels with no class label. In (g), we compute the Euclidean distance from the surface at each pixel to the optical center of the camera. In (h), we show the distribution of camera-space normals as a function of the normal’s  $x$  and  $y$  coordinates (i.e., how much the normal is pointing to the right and up in camera space), where color indicates log-probability, and the small inset indicates our mapping from normal values to RGB values (e.g., in Figure 1c). The horizontal axes in (c) and (g) are log-scaled.




Figure 8. Distributions of hue-saturation values in our disentangled image representation. We normalize each RGB value independently, convert it to HSV space, and show the resulting hue-saturation distributions. Color indicates log-probability.




Figure 9. Distributions of brightness values in our disentangled image representation. We store our image data in an unnormalized HDR format, so we observe brightness values outside the range  $[0,1]$ . The horizontal axes in these plots are log-scaled.




Figure 10. Costs of generating our dataset in terms of money (a), computation time (b), and annotation effort (c).

gles that are very far away from the camera, since removing these triangles tends to have a negligible impact on the user. Our tool guarantees that annotations on the decimated mesh are implicitly propagated back to the full-resolution mesh.

**Output.** Using our tool, we annotated our entire dataset of 461 scenes with instance-level NYU40 labels [28, 68].

**Post-processing** In a final post-processing step, we propagate our mesh annotations to our images using per-pixel metadata that we generate during rendering. See the supplementary material for details. We also generate a tight 9-DOF bounding box for each semantic instance [11, 51], so that our dataset can be applied directly to 3D object detection problems (e.g., [69]) (see Figure 6).

## 5. Analysis

In this section, we analyze our dataset at the level of scenes, objects, and pixels, and we analyze costs in terms of money, computation time, and annotation effort. We summarize the contents of our dataset in Figures 7, 8, 9, and we summarize costs in Figure 10.

**Scenes** Our dataset consists mostly of residential scenes (e.g., living rooms, kitchens, bedrooms, bathrooms), but commercial scenes (e.g., offices, restaurants) are also common (Fig. 7a). Our scenes are highly cluttered, containing 127.3 objects per scene on average. This level of clutter and annotation detail compares favorably to existing indoor scene datasets (e.g., ScanNet [19]: 14 objects per scene; Replica [72]: 84 objects per scene). Our scenes range in size from  $<0.1$  up to 11 million triangles.

**Objects** At the object level, the most common semantic classes in our dataset are  $\{otherprop, books, chair\}$  (Fig. 7b). The prevalence of these classes is roughly consistent with existing indoor scene datasets [19, 69, 72]. The distribution of bounding box volumes for our objects is bimodal, where the two modes correspond to the bounding box volumes of a coffee mug and an office chair (Fig. 7c).

**Segmentation Images** At the pixel level, the most common semantic classes in our dataset are  $\{wall, floor, ceiling\}$  (Fig. 7d), which tend to dominate even in cluttered images (Fig. 1e). 88.3% of pixels in our dataset have a semantic class label, and 52.0% of pixels have a semantic instance ID. This annotation density is lower than GTA5 [55] (98.3%), but higher than ScanNet [19] (76%).

The images in our dataset contain 8.9 classes on average, 49.9 objects on average, and 51.5% of our images contain 21 or more objects (Fig. 7e,f). This level of clutter and annotation detail compares favorably to existing semantic segmentation datasets (e.g., NYUv2 [68]: 23.5 objects per image; COCO [48]: 7.2 objects per image; ADE20K [85]: 19.6 objects per image), and offers quantitative evidence that our view sampling heuristic is successfully generating informative views of our scenes.

Training procedure		mIoU (13-class)	mIoU (40-class)
Pre-train	Fine-tune		
None	NYUv2 (100%)	45.2	31.4
Hypersim (ours)	NYUv2 (25%)	46.4	29.0
	NYUv2 (50%)	49.1	32.7
	NYUv2 (100%)	<b>51.6</b>	<b>36.4</b>

Table 2. Sim-to-real performance of our dataset on NYUv2 [68] semantic segmentation. Higher is better. In parentheses, we show the amount of NYUv2 training data used during training. For 13-class segmentation, pre-training on our dataset and fine-tuning on 25% of the NYUv2 training set outperforms training on the full NYUv2 training set. For 40-class segmentation, pre-training on our dataset and fine-tuning on 50% of the NYUv2 training set outperforms training on the full NYUv2 training set.

