

Floating No More: Object-Ground Reconstruction from a Single Image

Yunze Man¹, Yichen Sheng², Jianming Zhang³, Liang-Yan Gui¹, Yu-Xiong Wang¹

¹University of Illinois Urbana-Champaign ²Purdue University ³Adobe

{yunzem2, lgui, yxw}@illinois.edu, sheng30@purdue.edu, jianmzha@adobe.com

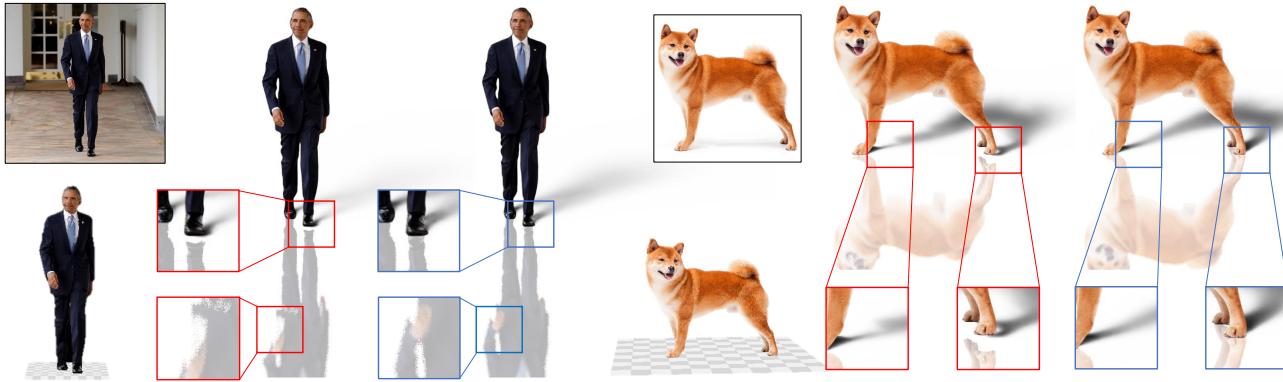


Figure 1. **ORG** (Object Reconstruction with Ground) simultaneously reconstructs the 3D object, the object-ground relationship, and the camera parameters from a monocular image. During shadow and reflection generation, prior methods have floating issues or unnatural shadows on the ground (red boxes). On the other hand, ORG achieves significantly more realistic editing and generation (blue boxes).

Abstract

Recent advancements in 3D object reconstruction from single images have primarily focused on improving the accuracy of object shapes. Yet, these techniques often fail to accurately capture the inter-relation between the object, ground, and camera. As a result, the reconstructed objects often appear floating or tilted when placed on flat surfaces. This limitation significantly affects 3D-aware image editing applications, like shadow rendering and object pose manipulation. To address this issue, we introduce **ORG** (Object Reconstruction with Ground), a novel task aimed at reconstructing 3D object geometry in conjunction with the ground surface. Our method uses two compact pixel-level representations to depict the relationship between camera, object, and ground. Experiments show that the proposed ORG model can effectively reconstruct object-ground geometry on unseen data, significantly enhancing the quality of shadow generation and pose manipulation compared to conventional single-image 3D reconstruction techniques. The project page can be found at [this website](#).

1. Introduction

The task of reconstructing an object in conjunction with a physically plausible ground, while not extensively explored,

is of significant importance. This is particularly relevant in the realm of image editing applications, where it influences key aspects like controllable shadow/reflection synthesis and object view manipulation. In this work, we aim at predicting an accurate and grounded representation of objects in 3D space from a single image, specifically under unrestricted camera conditions. Recent single-view approaches have demonstrated considerable promise in tackling object reconstruction [26, 34, 39, 57, 61]. However, due to the lack of integrated object-ground modeling, objects reconstructed using these methods often appear to be “floating” or tilted when placed on a flat surface, which greatly hinders the realistic rendering.

More specifically, recent works on *monocular depth estimation* [6, 7, 34, 61] has shown great performance. They aim to recover the 3D information of an object from a single-view image by directly estimating the pixel-level depth values. Their models have been trained on large-scale datasets, and thus can generalize well on in-the-wild images. However, as pointed out by [61], to project the depth map into 3D point clouds, additional camera parameters are needed. In some cases, off-the-shelf estimators can provide a rough estimate of these parameters, but this approach can limit the flexibility and effectiveness of object reconstruction in uncontrolled environments. Moreover, the unknown shift in the depth or disparity map will

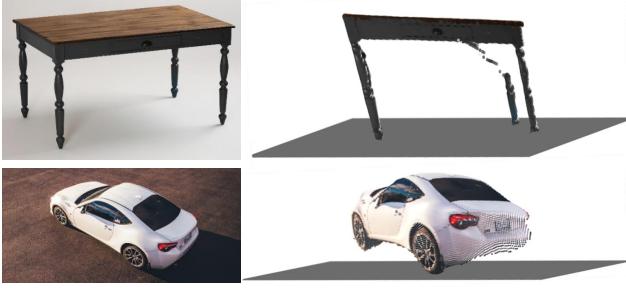


Figure 2. Without modeling object-ground correlation, existing single-view 3D estimation method [34] generates 3D models floating or tilted on the ground.

cause distortion in the 3D reconstruction (see Figure 2 top row). Without an explicit modeling of the object-ground relationship, recovered 3D objects are often hard to place on a flat support plane (see Figure 2 bottom row). These challenges are also present in recent category-specific 2D-to-3D methods that recover 3D shape from *latent embedding space* [9, 38, 39, 46, 53] and zero-shot novel-view synthesis methods [25, 26, 29, 33, 43, 57], where they often just assume a simple orthographic camera model, or assume the camera parameter being given as input to avoid over-complication of the problem, which on the other hand limits their application in unconstrained scenarios.

To address these challenges, we propose ORG (Object Reconstruction with the Ground), a new formulation for representing objects in relation to the ground. Given a single image, our objective is to simultaneously deduce the 3D shape of the object, its positioning relative to the ground plane, and the camera parameters. We compare our method with three existing research strands: depth estimation, latent embedding reconstruction, and diffusion-based novel-view synthesis, in addition to multi-view reconstruction techniques, as detailed in Table 1. Existing single-view methods often fail to maintain the object-ground relationship and usually presuppose known camera parameters or rely on overly simplistic camera models, leading to suboptimal performance for tasks like efficient shadow generation. In stark contrast, the output of our model supports the intricate interplay between object, ground, and camera (see Figure 1), facilitating superior shadow generation and pose-aware geometric reconstruction. To this end, we model the object as consisting of its front (visible) and back surfaces, and predict two pixel-level height map between the object and the ground [41], along with a dense camera parameter descriptor [16]. Our results demonstrate that such a simplified representation of objects is not only adequate for generating 3D-realistic shadows but also yields convincing reconstruction for a wide array of commonly encountered objects.

