

Digital Twin Catalog: A Large-Scale Photorealistic 3D Object Digital Twin Dataset

Zhao Dong¹ Ka Chen^{1,*} Zhaoyang Lv^{1,*} Hong-Xing Yu^{2,*} Yunzhi Zhang^{2,*} Cheng Zhang^{1,*}
 Yufeng Zhu^{1,*} Stephen Tian² Zhengqin Li¹ Geordie Moffatt¹ Sean Christofferson¹ James Fort¹
 Xiaqing Pan¹ Mingfei Yan¹ Jiajun Wu² Carl Yuheng Ren¹ Richard Newcombe¹

¹Meta Reality Labs Research ²Stanford University




Figure 1. The Digital Twin Catalog (DTC) dataset comprises 2,000 digital twins of physical-world objects (a), characterized by millimeter-level geometric accuracy and photorealistic PBR materials (b). DTC includes evaluation data captured using both DSLR cameras and egocentric Aria glasses, featuring captured images with precise foreground object masks and environment lighting for relighting evaluation.

Abstract

We introduce Digital Twin Catalog (DTC), a new large-scale photorealistic 3D object digital twin dataset. A digital twin of a 3D object is a highly detailed, virtually indistinguishable representation of a physical object, accurately capturing its shape, appearance, physical properties, and other attributes. Recent advances in neural-based 3D reconstruction and inverse rendering have significantly improved the quality of 3D object reconstruction. Despite these advancements, there remains a lack of a large-scale, digital twin quality real-world dataset and benchmark that can quantitatively assess and compare the performance of different reconstruction methods, as well as improve reconstruction quality through

training or fine-tuning. Moreover, to democratize 3D digital twin creation, it is essential to integrate creation techniques with next-generation egocentric computing platforms, such as AR glasses. Currently, there is no dataset available to evaluate 3D object reconstruction using egocentric captured images. To address these gaps, the DTC dataset features 2,000 scanned digital twin-quality 3D objects, along with image sequences captured under different lighting conditions using DSLR cameras and egocentric AR glasses. This dataset establishes the first comprehensive real-world evaluation benchmark for 3D digital twin creation tasks, offering a robust foundation for comparing and improving existing reconstruction methods. The DTC dataset is already released at <https://www.projectaria.com/datasets/dtc/> and we will also make the baseline evaluations open-source.

* Authors contributed equally and are listed in alphabetical order.




Figure 2. Example DTC models with photorealistic PBR materials.

1. Introduction

A digital twin of a 3D object is a highly detailed, virtually indistinguishable representation of a physical object, capturing its shape, appearance, physical properties and other attributes with precision. Such a digital twin enables visualization, analysis, and interaction as if it were the real object, supporting simulation, automation, and real-world problem-solving across a wide range of applications in AR/VR [3, 28], spatial/contextual AI [2], and robotics [23, 43]. As fundamental properties of an object, its shape and appearance form the basis for recognizing and interpreting the 3D object, enabling identification, manipulation, and realistic rendering. Recovering these attributes has long been a foundational topic in computer vision and graphics, inspiring extensive research in 3D reconstruction and inverse rendering. Recent breakthroughs in neural-based representation and reconstruction techniques, such as NeRF [52] and 3D Gaussian splatting (3DGS) [37], have significantly elevated the quality of novel view synthesis (NVS) to photorealistic levels. Many subsequent works [35, 62] integrate neural reconstruction with physically-based inverse rendering, enabling relightable appearances. Furthermore, leveraging priors from large reconstruction models (LRMs) [30], high-quality shape and appearance reconstruction can now be achieved with as few as one to four views [41, 70, 84].