**Depth Images** Our depth values are log-normally distributed with an average depth of 5.4 meters (Fig. 7g). This distribution is roughly consistent with existing indoor depth datasets [68, 74, 82], but is noticeably different from outdoor datasets [17, 25].

**Surface Normal Images** Our distribution of surface normals is biased towards planar surfaces, e.g., walls, floors, and ceilings viewed from an upright camera (Fig. 7g). This distribution disagrees with the *surface isotropy prior* used in the popular SIRFS model [12] for decomposing images into disentangled shape, lighting, and shading components. The authors of SIRFS note that surface isotropy is an appropriate prior for object-centric tasks, but is less appropriate for scene-centric tasks, and our data supports this assertion.

**Disentangled Lighting and Shading Images** We observe that some diffuse reflectance values are more likely than others (e.g., desaturated oranges are more likely than saturated greens), and that our distribution of diffuse reflectance is *speckled*, indicating a sparse palette of especially likely values (Fig. 8). Our distribution of diffuse illumination is especially biased towards natural lighting conditions (Fig. 8). These distributions roughly agree with common priors on reflectance and illumination found in the inverse rendering literature [12, 14].

The distributions of hue-saturation values in our non-diffuse residual images and final color images are similar (Fig. 8). (If our scenes consisted of half perfectly diffuse and half perfectly specular surfaces, we would expect these two distributions to be identical.) However, our residual images are much more sparse, i.e., close to zero brightness most of the time (Fig. 1h, Fig. 9). Nonetheless, we observe that our residual images contribute a non-trivial amount of energy to our final images, and this observation validates our decision to represent non-diffuse illumination explicitly in our dataset.

**Rendering Costs** In total, it cost \$57K to generate our dataset (\$6K to purchase 461 scenes, \$51K to render 77,400

Training procedure				
Pre-train	Fine-tune	AP <sup>mesh</sup>	AP <sup>mask</sup>	AP <sup>box</sup>
None	Pix3D	28.8	63.9	72.2
Hypersim (ours)	Pix3D	<b>29.6</b>	<b>64.6</b>	<b>72.7</b>

Table 3. Sim-to-real performance of our dataset on Pix3D [73] 3D shape prediction. Following [27], we report AP<sup>mesh</sup>, AP<sup>mask</sup>, and AP<sup>box</sup> on the  $\mathcal{S}_2$  test set. Higher is better. On the top row, we show the previous state-of-the-art, achieved by training Mesh R-CNN [27] on Pix3D. On the bottom row, we show our result, achieved by pre-training Mesh R-CNN on our dataset and fine-tuning on Pix3D. Pre-training on our dataset achieves state-of-the-art performance.

images), and took 231 vCPU years (2.4 years of wall-clock time on a large compute node). Although generating our entire dataset is undoubtedly expensive, it is considerably less expensive than some other popular learning tasks. For example, generating our dataset is 0.56 $\times$  as expensive as training the open-source Megatron-LM natural language processing model [67], which would cost \$103K to train from scratch using publicly available cloud computing services.

Per-image rendering costs in dollars are linearly related to compute times, and we observe that both are log-normally distributed (Fig. 10a,b). On average, rendering a single image in our dataset at 1024 $\times$ 768 resolution costs \$0.67 and takes 26 vCPU hours (20 minutes of wall-clock time on a large compute node). We include the cost of rendering each image in our public code release, so the marginal value and marginal cost of each image can be analyzed jointly in downstream applications [64].