We create our training data from *Objaverse* [10], rendering six images for each object with diverse focal length and camera viewpoint. We evaluate our method across two unseen datasets, including objects and humans, and show

	Multi-view	Latent	Depth	NVS	Ours
single image	✗	✓	✓	✓	✓
category-free	✓	✗	✓	✓	✓
camera-aware	✓	✗	✗	✗	✓
ground-aware	✗	✗	✗	✗	✓

Table 1. ORG processes multiple advantages from flexibility to generalization against multi-view reconstruction work and other single-view work including generation from latent embedding, monocular depth estimation, and novel-view synthesis (NVS) methods.

qualitative results on random unseen web images. Our proposed method outperforms existing methods in terms of accuracy, robustness, and efficiency in various scenarios. Results show that our method achieves superior performance and provides a more comprehensive and light-weight solution to the challenges of single-view object geometry estimation. In summary, our main contributions are as follows.

- A novel framework ORG, for in-the-wild single-view object-ground 3D geometry estimation. To the best of our knowledge, this is the first method to jointly model object, camera, and ground plane from single image.
- We propose a perspective field guided pixel height re-projection module to efficiently convert our estimated representations into depth maps and point clouds.
- ORG achieves outstanding shadow generation and reconstruction performance on unseen real-world images, demonstrating great robustness and generalization ability.

2. Related Work

Single-view Depth Estimation. There has been significant progress made in recent times in the estimation of monocular depth [7, 8, 12, 34, 44, 61]. Given metric depth supervision, some work directly trains their model to regress the depth objective [12, 24, 60, 61]. While these methods achieve great performance on various datasets, the difficulty of obtaining metric ground truth depth hinders the use of direct depth supervision. Instead, another line of work relies on ranking losses, which evaluates relative depth [6, 54], or scale- and shift-invariant losses [34, 44] for supervision. The latter methods produce particularly robust depth predictions without heavy annotation efforts, but the models are not able to reason object-ground relationship and often produce unrealistic results when using depth map for downstream image editing tasks. In light of this, a recent work [41] proposes another annotation-friendly representation, pixel height, for better object shadow generation. However, this method has strict constraints on the camera viewpoint. We repurpose the representation for monocular 3D reconstruction and loosen the viewpoint by joint modeling camera with object geometry.

Single-view 3D Geometry Reconstruction. Reconstructing object shapes from single-view image is a challeng-

ing but well-established problem, with one of its seminal work [37] optimizing for 6-DoF poses of objects with known 3D models. In the ensuing decades, learning-based methods have begun to propose category-specific networks for 3D estimation that span a wide range of objects with [4, 18] and without direct 3D supervision [14, 17, 20, 58], and using neural implicit representations [30, 59]. With robust 3D supervision, recent methods have demonstrated the feasibility of learning 3D geometry with limited memory. Pixel2Mesh [46] offers a method to reconstruct the 3D shape with mesh using a single image input. Meanwhile, PIFu [38, 39] offers an efficient implicit function to recover high-resolution surfaces of humans, including previously unseen and occluded regions. While achieving great performance, some of these works rely on learning priors specific to a certain object category, hindering its generalizability in the wild. Recently, advances in text-to-3D generation [5, 22, 32, 50] also inspire image-to-3D generation using diffusion prior [25, 26, 29, 33, 43, 57]. Masked autoencoders are also used to object reconstruction from single image [52]. In comparison, our method is the first one to model the object geometry with respect to the ground for efficient image editing and 3D reconstruction.

Camera Parameter Estimation. An essential aspect of single-view monocular 3D object comprehension is to retrieve the focal length of a camera and the camera pose relative to the object and the ground plane. Classic methods leverage reference image components, including calibration grids [62] or vanishing points [11], to estimate camera parameters. Recently, data-driven approaches have been proposed to use deep neural networks to infer the focal length [15, 51] and camera poses [19, 28, 55] directly from in-the-wild images, or to use dense representation [16] to encode camera parameters for a more robust estimation. In contrast, our method ORG jointly estimates intrinsic and extrinsic camera parameters together with object geometry and ground positions, achieving a self-contained pipeline for 3D-aware image editing and reconstruction.

3. Approach

ORG considers single-view object geometry estimation by joint *pixel height* and *perspective field* prediction. We provide an overview of our framework in Figure 3. Modeling object geometry and camera parameters as dense fields, we first introduce the background knowledge of the dense object-ground and dense camera representations (Sec 3.1). We learn a pyramid vision transformer (PVT) [48, 49] to predict the dense representation fields (Sec 3.2), and prove that they can be repurposed for reconstruction task by proposing a perspective-guided pixel height reprojection method (Sec 3.3).

3.1. Object, Ground, and Camera Representations

Pixel Height Representation. Proposed for single-image shadow generation [41, 42], pixel height is a dense representation defined as the pixel distance between a point on an object and its ground projection, namely its vertical projection on the ground in the image, as we can see in Figure 3. It is a pixel-level scalar which measures the distance between object and its supporting plane in the image coordinate (number of pixels, rather than meters). Pixel height possesses many advantages over the depth representation in modeling the object geometry. First, it is disentangled from the camera model, and thus can be directly inferred from the images context without additional camera information. Moreover, it models the object and ground relationship, which is pivotal in generating realistic 3D models for real-world image applications, as objects almost always have a canonical position on the ground plane.

While photo-realistic shadows can be generated from pixel height map with projective geometry, we see more potential in this new representation. Constraining the object location with respect to a 2D plane, the pixel height representation plays a critical role in reconstructing 3D shape of objects on top of the ground. Moreover, strict requirements are enforced on camera viewpoints for pixel height [41], and only the front surface of a object is considered. Therefore, we propose to loosen this condition by modeling both front and back surfaces of the object, and jointly predicting camera intrinsic and pose relative to the ground. In the end, the field-of-view (FoV) is used to lift pixel distances into metric distance, and camera viewpoint helps align the object into the canonical pose relative to the ground.

Perspective Field Representation. As shown in Figure 3, the perspective field representation of a given image is composed of two dense fields, a latitude field represented by blue contour lines, and a up-vector field represented by green arrows [16]. Specifically, assuming a camera-centered spherical coordinate system where the zenith direction is opposite to gravity. The camera model \mathcal{K} projects a 3D position $X \in \mathbb{R}^3$ in the spherical coordinate into the image frame $x \in \mathbb{R}^2$. For each pixel location x , the up-vector is defined as the projection of the tangential direction of X along the meridian towards the north pole, the latitude is defined as the angle between the vector pointing from the camera to X and the ground plane. In other words, the latitude field and the up-vector field encode the elevation angle and the roll angle of the points on the object, respectively. Both perspective fields and pixel height map are invariant or equivariant to image editing operations like cropping, rotation and translation. As a result, they are highly suitable for neural network models designed for dense prediction tasks.