Despite the rapid advancements in 3D object reconstruction, one question remains: *does the reconstruction quality truly meet the standard of a digital twin, where virtual representations are **indistinguishable** from reality?* This digital twin standard demands both **highly accurate shape matching** and **photorealistic appearance across different lighting**, which present significant acquisition challenges for real-world objects. Existing object-centric datasets for 3D recon-

struction or inverse rendering have focused on either dataset size [18] or quality of specific aspects [16, 21, 34, 46, 61, 74], often sacrificing comprehensive fidelity and limiting their application scope. This trade-off has led to a lack of datasets that fully satisfy the digital twin criteria, hindering current 3D reconstruction methods from achieving digital twin fidelity. To bridge this gap, we developed the Digital Twin Catalog (DTC) dataset, comprising 2,000 scanned 3D object models (Fig. 1(a)), each with millimeter geometry accuracy and photorealistic PBR materials (Fig. 1(b), Fig. 2).

In addition to 3D digital twin models, the DTC dataset includes evaluation data designed to support 3D object reconstruction research. This evaluation data features multi-view image sequences with precise foreground object masks and environment lighting information for relighting evaluation. Traditionally, high-quality HDR images captured with modern DSLR cameras have been the standard for 3D reconstruction research. Looking ahead, we encourage the integration of 3D reconstruction research with next-generation human-centric computing platforms, such as ego-centric AR glasses, aiming to democratize 3D reconstruction techniques and empower everyone to effortlessly create 3D digital twins. To this end, alongside DSLR-captured evaluation data (Fig. 1(c)), the DTC dataset also provides ego-centric evaluation data captured using Project Aria glasses (<https://www.projectaria.com>) (Fig. 1(d)).

The DTC dataset offers extensive opportunities for advancing research in object digital twin creation. We provide a benchmark for state-of-the-art 3D object reconstruction and inverse rendering methods. These benchmarks evaluate performance across novel view synthesis (NVS), shape reconstruction, and relightable appearance reconstruction. We further provide the evaluation of novel view synthesis meth-

Table 1. Comparison with existing object-centric inverse rendering datasets. *Objaverse [18] consists of both synthetic objects and real scans, only part of which contain PBR materials.

Dataset	# Objects	Real	Scene Type	Multi-view	Shape	PBR Mat.	Relit Image	Lighting	Egocentric Cap.
ShapeNet-Intrinsics [61]	31K	✗	synthetic	✓	✓	✗	✓	✓	✗
NeRD Synthetic [11]	3	✗	synthetic	✓	✓	✓	✓	✓	✗
ABO [16]	8K	✗	synthetic	✓	✓	✓	✓	✓	✗
MIT Intrinsic [25]	20	✓	studio	✓	✗	✗	✗	✗	✗
DTU-MVS [34]	80	✓	studio	✓	✓	✗	✗	✗	✗
Objaverse [18]	818K	(✓)*	studio	✓	✓	(✓)*	✗	✗	✗
DiLiGenT-MV [42]	5	✓	studio	✓	✓	✗	✗	✓	✗
ReNe [65]	20	✓	studio	✓	✗	✗	✗	✓	✗
OpenIllumination [46]	64	✓	studio	✓	✗	✓	✓	✓	✗
GSO [21]	1030	✓	studio	✓	✓	✗	✗	✗	✗
Lombardi <i>et al.</i> [47]	6	✓	in-the-wild	✗	✓	✗	✓	✓	✗
NeRD Real [11]	4	✓	in-the-wild	✓	✗	✓	✓	✗	✗
NeROIC [38]	3	✓	in-the-wild	✓	✗	✓	✓	✗	✗
Oxholm <i>et al.</i> [56]	4	✓	in-the-wild	✓	✓	✗	✓	✓	✗
OmniObject3D [74]	6k	✓	in-the-wild	✓	✓	✗	✗	✗	✗
Stanford Orb [39]	14	✓	in-the-wild	✓	✓	✓	✓	✓	✗
DTC (ours)	2k	✓	in-the-wild	✓	✓	✓	✓	✓	✓

ods using the egocentric aligned DTC data. Additionally, we explore the dataset’s potential in downstream robotics applications by assessing its effectiveness in training robotic policies for pushing and grasping tasks in simulation. These benchmarks and applications provide valuable insights, highlight existing challenges, and uncover promising directions for future research in 3D digital twin creation.

2. Related Work

We provide a comparison of our DTC dataset to existing object-centric datasets in Table 1. We provide the largest 3D dataset with PBR materials and real world multi-view recordings with digital twin counterparts. We further provide digital twin aligned egocentric recordings, the first of their kind in the egocentric domain. We will discuss the related datasets and methods they can empower as follows.