**Annotation Costs** In total, annotating our scenes took 369 hours, leading to an average annotation speed of 39.8K pixels per second. Our annotation speed is two orders of magnitude faster than the fine-grained annotations in Cityscapes [17] (0.3K pixels per second), but an order of magnitude slower than GTA5 [55] (279.5K pixels per second). However, our annotations provide complementary information that isn’t available in either of these datasets. For example, our mesh annotations provide segmentation information for occluded parts of the scene that aren’t visible in any image, which could be used for amodal segmentation problems (e.g., [42]). Per-scene annotation times are log-normally distributed, with a single scene taking 48 minutes to annotate on average (Fig. 10c). All of our manual filtering decisions took less than 8 hours.

## 6. Experiments

In this section, we evaluate the sim-to-real transfer performance of our dataset on semantic segmentation and 3D shape prediction. For both tasks, our methodology is to pre-train on our dataset, fine-tune on an appropriate real-world dataset, and evaluate performance on the real-world dataset.

We summarize our results in Tables 2 and 3, and we include additional details in the supplementary material.

**Data Splits** In our public code release, we provide standard {training, validation, test} splits consisting of {59,543, 7,386, 7,690} images each. We split our dataset by scene, i.e., every image of a given scene belongs to the same split. We use these splits in all our experiments.

**Semantic Segmentation** We evaluate semantic segmentation performance on the NYUv2 dataset [68], which consists of 795 training images and 654 test images. We evaluate performance on the 13-class [18] and 40-class [28] segmentation tasks, and we report mean intersection-over-union (mIoU) as our evaluation metric. In all our experiments, we use RGB images for training and testing.

We use a standard U-Net [57] as our model architecture, with a ResNet-34 [32] encoder initialized using ImageNet [59] weights. During training, we always apply random {cropping, flipping, resizing}, and we apply color jittering with probability 0.5. We use an identical training recipe during pre-training and fine-tuning, and we perform our evaluation at  $512 \times 512$  resolution.

We find that pre-training on our dataset significantly improves semantic segmentation performance on NYUv2 (see Table 2). Our semantic segmentation results (+6.2 mIoU 13-class; +5.0 mIoU 40-class) are better than the reported results for PBRs [70, 83] (+1.6 mIoU 40-class), but worse than the reported results for SceneNet-RGBD [52] (+8.1 mIoU 13-class), under similar experimental conditions. However, our dataset (77K images, 0.5K scenes) is an order of magnitude smaller than PBRs (568K images, 45K scenes), and two orders of magnitude smaller than SceneNet-RGBD (5,000K images, 16K scenes). Despite its smaller size, we attribute the competitive performance of our dataset to the increased photorealism of our images and scenes. This finding suggests that there are multiple viable strategies for achieving good sim-to-real performance on a fixed rendering budget – a small, more photorealistic dataset can be competitive with a large, less photorealistic dataset. Determining the optimal *portfolio* of more and less photorealistic images for a given downstream task, subject to a fixed rendering budget, is an exciting direction for future work [64].

**3D Shape Prediction** We evaluate 3D shape prediction performance on the Pix3D dataset [73], which consists of 10,069 images and 395 unique triangle meshes. We use the data splits defined in [27] for training, validation, and testing, and we perform our final evaluation on the  $S_2$  test set, which is the most challenging Pix3D test set. Following [27], we report  $AP^{\text{mesh}}$ ,  $AP^{\text{mask}}$ , and  $AP^{\text{box}}$ .

We use Mesh R-CNN [27] as our model architecture. We follow the authors’ training recipe exactly, except we tune the learning rate used during pre-training. When selecting

our learning rate, we follow the authors’ guidelines [1] for choosing hyperparameters.

## 7. Conclusions

We leveraged a large repository of synthetic scenes to create a new dataset for holistic indoor scene understanding. We introduced a novel computational pipeline to generate salient views of our scenes and render photorealistic images in the cloud, as well as a new interactive tool to efficiently annotate our scenes. We used our pipeline and annotation tool to create the first computer vision dataset to combine images, 3D assets, semantic instance segmentations, and a disentangled image representation. We analyzed the costs of generating our dataset, and we found that it is possible to generate our entire dataset from scratch, for roughly half the cost of training a popular open-source natural language processing model. Finally, we demonstrated that pre-training on our dataset improves performance on two real-world scene understanding tasks, and achieves state-of-the-art performance on the most challenging Pix3D test set.