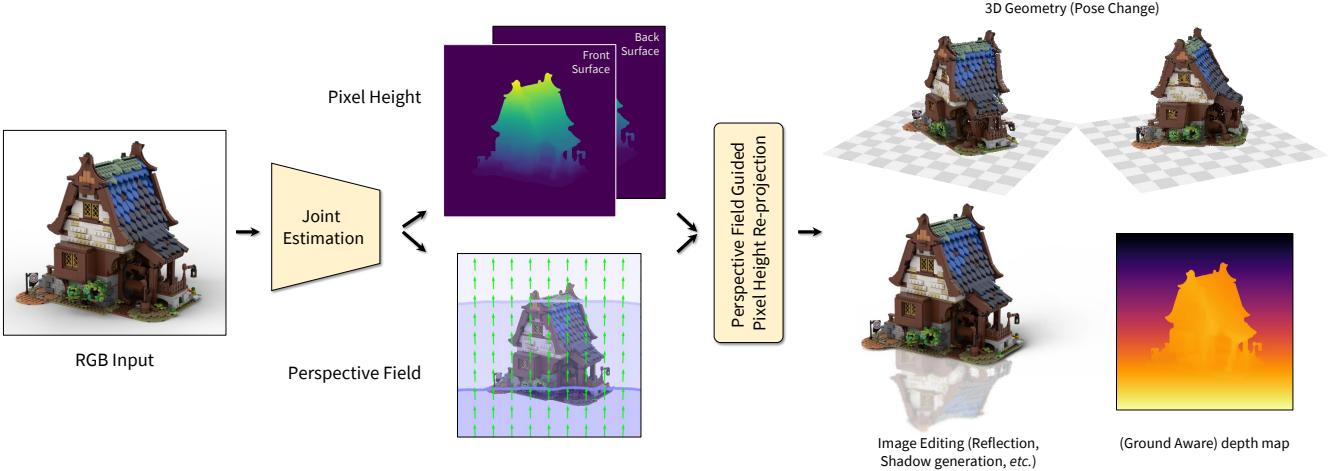


Figure 3. **ORG Paradigm.** Our proposed method is able to take a single-view object-centric image as input, and jointly estimate two dense representations, the pixel height and perspective field, encoding the object-ground relationship and camera parameters, respectively. A Perspective Field Guided Pixel Height Re-projection module is proposed to repurpose the two predicted dense fields into depth map estimation and point cloud generation.

3.2. Dense Field Estimation

We present a neural network model to estimate the two dense fields from a single image. The per-pixel structure and translation-invariant nature of the pixel height and perspective field representations make them highly suitable for neural network prediction. Following [34, 61], we formulate the dense field estimation task as a regression problem. Specifically, for each image pixel of the pixel height field, assuming a ray starting from a camera pointing towards the pixel travels through the object, there will be an entry point on the front surface of the object p_f and an exit point on the back surface of the object p_b . When the ray passes the surfaces of object multiple times, we only consider the first entry and last exit. The model is then asked to predict the pixel height for both p_f and p_b . Moreover, we normalize it with the height of the input image. For latitude field, we normalize the original $[-\pi/2, \pi/2]$ range into $[0, 1]$. And for the up-vector field, each angle θ can range from 0 to 2π , so direct normalization and regression pose ambiguity for the model since 0 and 2π represent the same angle. Hence, we represent each angle θ with a tuple $(\sin \theta, \cos \theta)$, and train the model to regress to a two-channel vector map. All regression tasks are trained with loss ℓ_2 .

Model Architecture and Training Details. We use the architecture of PVTv2-b3 [49] as our backbone to extract joint feature map. We use SegFormer [56] with the Mix Transformer-B3 as our decoder. Residual connection is added before the decoder to include lower-level context from the 2-layer CNN block. We find that transformer-based encoder is suitable for our task as it effectively maintains global consistency in the two dense representation

fields. We further make modifications to the decoder head, enabling it to produce a regression value for the pixel height map, up field map, and latitude field map. We use PVTv2-b3 pretrained on COCO dataset [23] as the backbone of our architecture. The model is trained with AdamW [27] optimizer with learning rate 0.0005 and weight decay $1e-2$ for 60K steps with batch size 8 on a 4-A100 machine. We schedule the multi-step training stages at step 30K, 40K, and 50K, with learning rate decreases by 10 time at each stage. We resize the image to $(512, 512)$ before using horizontal flipping, random cropping, and color jittering augmentation during training.

3.3. Perspective-Guided Pixel Height Reprojection

After predicting two dense representations, we prove that they encode sufficient information to be efficiently converted into depth maps and point clouds for downstream tasks and for fair comparison with existing methods. First, since the perspective field can be generated from camera parameters, we discretize the continuous parameter range and use a grid search optimization strategy to estimate camera field-of-view α and extrinsic rotation matrix \mathbf{R} as row and pitch angles. Afterwards, the camera focal length is calculated as $f = \frac{H}{2 \tan \alpha/2}$, where H is the height of the input image. Then the intrinsic matrix \mathbf{K} is also estimated as:

$$\mathbf{K} = \begin{bmatrix} f & 0 & c_x \\ 0 & f & c_y \\ 0 & 0 & 1 \end{bmatrix}, \quad (1)$$

where (c_x, c_y) is the principle point of the image and is usually estimated to be the center of the image.

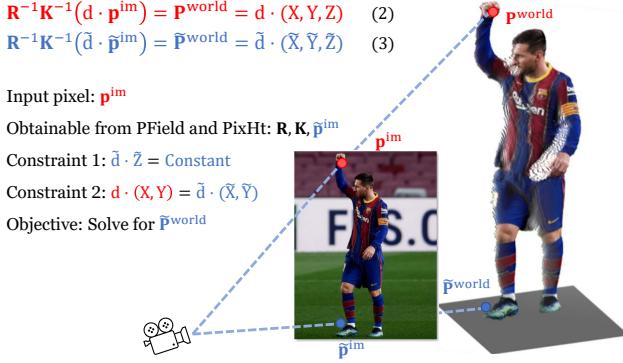


Figure 4. Perspective-Guided Pixel Height Reprojection. PFeld and PixHt are perspective field and pixel height, respectively.