3D Digital Twin Datasets Existing 3D digital twin datasets with PBR materials often serve as ground truth for evaluating inverse rendering results. Early efforts [6, 25] provide small-scale intrinsic image of real objects and do not provide shape or PBR material information. Synthetic datasets [11, 16, 45, 61, 71, 73] are widely used for evaluation but do not represent the complexity in a real world environment. For datasets that contain real objects, [18, 34, 42, 47, 56, 65], the reconstruction quality can vary, which leads to imprecise evaluations. Table 1 provides a comparison to the previous work in this domain. Compared to Objaverse [18, 19], which is a collection of existing 3D models with only a small subset containing PBR materials with varying quality, we offer a high quality collection of 3D object data that is also aligned with real world recordings. Compared to OmniObject3D [74], the DTC models

provide higher-quality shape and additional PBR materials that are necessary for high quality inverse rendering. We offer the largest quantity of 3D object models compared to all counterparts in various tasks. For real-world evaluation, Stanford-ORB [39] was the prior largest inverse rendering benchmark with in-the-wild lighting. In contrast, we provide more object diversity and higher quality for each object model. The Aria Digital Twin dataset [58] was the first dataset to provide digital twin aligned environments for the scenes and recorded using egocentric device. However, their scene environments are limited and the contained object ground truths inside do not have high quality geometries with PBR materials.

Object Reconstruction & Inverse Rendering. Using object-centric multi-view images as input, early object reconstruction methods focused on estimating individual object properties, such as shape from shading [5, 9, 85], material acquisition [44, 45, 56, 57, 75], and lighting estimation [68, 80]. Some approaches also aimed to recover reflectance and illumination assuming known object shapes [47, 48]. Inverse rendering, which seeks to invert the rendering equation [36], estimates an image’s intrinsic components—geometry, materials, and lighting. The advent of differentiable renderers [14, 15, 49] enabled full-fledged inverse rendering methods to simultaneously recover all these properties for object reconstruction [50].

Neural volumetric representations such as Neural Radiance Fields (NeRFs) [52] and the like [7, 11–13, 59, 81, 87] encode geometry and appearance as volumetric densities and radiance with a Multi-Layer Perceptron (MLP) network, and render images using the volume rendering equation [51]. 3D-GS [37] introduces 3D Gaussian primitives and rasterization and its following-up variants [31] demonstrates high quality

geometry prediction as well.

Other surface-based representations [29, 53, 62, 66, 72, 82, 83, 88] extract surfaces as the zero level set, for instance, of a signed distance function (SDF) or an occupancy field [55], allowing them to efficiently model the appearance on the surface with an explicit material model, such as bidirectional reflectance distribution functions (BRDFs). This also enables modeling more complex global illumination effects, such as self-shadows. Most of these methods focus on per-scene optimization and require dense multiple views as input. Recently, researchers have incorporated learning-based models, distilling priors from large training datasets for fast inference on limited test views [8, 10, 33, 45, 61, 71, 73, 89].

In this work, we provide the DTC dataset with real-world object recordings that can serve as the benchmark to evaluate object-centric inverse rendering tasks. We evaluate representative baselines from existing work.

3. Digital Twin Catalog: A Large-Scale Photorealistic 3D Object Digital Twin Dataset

3.1. Dataset Composition

Our DTC dataset contains: (1) **2,000** scanned 3D object models, featuring millimeter geometric accuracy relative to their physical counterparts, along with a full set of photorealistic PBR material maps (Fig. 1), (2) **100** DSLR-captured evaluation data of **50** objects under different lighting conditions, and (3) **200** egocentric Aria-captured evaluation data of **100** objects with both **active** and **casual** observation modes.