We believe our dataset can enable progress on a wide range of computer vision problems where obtaining real-world ground truth is difficult or impossible. In particular, our dataset is well-suited for geometric learning problems that require 3D supervision, multi-task learning problems, and inverse rendering problems. Our labeled scenes could be used to train automatic mesh segmentation systems, as well as generative modeling systems that synthesize realistic scene variations on-demand. Moving beyond our specific dataset, we see many potential applications for photorealistic synthetic data in computer vision, and we believe there are abundant opportunities to co-design rendering algorithms and learning algorithms to amortize rendering costs more effectively.

## Acknowledgements

We thank the professional artists at Evermotion for making their Archinteriors Collection available for purchase; Danny Nahmias for helping us to acquire data; Max Horton for helping us to prototype our annotation tool; Momchil Lukanov and Vlado Koylazov at Chaos Group for their excellent support with V-Ray; David Antler, Hanlin Goh, and Brady Quist for proofreading the paper; Ali Farhadi, Zhile Ren, Fred Schaffalitzky, and Qi Shan for the helpful discussions; and Jia Zheng for catching and correcting an error in Table 1.

## References

- [1] <http://github.com/facebookresearch/meshrcnn/issues/31>. 8
- [2] Chaos Group V-Ray. <http://www.chaosgroup.com>. 4