Derivation. An illustration is provided in Figure 4. Given one pixel $\mathbf{p}^{\text{im}} = (x, y) \in \mathbb{R}^2$, we known its vertical projection point $\tilde{\mathbf{p}}^{\text{im}} = (\tilde{x}, \tilde{y}) \in \mathbb{R}^2$ on the ground in the image frame, given by the estimated pixel height map. Recall that intrinsic and extrinsic matrices can be used to project a 3D point \mathbf{P}_i in the world coordinate into an image pixel \mathbf{p}_i . More specifically, given intrinsic matrix \mathbf{K} and extrinsic rotation matrix \mathbf{R} , we have the following equations describing the correspondence between a pixel \mathbf{p}^{im} on the object and its corresponding 3D points $\mathbf{P}^{\text{world}}$ in the world coordinate:

$$\text{object: } \mathbf{R}^{-1} \mathbf{K}^{-1} (d \cdot \mathbf{p}^{\text{im}}) = \mathbf{P}^{\text{world}} = d \cdot (X, Y, Z) \quad (2)$$

$$\text{ground: } \mathbf{R}^{-1} \mathbf{K}^{-1} (\tilde{d} \cdot \tilde{\mathbf{p}}^{\text{im}}) = \tilde{\mathbf{P}}^{\text{world}} = \tilde{d} \cdot (\tilde{X}, \tilde{Y}, \tilde{Z}) \quad (3)$$

where d is the depth value of the point. Here, the point $\tilde{\mathbf{p}}^{\text{im}}$ in Eq. (3) is the vertical projection of the ground of \mathbf{p}^{im} in Eq. (2). For a given pixel \mathbf{p}^{im} , its corresponding $\tilde{\mathbf{p}}^{\text{im}}$ can be obtained from our estimated vertical direction (*perspective field*) and the estimated *pixel height*. Note that the world coordinate has its Z axis pointing vertically upwards, and its XY plane parallel to the ground plane. The objective is to obtain the location of the reconstructed 3D point $\mathbf{P}^{\text{world}} = d \cdot (X, Y, Z)$, and to eliminate the unknown depth d , we need two additional constraints with the help of Eq. (3). The constraint one is that *all 3D points $\tilde{\mathbf{P}}^{\text{world}}$ on the ground have a constant z -axis value*. Without loss of generality, we assume that the constant is one, to obtain a **scale-invariant** 3D point cloud. This gives us $\tilde{d} = 1/\tilde{Z}$, which then leads to the normalized $\tilde{\mathbf{P}}^{\text{world}}$:

$$\tilde{\mathbf{P}}^{\text{world}} = (\tilde{X}/\tilde{Z}, \tilde{Y}/\tilde{Z}, 1) = (X_n, Y_n, 1) \quad (4)$$

Then the constraint two is that *the 3D point $\mathbf{P}^{\text{world}}$ and its vertical ground projection $\tilde{\mathbf{P}}^{\text{world}}$ have identical XY coordinates*. With this, we know that $d = \frac{X_n}{X} = \frac{Y_n}{Y}$. We calculate $d = \frac{X_n Y_n}{X Y}$ for numerical robustness, and the final normalized 3D point is

$$\mathbf{P}_n^{\text{world}} = \frac{X_n Y_n}{X Y} \cdot (X, Y, Z) \quad (5)$$

where X, Y, Z, X_n, Y_n are calculated from Eqs. (2) to (4).

4. Experiments

In this section, we conduct extensive qualitative and quantitative experiments to demonstrate the effectiveness and generalizability of ORG. We evaluate our model with the classic depth estimation metric and point cloud reconstruction metric on both object-centric images and human-centric images. We show that repurposing two dense representation predictions leads to a very robust 3D reconstruction framework for diverse categories and viewpoints of images.

4.1. Data Rendering

Existing object-centric datasets [1, 36] do not provide accurate depth map and object-ground rotation information simultaneously. Hence, we render a large-scale dataset from *Objaverse* [10]. Objaverse is a large-scale object-centric dataset consisting of over 800K high-quality 3D models. For each object in the dataset, we randomly sample 6 sets of camera intrinsic and extrinsic parameters (FoV and rotation matrix), each is used to render an RGB image with pixel height and perspective field ground truth maps. The image dimension is (512, 512). The camera always points at the center of the object and the z -axis of world coordinates points orthogonally to the ground plane. We use a physically-based renderer Blender [3] to render realistic surface appearance and develop a CUDA-based ray-tracer to efficiently render front and back surface pixel heights. We perform corrupt data filtering to remove images with incorrect annotations and images with objects that are too small on the canvas. This results in 3,364,052 images in the dataset in total. We split the objects into train/val/test sets in 8:1:1. We also randomize the intensity, position, number of light sources, and distance between camera and object to increase the diversity of the dataset. We will release our *data rendering script* and *rendered dataset*. More details of the implementation and the dataset are in the supplementary.

4.2. Baselines

We compare our method with single-view depth estimation, image-to-3D reconstruction, and camera parameter estimation work. For depth estimation work, we compare with LeReS [61], MiDaS [2, 34, 35], and MegaDepth [21], which are single-view generic depth estimation methods pretrained on large-scale datasets. For image-to-3D reconstruction work, we compare with Zero-123 [26], a single-image novel-view synthesis and reconstruction method also pretrained on Objaverse dataset [10]. For camera parameter estimation, we compare with the state-of-the-art off-the-shelf camera estimator CTRL-C [19] and a heuristic method we implemented by eyeballing a rough FoV and pitch angle for all evaluation samples in the test set to get the camera focal length and rotation matrix. Using estimated camera parameters, we can convert the predicted depth map into point clouds. Note that in order to generate depth map and