3.2. Creation of 3D Object Models

Utilizing the state-of-the-art industrial 3D object scanner [1], we selected 2,000 physical-world objects spanning 40 LVIS [27] categories, carefully chosen to ensure both category diversity and compatibility with the scanner’s capabilities. As illustrated in Fig. 3(a) (c), the scanner features a fixed lighting-camera setup within an upper-hemisphere dome, equipped with 8 structured lights for geometry scanning, and 29 spotlights and 29 cameras for material acquisition. During the scanning process, the object is placed on a central holder, and its pose can be adjusted with multi-round scanning to achieve a complete 360-degree scan. For our dataset, each object typically undergoes three pose changes, with a total scanning time of approximately 20 minutes per object.

After scanning, a proprietary post-processing pipeline reconstructs both the geometry and the PBR material maps. For 4K-resolution PBR material maps, the post-processing requires approximately 4 hours per object. In terms of quality, the structured-light-based shape reconstruction (Fig. 3 (b)) in the post-processing achieves millimeter-level geometric accuracy. However, the material optimization process




Figure 3. 3D object scanner by Covision Media[®].



Figure 4. Rendered DTC models (left) v.s. Photo (Right).

performs best for diffuse objects and often struggles with glossy or shiny surfaces. To address this limitation, we hired technical artists to develop a workflow to refine materials for glossy and shiny objects, ensuring that the material quality meets the standards of a digital twin.




Figure 5. Shape and material (albedo) quality comparison between Stanford-ORB [39] (left) and our DTC (right).

3D Model Accuracy. To validate the material and geometry accuracy of the 3D models, we compared a rendered image of our scanned and processed model with a photograph of the same object taken inside a light box. A virtual light box was meticulously modeled to replicate the light intensity and color temperature of the real light box. The scanned object was then placed in the virtual light box to generate the rendered image. The side-by-side comparison demonstrates a remarkable match between the rendered and real images (Fig. 4).

Comparison Against Stanford-ORB. We also scanned the objects used in Stanford-ORB [39] to compare the shape and appearance quality. As illustrated in Fig. 5, the Stanford-ORB models exhibit shape artifacts and noisy, lower-quality materials compared to our models.

3.3. DSLR Evaluation Data

Within DTC, we include a DSLR-captured evaluation dataset of 50 objects from Sec. 3.2 captured under two different lighting conditions, resulting in 100 distinct image sequences. For every sequence, we provide (a) approximately 120 HDR and LDR images from different viewing directions, (b) one object pose and (c) per-image camera pose. The two lighting conditions are represented using two environment maps.

Data Capture. To ensure the DSLR evaluation data quality, we designed and built a DSLR camera rig to automate the capture process (Fig. 6). The rig is designed to rotate the cameras around the centralized object, assuming the environment lighting remains unchanged during the capture. It features a motorized rotary stage with a centrally mounted stationary platform. Attached to the rotary stage is an extrusion frame that forms the gantry arm, supported by a set of castors to bear its weight and enable smooth rotation around the central axis. The extrusion frame is equipped with adjustable camera mounts, allowing DSLR cameras to be positioned flexibly to optimize the capture setup. For our capture process, we utilized three DSLR cameras to perform a 360-degree rotation around the object, capturing images at 9-degree intervals, resulting in 120 photos per object. To ensure precise camera pose estimation, a ChArUco board




Figure 6. DSLR rig for capturing evaluation data.

was placed beneath the object during the capture. Example images from this setup are shown in Fig. 6.

Environment Maps. Following a similar approach in Stanford-ORB [39], we capture the two environment maps using chrome ball images obtained with the capture rig described earlier. With precise camera poses provided by a ChArUco board placed beneath the chrome ball, we first fit a synthetic 3D sphere to the chrome ball by optimizing its 3D position using a geometry-friendly differentiable renderer [40, 76]. Subsequently, using a differentiable Monte Carlo-based renderer [32], we refine the environment map to match the reflection on the chrome ball, employing the single-view light estimation method proposed in [80]. The coordinate system of the environment map is determined by the ChArUco detection.