- [3] Evermotion Archinteriors Collection. <http://www.evermotion.org>. 4
- [4] TurboSquid. <http://www.turbosquid.com>. 4
- [5] Tomas Akenine-Möller, Eric Haines, Naty Hoffman, Angelo Pesce, Michał Iwanicki, , and Sébastien Hillaire. *Real-Time Rendering, Fourth Edition*. CRC Press, 2018. 5
- [6] Ibraheem Alhashim and Peter Wonka. High quality monocular depth estimation via transfer learning. arXiv 2018. 1
- [7] Peter Anderson, Qi Wu, Damien Teney, Jake Bruce, Mark Johnson, Niko Sünderhauf, Ian Reid, Stephen Gould, and Anton van den Hengel. Vision-and-language navigation: Interpreting visually-grounded navigation instructions in real environments. In *CVPR 2018*. 1, 2
- [8] Matt Angus, Mohamed ElBalkini, Samin Khan, Ali Harakeh, Oles Andrienko, Cody Reading, Steven Waslander, and Krzysztof Czarnecki. Unlimited road-scene synthetic annotation (URSA) dataset. arXiv 2018. 2, 3, 5
- [9] Iro Armeni, Alexander Sax, Amir R. Zamir, and Silvio Savarese. Joint 2D-3D-semantic data for indoor scene understanding. arXiv 2017. 1, 2
- [10] Iro Armeni, Ozan Sener, Amir R. Zamir, Helen Jiang, Ioannis Brilakis, Martin Fischer, and Silvio Savarese. 3D semantic parsing of large-scale indoor spaces. In *CVPR 2016*. 1, 2
- [11] Gill Barequet and Sarel Har-Peled. Efficiently approximating the minimum-volume bounding box of a point set in three dimensions. In *ACM-SIAM Symposium on Discrete Algorithms 2001*. 6
- [12] Jonathan T. Barron and Jitendra Malik. Shape, illumination, and reflectance from shading. *PAMI*, 2015. 7
- [13] Sean Bell, Kavita Bala, and Noah Snavely. Intrinsic images in the wild. In *SIGGRAPH 2014*. 2
- [14] Nicolas Bonneel, Balazs Kovacs, Sylvain Paris, and Kavita Bala. Intrinsic decompositions for image editing. *Computer Graphics Forum*, 36(2), 2017. 1, 2, 7
- [15] Angel Chang, Angela Dai, Thomas Funkhouser, Maciej Habber, Matthias Niessner, Manolis Savva, Shuran Song, Andy Zeng, and Yinda Zhang. Matterport3D: Learning from RGB-D data in indoor environments. In *3DV 2017*. 1, 2
- [16] Angel X. Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, Jianxiong Xiao, Li Yi, and Fisher Yu. ShapeNet: An information-rich 3D model repository. arXiv 2015. 3
- [17] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The Cityscapes dataset for semantic urban scene understanding. In *CVPR 2016*. 7
- [18] Camille Couprie, Clement Farabet, Laurent Najman, and Yann LeCun. Indoor semantic segmentation using depth information. In *ICLR 2013*. 8
- [19] Angela Dai, Angel X. Chang, Manolis Savva, Maciej Habber, Thomas Funkhouser, and Matthias Niessner. ScanNet: Richly-annotated 3D reconstructions of indoor scenes. In *CVPR 2017*. 1, 2, 3, 5, 6
- [20] Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. CARLA: An open urban driving simulator. In *CoRL 2017*. 2
- [21] Pedro F. Felzenszwalb and Daniel P. Huttenlocher. Efficient graph-based image segmentation. *IJCV*, 59(2), 2004. 3
- [22] Huan Fu, Rongfei Jia, Lin Gao, Mingming Gong, Binqiang Zhao, Steve Maybank, and Dacheng Tao. 3D-FUTURE: 3D furniture shape with TextURE. arXiv 2020. 1, 2
- [23] Adrien Gaidon, Qiao Wang, Yohann Cabon, and Eleonora Vig. Virtual worlds as proxy for multi-object tracking analysis. In *CVPR 2016*. 2
- [24] Alberto Garcia-Garcia, Pablo Martinez-Gonzalez, Sergiu Oprea, John Alejandro Castro-Vargas, Sergio Orts-Escolano, Jose Garcia-Rodriguez, and Alvaro Jover-Alvarez. The RobotriX: An extremely photorealistic and very-large-scale indoor dataset of sequences with robot trajectories and interactions. In *IROS 2018*. 1, 2
- [25] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The KITTI dataset. *IJRR*, 2013. 7
- [26] Kyle Genova, Manolis Savva, Angel X. Chang, and Thomas Funkhouser. Learning where to look: Data-driven viewpoint set selection for 3D scenes. arXiv 2017. 3
- [27] Georgia Gkioxari, Jitendra Malik, and Justin Johnson. Mesh R-CNN. In *ICCV 2019*. 2, 7, 8
- [28] Saurabh Gupta, Pablo Arbelaez, and Jitendra Malik. Perceptual organization and recognition of indoor scenes from RGB-D images. In *CVPR 2013*. 6, 8
- [29] Ankur Handa, Viorica Patraucean, Vijay Badrinarayanan, Simon Stent, and Roberto Cipolla. Understanding real world indoor scenes with synthetic data. In *CVPR 2016*. 3
- [30] Ankur Handa, Viorica Patraucean, Simon Stent, and Roberto Cipolla. SceneNet: An annotated model generator for indoor scene understanding. In *ICRA 2016*. 3
- [31] Ankur Handa, Thomas Whelan, John McDonald, and Andrew J. Davison. A benchmark for RGB-D visual odometry, 3D reconstruction and SLAM. In *ICRA 2014*. 1
- [32] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR 2016*. 8
- [33] Binh-Son Hua, Quang-Hieu Pham, Duc Thanh Nguyen, Minh-Khoi Tran, Lap-Fai Yu, and Sai-Kit Yeung. SceneNN: A scene meshes dataset with aNnotations. In *3DV 2016*. 1, 2
- [34] Braden Hurl, Krzysztof Czarnecki, and Steven Waslander. Precise synthetic image and LiDAR (PreSIL) dataset for autonomous vehicle perception. arXiv 2019. 2
- [35] Martin Isenburg and Peter Lindstrom. Streaming meshes. In *Visualization 2005*. 5
- [36] Chenfanfu Jiang, Siyuan Qi, Yixin Zhu, Siyuan Huang, Jenny Lin, Lap-Fai Yu, Demetri Terzopoulos, and Song-Chun Zhu. Configurable 3D scene synthesis and 2D image rendering with per-pixel ground truth using stochastic grammars. *IJCV*, 126(9), 2018. 1, 2
- [37] Lei Jin, Yanyu Xu, Jia Zheng, Junfei Zhang, Rui Tang, Shugong Xu, Jingyi Yu, and Shenghua Gao. Geometric structure based and regularized depth estimation from 360 degree indoor imagery. In *CVPR 2020*. 1