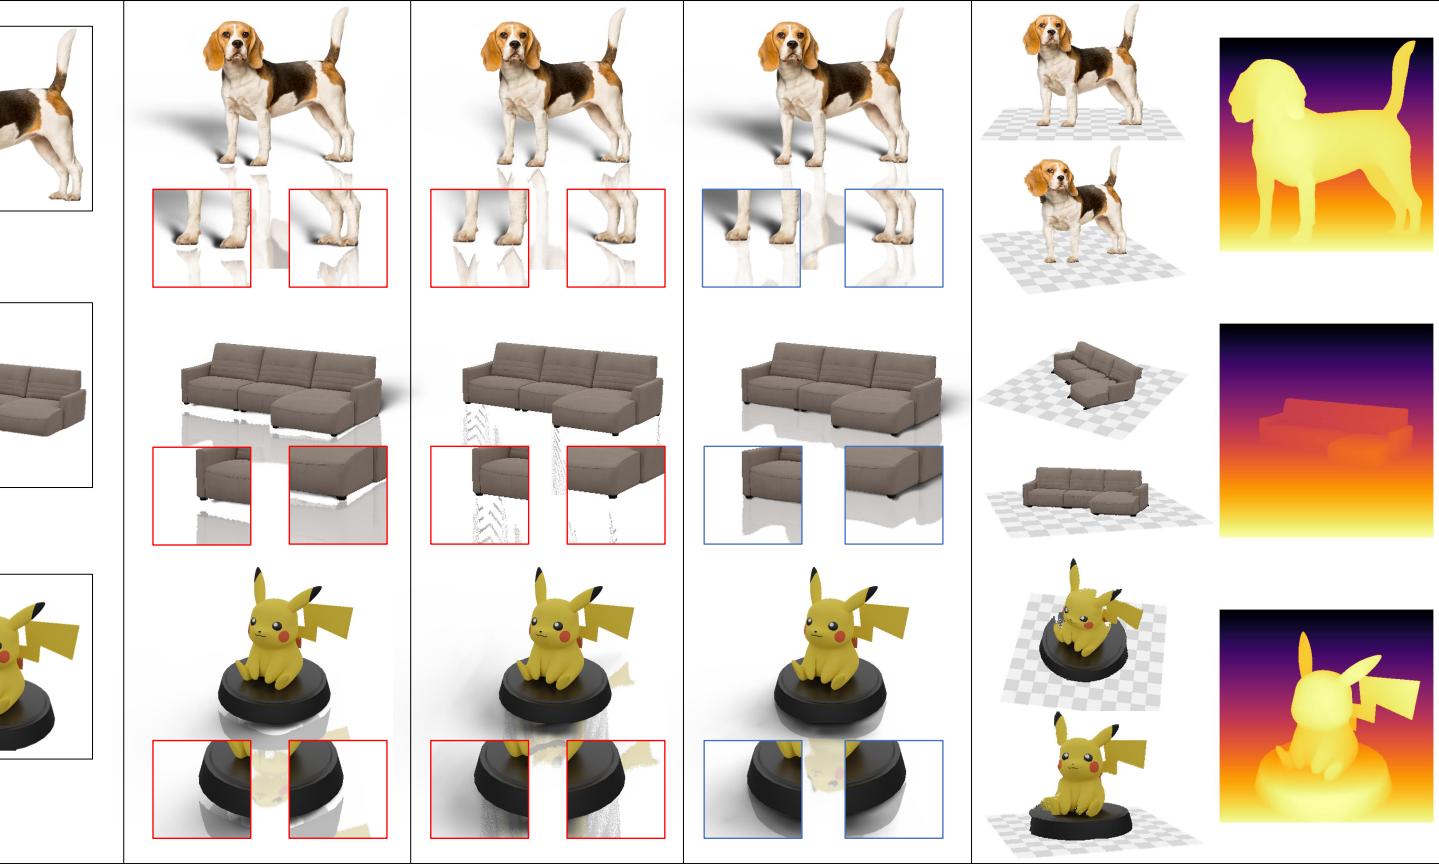


Figure 5. Qualitative results of shadow and reflection generation on the ground, as well as object-ground reconstruction and depth estimation. We show comparison with the depth-based estimation method LeReS [61] and monocular novel view synthesis method Zero-123 [26]. ORG maintains great object-ground relationship compared with prior methods which leads to much more realistic shadow and reflection generation, as shown in the blue boxes. Our method runs very fast and can easily output representations like depth map and point cloud.

point clouds for objects, we use image mask to remove the background region of our prediction, as well as for existing methods, as we can see in Figure 5. More details are provided in the supplementary.

Metrics. For a fair comparison with existing methods, we evaluate our method on depth estimation and point cloud reconstruction tasks. In the meantime, we visualize the estimated ground plane together with reconstruction objects to validate the object-ground correlation. For depth estimation, following previous methods [34, 61], we use the absolute mean relative error (AbsRel) and the percentage of pixels with $\delta_1 = \max(\frac{d_i}{d_i^*}, \frac{d_i^*}{d_i}) < 1.25$. We follow MiDaS [34] and LeReS [61] to align the scale and shift before evaluation. For point cloud estimation, following prior work [8, 61], we use Locally Scale Invariant RMSE (LSIV) and Chamfer Distance (CD). In addition, we also evaluate our direct estimation on pixel height, latitude-vector field, and up-vector field using mean-square error (MSE) and absolute error (L1).

	small	medium	large
Baseline	0.23	0.37	0.72
ORG (Ours)	0.21	0.28	0.45
<i>diff</i>	-0.02	-0.09	-0.27

Table 2. ORG achieves higher improvement against baseline model (DPT-BeiT [2, 35] + Ctrl-C [19]) when objects have larger viewpoint diversity. We report results on point clouds LSIV metrics on validation set. *Small*, *medium*, and *large* stand for different levels of viewpoint diversity of the samples.

4.3. Shadow, Reflection, and Reconstruction

We show results for 3D reconstruction, shadow generation, and reflection generation on unseen objects in Figure 5. We compare generation performance with the monocular depth estimation method [61] and the novel view synthesis method [26]. For both methods, we use Ctrl-C [19] to predict camera parameters. Since these methods do not model the ground explicitly, we use the estimated pitch angle to

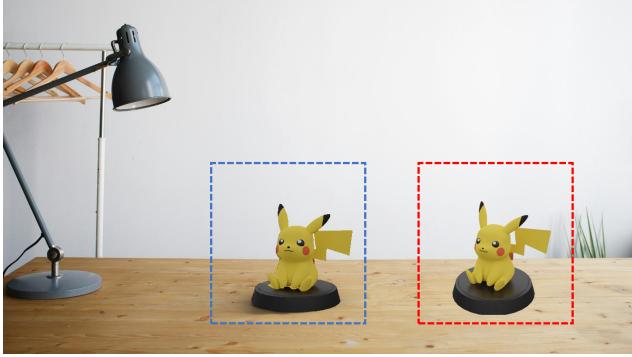


Figure 6. Reconstructed object in the realistic background. *Blue*: novel view synthesis and realistic background composition with our method. *Red*: direct background composition.

Object Geometry	Camera Parameters	LSIV↓	diff
depth	OFS estimator	1.25	0
depth	perspective field	1.01	-0.24
pixel height	OFS estimator	0.98	-0.27
pixel height	perspective field	0.81	-0.44

Table 3. Our proposed joint estimation of pixel height and perspective field contribute to the final performance. We report results on point clouds LSIV metrics. OFS stands for off-the-shelf, and we use Ctrl-C as the OSF estimation in this experiment.

obtain the ground plane by assuming that it passes through the object’s lowest point (point with smallest height value). For the baseline for novel view synthesis, we use SJC [45] to reconstruct the shape of the object. As depicted in Figure 5, notably, there is a marked improvement in the quality of shadows and reflections, particularly at contact points on the ground, as highlighted in the designated boxes. Our research also includes object-ground reconstructions and depth map conversions. The 3D shape of the reconstructed models in our work is not only realistic but also maintains an accurate vertical alignment with the ground plane. This visualization effectively demonstrates our model’s versatility, showcasing its exceptional performance across a wide array of object categories, poses, and viewpoints.

4.4. Novel View Synthesis and Image Composition

We demonstrate applications such as object view manipulation, shadow generation, and image composition in Figure 6. In the red box, we show direct copy-and-paste composition as a comparison, and performance of ORG is shown in the blue box. We notice that the simple copy-pasting method does not match the camera perspective of the new object and its supporting plane in the background, creating unrealistic visual effects. Our method, on the other hand, estimate the background perspective, reconstruct the object into 3D and re-render it from the target perspective, and generate photo-realistic shadow from the estimated object shape, achieving better visual alignment and realism.