Pose Registration for Camera and Object. We first obtain the initial camera poses using the ChArUco board. In most cases, for images captured from the top and middle views, the pose estimates are typically accurate. However, for bottom-view images, inaccuracies arise due to pattern distortion at grazing angles. To mitigate this issue, we refine the camera poses by fitting the rendering of a virtual ChArUco board to the real captured photo using a differentiable renderer [32, 40, 76]. Once the camera poses have been accurately refined, we optimize the object pose by minimizing the mask loss between the rendered mask and the reference one generated by [62]. By providing separate poses for camera and object anchored by ChArUco, the cameras and object are automatically aligned with the optimized environment maps.

3.4. Egocentric Evaluation Data

We include an evaluation dataset of real world recordings captured by an egocentric device paired with 100 objects from the DTC dataset. We capture the egocentric recordings using the open-sourced Project Aria device [22] and acquire the additional 3D ground truth using its machine perception tool, which includes online device calibration, device trajectory and semi-dense point clouds for each recording. Each recording contains a single 3D object and precisely aligned object poses in the Aria trajectory coordinate frame.

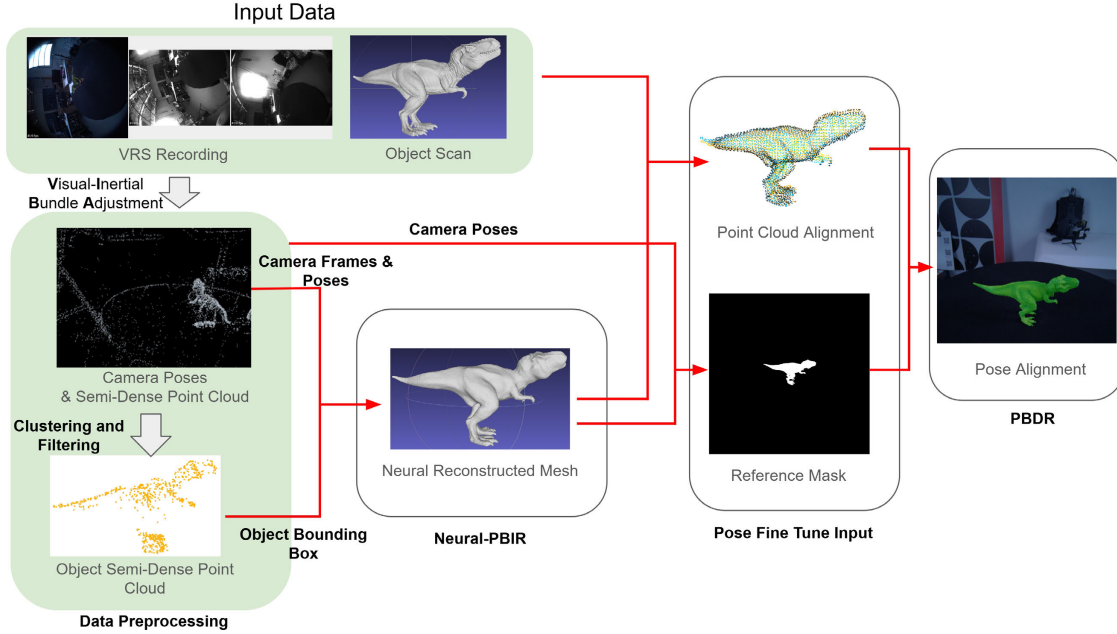


Figure 7. The workflow to align egocentric video and 3D objects. We acquire the object’s semi-dense point cloud from the egocentric recording and the neural reconstructed mesh using Neural-PBIR [62]. Finally, we align the 3D object with the rendered mask from the neural reconstructed mesh using physics-based differentiable rendering (PBDR).

Both the object and the trajectory are in metric scale. We can derive additional ground truth for each video from the aligned 3D object properties.

Data Capture. To feature real world recordings observed from human perspectives that can be representative of 3D object reconstruction for AR/MR devices or robots, we provide two types of recording trajectories for selected objects, termed *active* and *passive* respectively, collected by human wearers. The active recording features a complete 360 view of the objects, which is similar to existing object 3D novel view synthesis dataset. The *passive* recording features casual looks from human wearer, which only contain partial views of the object from certain viewing angles along the trajectory. In dataset creation, we collected the active and passive recordings in the same environment and generate their 3D information in the shared 3D space. This helps reduce potential failures when aligning the object to passive 3D recordings, which are shorter and contains less 3D information. To reduce the effect of noises and motion blur, which are common in egocentric videos in indoor low light environments, we light the capture environment to 3K+ lux illumination and used a fixed low-exposure and gain profile to collect each recording with the appropriate brightness.