- [38] Samin Khan, Buu Phan, Rick Salay, and Krzysztof Czarnecki. ProcSy: Procedural synthetic dataset generation towards influence factor studies of semantic segmentation networks. In *CVPR 2019 Workshops*. 2
- [39] Eric Kolve, Roozbeh Mottaghi, Winson Han, Eli VanderBilt, Luca Weihs, Alvaro Herrasti, Daniel Gordon, Yuke Zhu, Abhinav Gupta, and Ali Farhadi. AI2-THOR: An interactive 3D environment for visual AI. arXiv 2017. 1, 2
- [40] Balazs Kovacs, Sean Bell, Noah Snavely, and Kavita Bala. Shading annotations in the wild. In *CVPR 2017*. 2
- [41] Philipp Krähenbühl. Free supervision from video games. In *CVPR 2018*. 2
- [42] Ke Li and Jitendra Malik. Amodal instance segmentation. In *ECCV 2016*. 7
- [43] Wenbin Li, Sajad Saeedi, John McCormac, Ronald Clark, Dimos Tzoumanikas, Qing Ye, Yuzhong Huang, Rui Tang, and Stefan Leutenegger. InteriorNet: Mega-scale multi-sensor photo-realistic indoor scenes dataset. In *BMVC 2018*. 1, 2, 3
- [44] Zhengqin Li, Mohammad Shafiei, Ravi Ramamoorthi, Kalyan Sunkavalli, and Manmohan Chandraker. Inverse rendering for complex indoor scenes: Shape, spatially-varying lighting and SVBRDF from a single image. In *CVPR 2020*. 1, 2
- [45] Zhengqi Li and Noah Snavely. CGIntrinsics: Better intrinsic image decomposition through physically-based rendering. In *ECCV 2018*. 1, 2
- [46] Zhengqin Li, Ting-Wei Yu, Shen Sang, Sarah Wang, Sai Bi, Zexiang Xu, Hong-Xing Yu, Kalyan Sunkavalli, Miloš Hašan, Ravi Ramamoorthi, and Manmohan Chandraker. OpenRooms: An end-to-end open framework for photorealistic indoor scene datasets. arXiv 2020. 1, 2
- [47] Joseph J. Lim, Hamed Pirsiavash, and Antonio Torralba. Parsing IKEA objects: Fine pose estimation. In *ICCV 2013*. 3
- [48] Tsung-Yi Lin, Michael Maire, Serge Belongie, Lubomir Bourdev, Ross Girshick, James Hays, Pietro Perona, Deva Ramanan, C. Lawrence Zitnick, and Piotr Dollár. Microsoft COCO: Common objects in context. In *ECCV 2014*. 6
- [49] Bingyuan Liu, Jiantao Zhang, Xiaoting Zhang, Wei Zhang, Chuanhui Yu, and Yuan Zhou. Furnishing your room by what you see: An end-to-end furniture set retrieval framework with rich annotated benchmark dataset. arXiv 2019. 1, 2
- [50] David Luebke, Martin Reddy, Jonathan D. Cohen, Amitabh Varshney, Benjamin Watson, and Robert Huebner. *Level of Detail for 3D Graphics*. Morgan Kaufmann, 2002. 5
- [51] Grégoire Malandain and Jean-Daniel Boissonnat. Computing the diameter of a point set. *International Journal of Computational Geometry & Applications*, 12(6), 2002. 6
- [52] John McCormac, Ankur Handa, Stefan Leutenegger, and Andrew J. Davison. SceneNet RGB-D: Can 5M synthetic images beat generic ImageNet pre-training on indoor segmentation? In *ICCV 2017*. 1, 2, 3, 8