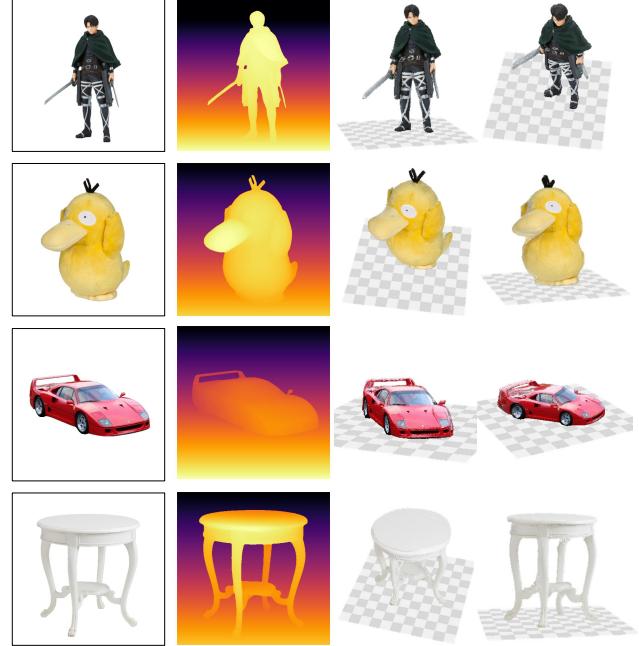


Figure 7. More qualitative results of ORG in depth map generation and object-ground reconstruction. Our method generalizes well to unseen in-the-wild images.

More Qualitative Results. Moreover, Figure 7 illustrates additional qualitative results from our study, focusing on depth map generation and object-ground reconstruction. Our methodology exhibits remarkable proficiency in reconstructing ground-supported objects of various types, underscoring the robustness of our approach.

4.5. Object with Diverse Viewpoints

We also break down the evaluation into subsets of samples with different range of camera angles. More specifically, we divide the difficulty level by the pitch angle because natural images usually have more diverse pitch angles but close to zero roll angles. Taking the mean pitch angle of the entire dataset, samples with a pitch angle smaller than 10 degrees difference than the mean angle are marked as *small* viewpoint diversity. Samples with a pitch angle difference between 10 and 30 degrees are marked as *medium* viewpoint diversity, and samples with a pitch angle difference greater than 30 degrees are marked as *large* viewpoint diversity. The results in Table 2 show that ORG achieves a higher improvement compared to the baseline model (LeReS [61] + Ctrl-C [19]) when objects have greater viewpoint diversity. This is because the traditional viewpoint estimation model struggles for object-centric images, especially for samples with extreme pitch angles.

4.6. Importance of Joint Estimation

The results in Table 3 show that the joint learning of pixel height and perspective field leads to the best reconstruction

camera parameters	Depth Map		Point Clouds		Pixel Height	Lati-Vector	Up-Vector
	AbsRel \downarrow	$\delta_1 \uparrow$	LSIV \downarrow	CD \downarrow	L1 \downarrow	L1 \downarrow	L1 \downarrow
MegaDepth [21]	39.4	53.7	1.60	1.73	36.8		
NDDepth [40]	35.8	54.2	1.49	1.65	30.9		
MiDaS [2, 34, 35]	heuristic const.	22.7	77.9	1.31	1.45	26.0	8.77
LeReS [61]		30.0	63.1	1.11	1.34	24.5	
ORG (Ours)		24.6	71.2	1.07	1.39	15.4	
MegaDepth [21]	39.4	53.7	1.51	1.64	31.1		
NDDepth [40]	35.8	54.2	1.46	1.60	28.3		
MiDaS [2, 34, 35]	Ctrl-C [19]	22.7	77.9	1.22	1.39	20.7	5.45
LeReS [61]		30.0	63.1	1.05	1.31	20.4	
DUSt3R [47]		17.9	79.5	0.95	1.29	15.9	
ORG (Ours)		21.1	76.0	0.99	1.27	15.4	
ORG (Ours)	Ours	19.1	81.2	0.93	1.26	15.4	4.94
							1.45

Table 4. ORG perform consistently the best in both depth estimation and point cloud estimation tasks of object-centric images under all metrics. We use two off-the-shelf camera estimation algorithms to make up for the unknown camera parameters. Pixel Height metric is reported in absolute error of number of pixels, Latitude-vector Field and Up-Vector Field are reported in degrees.

performance compared to the depth estimation and off-the-shelf camera parameter estimator. More specifically, without modifying the model architecture, we change the objective of our model from pixel height estimation to depth estimation following the loss used in LeReS [61]. Trained with the same dataset and scheduler, the pixel height representation is able to achieve better point-cloud reconstruction than depth-based learning. We argue that this is because the representation focuses more on object-ground geometry rather than object-camera geometry, which is more natural and easier to infer from object-centric images. This observation further validates that the superior generalizability of ORG comes from the better representation design and the joint training strategy, rather than the dataset.

4.7. Qualitative Evaluation on Reconstruction

We compare the depth map estimation, point cloud generation, and the prediction of our representations with four state-of-the-art monocular depth estimation and 3D reconstruction methods on the held-out test set. We use the state-of-the-art camera parameter estimation model Ctrl-C [19] and a heuristic estimation to compensate for missing intrinsic and extrinsic information from previous methods. We convert the raw output into depth map and point clouds for a fair comparison with existing methods. As shown in Table 4, our method performs consistently the best in both depth estimation and point cloud estimation tasks for object-centric images under all metrics. We also try using the other two alternative camera parameter estimators to reconstruct the point cloud from the pixel height estimation. And we can see that using the same off-the-shelf camera estimator, ORG can still outperform existing methods on both two tasks. We make sure that no samples in the evaluation dataset are seen by prior methods or our method during the

training phase, in order to create a zero-shot evaluation scenario. The results show that ORG achieves a great generalization ability in the object-centric 3D reconstruction task.

Furthermore, we also break down the evaluation into pixel height, latitude vector, and up-vector estimation, and evaluate with mean absolute error in the generic space of all three predictions (number of pixels for pixel height and degrees for two perspective fields). For prior methods, we use Ctrl-C and the heuristic constant (by grid search) to estimate elevation angle, roll angle, and camera FoV, and convert them into perspective field representations for comparison. Their pixel height estimations are also converted using depth estimation and camera parameter estimations. As we can see in Table 4, our method outperforms the baselines in the three tasks. These experiments demonstrate the robustness and generalizability of ORG over prior methods in object 3D estimation and reconstruction.

For evaluations on non-symmetric object front- and back-surface reconstruction, and the pixel height, depth map, and perspective field visualization, please refer to the supplementary material.

5. Conclusion

In this paper, we proposed ORG, to our knowledge, the first data-driven architecture that simultaneously reconstructs 3D object, estimates camera parameter, and models the object-ground relationship from a monocular image. To achieve this, we propose a new formulation for representing objects in relation to the ground. Qualitative and quantitative results on unseen object and human datasets as well as web images demonstrate the robustness and flexibility of our model, which marks a significant step towards in-the-wild single-image object geometry estimation.