Alignment between Egocentric Video and Object. We provide an illustration of the object alignment to the egocentric video in Fig. 7. Given the images, camera poses and semi-dense point cloud acquired from Project Aria tools, we employ a neural-based mesh reconstruction method [62] to

create reference meshes for generating high-quality reference masks and used that to align with the corresponding 3D object mesh. This alignment step serves as an initialization for a more precise pose refinement, which leverages differentiable rendering. For certain objects with symmetric geometry, we observe this process can introduce ambiguities in point cloud registration and subsequent failures. To address such cases, we provide a GUI to manually align and correct the object alignment. Finally, akin to the DSLR camera pose registration phase, we optimize a mask loss over object poses to achieve fine-tuned pose registration.

4. Benchmarking and Applications

We first use our DLSR and egocentric dataset as a benchmark for existing state-of-the-art methods. For inverse rendering, we design metrics to evaluate the shape and material quality of the recovered 3D object digital twin from three perspectives. For egocentric recording, we evaluate the novel-view synthesis as the initial evaluation. We include additional tasks, e.g. sparse view reconstruction in the supplementary materials for both DSLR and egocentric recordings. Finally we can demonstrate our high quality 3D digital twin models can be beneficial to robotics domain using an application in robotics manipulation.

4.1. Application to Inverse Rendering for DSLR

The DSLR dataset in DTC provides accurate ground truth, including poses, lighting and 3D models, for inverse rendering

Table 2. Benchmark comparison of existing methods on inverse rendering for DSLR. Depth SI-MSE and Shape Chamfer distance $\times 10^{-3}$.

	Geometry			Novel Scene Relighting				Novel View Synthesis			
	Depth↓	Normal↓	Shape↓	PSNR-H↑	PSNR-L↑	SSIM↑	LPIPS↓	PSNR-H↑	PSNR-L↑	SSIM↑	LPIPS↓
Neural-PIL [12]	5.71	0.25	25.02			N/A		28.42	35.76	0.882	0.096
PhySG [82]	0.31	0.16	11.31	27.28	32.86	0.959	0.049	28.54	34.46	0.964	0.045
NVDiffRec [53]	0.02	0.07	1.64	26.99	33.27	0.951	0.037	28.95	34.92	0.967	0.029
NeRD [11]	4.55	0.45	108.20	26.10	32.60	0.948	0.061	26.80	33.40	0.882	0.102
InvRender [72]	0.22	0.03	0.75	29.52	35.98	0.961	0.037	31.64	37.82	0.970	0.033
NVDiffRecMC [29]	0.02	0.06	1.34	27.78	34.55	0.952	0.042	31.27	38.17	0.972	0.032

Table 3. Benchmark on the egocentric aligned recordings.

	PSNR ↑	LPIPS ↓	SSIM ↑	Depth ↓	Normal ↓
3D-GS [37]	28.81	0.020	0.9888	0.1768	0.3301
2D-GS [31]	28.75	0.020	0.9886	0.1755	0.2112

tasks and serves as an evaluation suite to benchmark the performance of inverse rendering methods. We select six prior methods for this task and evaluate their performance using the ground truth provided by our dataset. In the following sections, we describe the data splitting strategy, evaluation metrics, and baselines.

Data Splitting. For benchmarking purposes, we select 15 objects from the DSLR dataset captured under two distinct lighting environments, resulting in a total of 30 image sequences. The selected objects encompass a diverse range of geometric and material properties to ensure a comprehensive evaluation. For each scene, 8 views are selected for testing, while the remaining views are reserved for training.