- [53] Duc Thanh Nguyen, Binh-Son Hua, Lap-Fai Yu, and Sai-Kit Yeung. A robust 3D-2D interactive tool for scene segmentation and annotation. In *Pacific Graphics 2017*. 3, 5
- [54] Sergey I. Nikolenko. Synthetic data for deep learning. arXiv 2019. 3
- [55] Stephan Richter, Vibhav Vineet, Stefan Roth, and Vladlen Koltun. Playing for data: Ground truth from computer games. In *ECCV 2016*. 2, 3, 5, 6, 7
- [56] Stephan R. Richter, Zeeshan Hayder, and Vladlen Koltun. Playing for benchmarks. In *ICCV 2017*. 2
- [57] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional networks for biomedical image segmentation. In *MICCAI 2015*. 8
- [58] German Ros, Laura Sellart, Joanna Materzynska, David Vazquez, and Antonio M. Lopez. The SYNTHIA Dataset: A large collection of synthetic images for semantic segmentation of urban scenes. In *ECCV 2016*. 2
- [59] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet large scale visual recognition challenge. *IJCV*, 2015. 8
- [60] Fatemeh Sadat Saleh, Mohammad Sadegh Aliakbarian, Mathieu Salzmann, Lars Petersson, and Jose M. Alvarez. Effective use of synthetic data for urban scene semantic segmentation. In *ECCV 2018*. 2
- [61] Manolis Savva, Angel X. Chang, Alexey Dosovitskiy, Thomas Funkhouser, and Vladlen Koltun. MINOS: Multimodal indoor simulator for navigation in complex environments. arXiv 2017. 1, 2
- [62] Manolis Savva, Abhishek Kadian, Oleksandr Maksymets, Yili Zhao, Erik Wijmans, Bhavana Jain, Julian Straub, Jia Liu, Vladlen Koltun, Jitendra Malik, Devi Parikh, and Dhruv Batra. Habitat: A platform for embodied AI research. In *ICCV 2019*. 1, 2
- [63] Soumyadip Sengupta, Jinwei Gu, Kihwan Kim, Guilin Liu, David W. Jacobs, and Jan Kautz. Neural inverse rendering of an indoor scene from a single image. In *ICCV 2019*. 1, 2
- [64] Burr Settles. *Active Learning*. Morgan & Claypool, 2012. 7, 8
- [65] Shital Shah, Debadepta Dey, Chris Lovett, and Ashish Kapoor. AirSim: High-fidelity visual and physical simulation for autonomous vehicles. In *Field and Service Robotics 2017*. 1, 2
- [66] Jian Shi, Yue Dong, Hao Su, and Stella X. Yu. Learning non-Lambertian object intrinsics across ShapeNet categories. In *CVPR 2017*. 3
- [67] Mohammad Shoeybi, Mostofa Patwary, Raul Puri, Patrick LeGresley, Jared Casper, and Bryan Catanzaro. Megatron-LM: Training multi-billion parameter language models using model parallelism. arXiv 2019. 7
- [68] Nathan Silberman, Pushmeet Kohli, Derek Hoiem, and Rob Fergus. Indoor segmentation and support inference from RGBD images. In *ECCV 2012*. 2, 6, 7, 8
- [69] Shuran Song, Samuel P. Lichtenberg, and Jianxiong Xiao. SUN RGB-D: A RGB-D scene understanding benchmark suite. In *CVPR 2015*. 5, 6
- [70] Shuran Song, Fisher Yu, Andy Zeng, Angel X. Chang, Manolis Savva, and Thomas Funkhouser. Semantic scene completion from a single depth image. In *CVPR 2017*. 1, 2, 3, 8