Acknowledgments. This work was supported in part by NSF Grant 2106825, NIFA Award 2020-67021-32799, and Adobe. This work used computational resources, including the NCSA Delta and DeltaAI supercomputers through allocations CIS230012, CIS230013, and CIS240370 from the Advanced Cyberinfrastructure Coordination Ecosystem: Services & Support (ACCESS) program, as well as the TACC Frontera supercomputer and Amazon Web Services (AWS) through the National Artificial Intelligence Research Resource (NAIRR) Pilot.

References

[1] Adel Ahmadyan, Liangkai Zhang, Artsiom Ablavatski, Jianing Wei, and Matthias Grundmann. Objectron: A large scale dataset of object-centric videos in the wild with pose annotations. In *CVPR*, 2021. 5

[2] Reiner Birkl, Diana Wofk, and Matthias Müller. Midas v3.1 – a model zoo for robust monocular relative depth estimation. *arXiv preprint arXiv:2307.14460*, 2023. 5, 6, 8

[3] Blender Online Community. Blender - a 3D modelling and rendering package, 1994. 5, 1

[4] Thomas J Cashman and Andrew W Fitzgibbon. What shape are dolphins? building 3D morphable models from 2D images. *TPAMI*, 2012. 3

[5] Rui Chen, Yongwei Chen, Ningxin Jiao, and Kui Jia. Fantasia3D: Disentangling geometry and appearance for high-quality text-to-3D content creation. In *ICCV*, 2023. 3

[6] Weifeng Chen, Zhao Fu, Dawei Yang, and Jia Deng. Single-image depth perception in the wild. In *NeurIPS*, 2016. 1, 2

[7] Weifeng Chen, Shengyi Qian, and Jia Deng. Learning single-image depth from videos using quality assessment networks. In *CVPR*, 2019. 1, 2

[8] Weifeng Chen, Shengyi Qian, David Fan, Noriyuki Kojima, Max Hamilton, and Jia Deng. OASIS: A large-scale dataset for single image 3D in the wild. In *CVPR*, 2020. 2, 6

[9] Yen-Chi Cheng, Hsin-Ying Lee, Sergey Tulyakov, Alexander Schwing, and Liangyan Gui. SDFusion: Multimodal 3D Shape Completion, Reconstruction, and Generation. In *CVPR*, 2023. 2

[10] Matt Deitke, Dustin Schwenk, Jordi Salvador, Luca Weihs, Oscar Michel, Eli VanderBilt, Ludwig Schmidt, Kiana Ehsani, Aniruddha Kembhavi, and Ali Farhadi. Objaverse: A universe of annotated 3D objects. In *CVPR*, 2023. 2, 5

[11] Jonathan Deutscher, Michael Isard, and John MacCormick. Automatic camera calibration from a single manhattan image. In *ECCV*, 2002. 3

[12] David Eigen, Christian Puhrsch, and Rob Fergus. Depth map prediction from a single image using a multi-scale deep network. In *NeurIPS*, 2014. 2

[13] Daniel Gatis. Rembg. <https://github.com/danielgatis/rembg>, 2023. 1

[14] Shubham Goel, Angjoo Kanazawa, and Jitendra Malik. Shape and viewpoint without keypoints. In *ECCV*, 2020. 3

[15] Yannick Hold-Geoffroy, Kalyan Sunkavalli, Jonathan Eisenmann, Matthew Fisher, Emiliano Gambaretto, Sunil Hadap, and Jean-François Lalonde. A perceptual measure for deep single image camera calibration. In *CVPR*, 2018. 3

[16] Linyi Jin, Jianming Zhang, Yannick Hold-Geoffroy, Oliver Wang, Kevin Blackburn-Matzen, Matthew Sticha, and David F Fouhey. Perspective fields for single image camera calibration. In *CVPR*, 2023. 2, 3

[17] Angjoo Kanazawa, Shubham Tulsiani, Alexei A Efros, and Jitendra Malik. Learning category-specific mesh reconstruction from image collections. In *ECCV*, 2018. 3

[18] Abhishek Kar, Shubham Tulsiani, Joao Carreira, and Jitendra Malik. Category-specific object reconstruction from a single image. In *CVPR*, 2015. 3

[19] Jinwoo Lee, Hyunsung Go, Hyunjoon Lee, Sunghyun Cho, Minhyuk Sung, and Junho Kim. Ctrl-C: Camera calibration transformer with line-classification. In *ICCV*, 2021. 3, 5, 6, 7, 8

[20] Xueteng Li, Sifei Liu, Kihwan Kim, Shalini De Mello, Varun Jampani, Ming-Hsuan Yang, and Jan Kautz. Self-supervised single-view 3D reconstruction via semantic consistency. In *ECCV*, 2020. 3

[21] Zhengqi Li and Noah Snavely. MegaDepth: Learning single-view depth prediction from internet photos. In *CVPR*, 2018. 5, 8

[22] Chen-Hsuan Lin, Jun Gao, Luming Tang, Towaki Takikawa, Xiaohui Zeng, Xun Huang, Karsten Kreis, Sanja Fidler, Ming-Yu Liu, and Tsung-Yi Lin. Magic3D: High-resolution text-to-3D content creation. In *CVPR*, 2023. 3

[23] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft COCO: Common objects in context. In *ECCV*, 2014. 4, 1

[24] Fayah Liu, Chunhua Shen, Guosheng Lin, and Ian Reid. Learning depth from single monocular images using deep convolutional neural fields. *TPAMI*, 2015. 2

[25] Minghua Liu, Chao Xu, Haian Jin, Linghao Chen, Mukund Varma T, Zexiang Xu, and Hao Su. One-2-3-45: Any single image to 3D mesh in 45 seconds without per-shape optimization. In *NeurIPS*, 2023. 2, 3

[26] Ruoshi Liu, Rundi Wu, Basile Van Hoorick, Pavel Tokmakov, Sergey Zakharov, and Carl Vondrick. Zero-1-to-3: Zero-shot one image to 3D object. In *ICCV*, 2023. 1, 2, 3, 5, 6

[27] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *ICLR*, 2017. 4, 1

[28] Yunze Man, Xinshuo Weng, Xi Li, and Kris Kitani. Ground-Net: Monocular ground plane normal estimation with geometric consistency. In *ACMMM*, 2019. 3

[29] Luke Melas-Kyriazi, Christian Rupprecht, Iro Laina, and Andrea Vedaldi. RealFusion: 360 reconstruction of any object from a single image. In *CVPR*, 2023. 2, 3

[30] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3D reconstruction in function space. In *CVPR*, 2019. 3

[31] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An

imperative style, high-performance deep learning library. In *NeurIPS*, 2019. 1

[32] Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. DreamFusion: Text-to-3D using 2D diffusion. In *ICLR*, 2022. 3

[33] Guocheng Qian, Jinjie Mai, Abdullah Hamdi, Jian Ren, Aliaksandr Siarohin, Bing Li, Hsin-Ying Lee, Ivan Skorokhodov, Peter Wonka, Sergey Tulyakov, et al. Magic123: One image to high-quality 3D object generation using both 2D and 3D diffusion priors. *arXiv preprint arXiv:2306.17843*, 2023. 2, 3