Evaluation Metrics. The metrics measure the accuracy of three aspects of baseline performance: geometry estimation, relighting, and novel view synthesis. For geometry estimation, we evaluate the accuracy of predicted depth and normal maps under held-out test views, as well as 3D meshes extracted from baseline methods, compared with the ground truth from our dataset. Relighting metrics evaluate the material decomposition quality of baselines by measuring the accuracy of predicted images under held-out lighting conditions. For view synthesis, we compare the predicted images from viewpoints unseen during training to ground truth captures. We refer to Kuang et al. [39] for metric details.

Baselines. We include the following baselines: NVDiffRec [53] and NVDiffRecMC [29], with a hybrid shape representation DMTet [60]; InvRender [88] and PhySG [82], which adopt signed distance functions (SDFs) to represent object geometry [77] and utilize implicit neural fields for material decomposition; Neural-PIL [12] and NeRD [11], which use NeRFs [52] as scene representations.

4.2. Application to Egocentric Reconstruction

Our digital twin models, aligned with real world video using the method described in Sec. 3.4, can help obtain accurate

ground truth for object-centric images that were previously difficult to acquire. We provide the first evaluation of object-centric novel view synthesis recorded from an egocentric device. We use the projected object shape given the 3D pose of the object and cameras in scene coordinates to acquire the image masks, depth and normal for each object. We selected 15 recordings from the egocentric recording sessions as the evaluation and used the *active* recordings to benchmark novel view reconstruction. For each recording, we hold out every 8th image view as a testing view.

Evaluations. We build our baselines based on the gsplats [79] implementation of the 3D Gaussian Splatting (GS) [37] and 2D GS [31], and handle the effect of lens shading from the Project Aria lens [26]. We calculate PSNR, depth and normal based on the observed objects with masks and provide SSIM and LPIPS score on images by masking out the non-object areas as black. Table 3 shows the benchmark results of the baselines. We use the same depth and normal metric in DSLR evaluation. We provide additional qualitative evaluations and analyses on egocentric data towards sparser view settings in the supplementary materials.

4.3. Application to Robotic Manipulation

High-quality object models have been leveraged in prior work to train real-world robotic agents in scenes represented explicitly [54] or implicitly [64]. These object models have also been shown to facilitate object-centric pose and lighting parameter estimation, enabling model-based planning [63]. In this section, we empirically evaluate the effectiveness of using DTC dataset objects in training robotic policies. Specifically, we consider learning robotic pushing and grasping skills in simulation.

First, we sample a subset of 24 cup category objects from the DTC dataset and 24 cup objects from Objaverse-XL [19]¹. Since not all Objaverse-XL objects come with textures, we randomize the colors of those objects uniformly in RGB space. To compute collision meshes for physical simulation, we perform convex decomposition on each object with CoACD [69]. We import these objects along with

¹The version of Objaverse-XL used in this work excludes all 3D models sourced from Sketchfab. Further, no Polycam assets were obtained from the Polycam source site

Task	DTC (ours)	Objaverse-XL [19]
Pushing @ 2cm	36.3% \pm 1.5%	25.3% \pm 6.0%
Pushing @ 3cm	43.7% \pm 1.2%	29.7% \pm 6.0%
Pushing @ 5cm	47.0% \pm 2.6%	40.3% \pm 5.5%
Grasping	42.7% \pm 4.7%	38.6% \pm 11.0%

Table 4. Success rate of policies trained on data collected using objects from our DTC dataset and sampled from Objaverse-XL when evaluated on an unseen test object. Errors indicate sample standard deviation over three policy training seeds.

a UR5e robot equipped with a Robotiq 2F-85 (pushing) or Robotiq 2F-140 (grasping) gripper into the PyBullet simulator [17] and collect data for each robotic task as described below. After training policies on data from each object set, we evaluate policy performance on a relatively high-quality unseen test object from the StanfordORB dataset [39].

Pushing. For the pushing task, we collect 5000 trajectories of pseudo-random robotic interaction data for each object set. For each trajectory, a single object from the considered object set and its initial position are randomly selected. Then we train a goal-conditioned neural network policy $\pi(a|o, o_g)$ where o is an image observation of the current scene and o_g is a goal image indicating the desired final object and robot position, $o, o_g \in \mathbb{R}^{256 \times 256 \times 3}$, and the action $a \in \mathbb{R}^2$ represents a change in the robot’s end-effector position in the x and y axes. The z axis end-effector height is held fixed. We sample goals for training via hindsight relabeling [4, 20, 24]. We then perform evaluation on 100 randomly sampled test goals manipulating an unseen test cup from StanfordORB [39].