- [71] Trevor Standley, Amir R. Zamir, Dawn Chen, Leonidas Guibas, Jitendra Malik, and Silvio Savarese. Which tasks should be learned together in multi-task learning? *arXiv 2019*. [2](#)
- [72] Julian Straub, Thomas Whelan, Lingni Ma, Yufan Chen, Erik Wijmans, Simon Green, Jakob J. Engel, Raul Mur-Artal, Carl Ren, Shobhit Verma, Anton Clarkson, Mingfei Yan, Brian Budge, Yajie Yan, Xiqing Pan, June Yon, Yuyang Zou, Kimberly Leon, Nigel Carter, Jesus Briales, Tyler Gillingham, Elias Mueggler, Luis Pesqueira, Manolis Savva, Dhruv Batra, Hauke M. Strasdat, Renzo De Nardi, Michael Goesele, Steven Lovegrove, and Richard Newcombe. The Replica Dataset: A digital replica of indoor spaces. *arXiv 2019*. [1](#), [2](#), [3](#), [5](#), [6](#)
- [73] Xingyuan Sun, Jiajun Wu, Xiuming Zhang, Zhoutong Zhang, Chengkai Zhang, Tianfan Xue, Joshua B. Tenenbaum, and William T. Freeman. Pix3D: Dataset and methods for single-image 3D shape modeling. In *CVPR 2017*. [2](#), [3](#), [7](#), [8](#)
- [74] Igor Vasiljevic, Nick Kolkin, Shanyi Zhang, Ruotian Luo, Haochen Wang, Falcon Z. Dai, Andrea F. Daniele, Mohammadreza Mostajabi, Steven Basart, Matthew R. Walter, and Gregory Shakhnarovich. DIODE: A Dense Indoor and Outdoor DEpth Dataset. *arXiv 2019*. [7](#)
- [75] Qiang Wang, Shizhen Zheng, Qingsong Yan, Fei Deng, Kaiyong Zhao, and Xiaowen Chu. IRS: A large synthetic indoor robotics stereo dataset for disparity and surface normal estimation. *arXiv 2019*. [1](#)
- [76] Wenshan Wang, DeLong Zhu, Xiangwei Wang, Yaoyu Hu, Yuheng Qiu, Chen Wang, Yafei Hu, Ashish Kapoor, and Sebastian Scherer. TartanAir: A dataset to push the limits of visual SLAM. *arXiv 2020*. [1](#), [2](#), [3](#)
- [77] Magnus Wrenninge and Jonas Unger. Synscapes: A photorealistic synthetic dataset for street scene parsing. *arXiv 2018*. [2](#)
- [78] Yi Wu, Yuxin Wu, Georgia Gkioxari, and Yuandong Tian. Building generalizable agents with a realistic and rich 3D environment. *arXiv 2018*. [3](#)
- [79] Fei Xia, Amir R. Zamir, Zhi-Yang He, Alexander Sax, Jitendra Malik, and Silvio Savarese. Gibson Env: Real-world perception for embodied agents. In *CVPR 2018*. [1](#), [2](#)
- [80] Yu Xiang, Wonhui Kim, Wei Chen, Jingwei Ji, Christopher Choy, Hao Su, Roozbeh Mottaghi, Leonidas Guibas, and Silvio Savarese. A large scale database for 3D object recognition. In *ECCV 2016*. [3](#)
- [81] Yu Xiang, Roozbeh Mottaghi, and Silvio Savarese. Beyond PASCAL: A benchmark for 3D object detection in the wild. In *WACV 2014*. [3](#)
- [82] Jianxiong Xiao, Andrew Owens, and Antonio Torralba. SUN3D: A database of big spaces reconstructed using SfM and object labels. In *ICCV 2013*. [7](#)
- [83] Yinda Zhang, Shuran Song, Ersin Yumer, Manolis Savva, Joon-Young Lee, Hailin Jin, and Thomas Funkhouser. Physically-based rendering for indoor scene understanding using convolutional neural networks. In *CVPR 2017*. [1](#), [2](#), [3](#), [8](#)
- [84] Jia Zheng, Junfei Zhang, Jing Li, Rui Tang, Shenghua Gao, and Zihan Zhou. Structured3D: A large photo-realistic dataset for structured 3D modeling. *arXiv 2019*. [1](#), [2](#), [3](#)
- [85] Bolei Zhou, Hang Zhao, Xavier Puig, Tete Xiao, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Semantic understanding of scenes through ADE20K dataset. *IJCV*, 2019. [6](#)