[34] René Ranftl, Katrin Lasinger, David Hafner, Konrad Schindler, and Vladlen Koltun. Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *TPAMI*, 2020. 1, 2, 4, 5, 6, 8

[35] René Ranftl, Alexey Bochkovskiy, and Vladlen Koltun. Vision transformers for dense prediction. In *ICCV*, 2021. 5, 6, 8

[36] Jeremy Reizenstein, Roman Shapovalov, Philipp Henzler, Luca Sbordone, Patrick Labatut, and David Novotny. Common Objects in 3D: Large-scale learning and evaluation of real-life 3D category reconstruction. In *ICCV*, 2021. 5

[37] Lawrence G Roberts. *Machine perception of three-dimensional solids*. PhD thesis, MIT, 1963. 3

[38] Shunsuke Saito, Zeng Huang, Ryota Natsume, Shigeo Morigima, Angjoo Kanazawa, and Hao Li. PIFu: Pixel-aligned implicit function for high-resolution clothed human digitization. In *ICCV*, 2019. 2, 3, 1

[39] Shunsuke Saito, Tomas Simon, Jason Saragih, and Hanbyul Joo. PIFuHD: Multi-level pixel-aligned implicit function for high-resolution 3D human digitization. In *CVPR*, 2020. 1, 2, 3

[40] Shuwei Shao, Zhongcai Pei, Weihai Chen, Xingming Wu, and Zhengguo Li. NDDepth: Normal-distance assisted monocular depth estimation. In *ICCV*, 2023. 8

[41] Yichen Sheng, Yifan Liu, Jianming Zhang, Wei Yin, A Cengiz Oztireli, He Zhang, Zhe Lin, Eli Shechtman, and Bedrich Benes. Controllable shadow generation using pixel height maps. In *ECCV*, 2022. 2, 3

[42] Yichen Sheng, Jianming Zhang, Julien Philip, Yannick Hold-Geoffroy, Xin Sun, He Zhang, Lu Ling, and Bedrich Benes. PixHt-Lab: Pixel height based light effect generation for image compositing. In *CVPR*, 2023. 3

[43] Junshu Tang, Tengfei Wang, Bo Zhang, Ting Zhang, Ran Yi, Lizhuang Ma, and Dong Chen. Make-it-3D: High-fidelity 3D creation from a single image with diffusion prior. In *ICCV*, 2023. 2, 3

[44] Chaoyang Wang, Simon Lucey, Federico Perazzi, and Oliver Wang. Web stereo video supervision for depth prediction from dynamic scenes. In *3DV*, 2019. 2

[45] Haochen Wang, Xiaodan Du, Jiahao Li, Raymond A Yeh, and Greg Shakhnarovich. Score jacobian chaining: Lifting pretrained 2D diffusion models for 3D generation. In *CVPR*, 2023. 7

[46] Nanyang Wang, Yinda Zhang, Zhuwen Li, Yanwei Fu, Wei Liu, and Yu-Gang Jiang. Pixel2Mesh: Generating 3D mesh models from single RGB images. In *ECCV*, 2018. 2, 3

[47] Shuzhe Wang, Vincent Leroy, Yohann Cabon, Boris Chidlovskii, and Jerome Revaud. DUST3R: Geometric 3D vision made easy. In *CVPR*, 2024. 8

[48] Wenhui Wang, Enze Xie, Xiang Li, Deng-Ping Fan, Kaitao Song, Ding Liang, Tong Lu, Ping Luo, and Ling Shao. Pyramid vision transformer: A versatile backbone for dense prediction without convolutions. In *ICCV*, 2021. 3, 1

[49] Wenhui Wang, Enze Xie, Xiang Li, Deng-Ping Fan, Kaitao Song, Ding Liang, Tong Lu, Ping Luo, and Ling Shao. Pvt v2: Improved baselines with pyramid vision transformer. *CVMJ*, 2022. 3, 4

[50] Zhengyi Wang, Cheng Lu, Yikai Wang, Fan Bao, Chongxuan Li, Hang Su, and Jun Zhu. ProlificDreamer: High-fidelity and diverse text-to-3D generation with variational score distillation. In *NeurIPS*, 2023. 3

[51] Scott Workman, Connor Greenwell, Menghua Zhai, Ryan Baltenberger, and Nathan Jacobs. Deepfocal: A method for direct focal length estimation. In *ICIP*, 2015. 3

[52] Chao-Yuan Wu, Justin Johnson, Jitendra Malik, Christoph Feichtenhofer, and Georgia Gkioxari. Multiview compressive coding for 3D reconstruction. In *CVPR*, 2023. 3

[53] Jiajun Wu, Chengkai Zhang, Xiuming Zhang, Zhoutong Zhang, William T Freeman, and Joshua B Tenenbaum. Learning shape priors for single-view 3D completion and reconstruction. In *ECCV*, 2018. 2

[54] Ke Xian, Chunhua Shen, Zhiguo Cao, Hao Lu, Yang Xiao, Ruibo Li, and Zhenbo Luo. Monocular relative depth perception with web stereo data supervision. In *CVPR*, 2018. 2

[55] Wenqi Xian, Zhengqi Li, Matthew Fisher, Jonathan Eisenmann, Eli Shechtman, and Noah Snavely. UprightNet: Geometry-aware camera orientation estimation from single images. In *ICCV*, 2019. 3

[56] Enze Xie, Wenhui Wang, Zhiding Yu, Anima Anandkumar, Jose M Alvarez, and Ping Luo. Segformer: Simple and efficient design for semantic segmentation with transformers. In *NeurIPS*, 2021. 4, 1

[57] Dejia Xu, Yifan Jiang, Peihao Wang, Zhiwen Fan, Yi Wang, and Zhangyang Wang. NeuralLift-360: Lifting an in-the-wild 2D photo to a 3D object with 360° views. In *CVPR*, 2023. 1, 2, 3

[58] Yufei Ye, Shubham Tulsiani, and Abhinav Gupta. Shelf-supervised mesh prediction in the wild. In *CVPR*, 2021. 3

[59] Yufei Ye, Abhinav Gupta, and Shubham Tulsiani. What's in your hands? 3D reconstruction of generic objects in hands. In *CVPR*, 2022. 3

[60] Wei Yin, Yifan Liu, Chunhua Shen, and Youliang Yan. Enforcing geometric constraints of virtual normal for depth prediction. In *ICCV*, 2019. 2

[61] Wei Yin, Jianming Zhang, Oliver Wang, Simon Niklaus, Long Mai, Simon Chen, and Chunhua Shen. Learning to recover 3D scene shape from a single image. In *CVPR*, 2021. 1, 2, 4, 5, 6, 7, 8

[62] Zhengyou Zhang. A flexible new technique for camera calibration. *TPAMI*, 2000. 3