Grasping. For grasping, we collect 5000 successful grasp examples for each object set by first placing a single object into the scene, randomizing the object identity and initial position. We then randomly sample candidate grasp poses in a radius around the object’s position and simulating their outcomes, rejecting unsuccessful grasps. We train a grasping policy $\pi(a|o)$ where $o \in \mathbb{R}^{256 \times 256 \times 3}$ is an image observation of the scene and $a \in \mathbb{R}^4$ represents the x, y, z position and θ yaw rotation of the robot end-effector at which to attempt the grasp. Again we use 100 test cup object poses.

Results. We report the results in Table 4 and Fig. 8. We find that across both pushing and grasping tasks, policies trained on DTC dataset objects outperform those trained on Objaverse-XL objects when evaluated on the unseen test object. For pushing, we report performance by defining binary success thresholds based on the final Euclidean distance of the object position to the goal position. Training on DTC objects appears to be especially helpful at enabling policies to make finer adjustments, improving pushing success rates at stricter thresholds. Additional experimental details can be found in the supplementary.

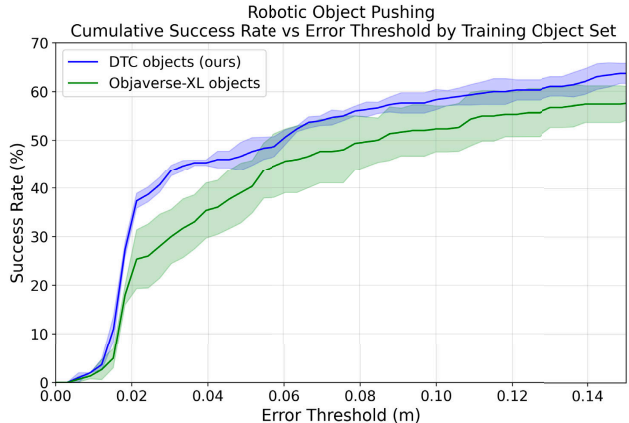


Figure 8. Success rates on robotic pushing task when training with our DTC objects and sampled Objaverse-XL objects. Particularly at lower error thresholds, policies trained on DTC objects outperform those trained using Objaverse-XL objects. Shaded bars represent sample standard deviations over policy training random seeds.

5. Conclusion

We presented a new large scale photorealistic 3D digital twin dataset with the real world recordings that contain its real world counterpart. We provide extensive evaluations of baselines on our DTC dataset serving as new benchmark for inverse rendering and novel view synthesis task. We also demonstrated that high quality digital twin models can be beneficial to applications in robotics domain. We believe our efforts can empower the research community to build and leverage digital twin models for future applications.

Limitations. Achieving high quality digital twin models currently requires deliberate hardware setup and human efforts in refinement. Solving this challenge without sacrificing quality can significantly further enhance the volume of digital twin models. Our hardware is also limited to objects within a certain size and can not yet recover objects that are deformable, highly specular, or transparent.

Future work. The existing digital twin model creation in DTC dataset involves lengthy post-processing and may require subjective human refinement, hindering the automation of model generation. However, recent advancements in physics-based differentiable rendering hold promise for enabling faster and more accurate creation of digital twins, especially for material reconstruction. Furthermore, building large-scale digital twins for applications will necessitate efforts to enhance the diversity in object appearance (e.g., transparent objects) and to capture additional attributes, such as physical properties and functionalities.

For robotics applications, while deploying manipulation policies learned in simulation to the real world remains generally challenging, we hope that the high-quality digital twin data provided by DTC can serve as a stepping stone towards effective sim-to-real transfer.

Acknowledgments. This work is in part supported by NSF CCRI #2120095, RI #2211258, and RI #2338203. YZ is in part supported by the Stanford Interdisciplinary Graduate Fellowship.

References

- [1] Covision 3D object scanner. <https://www.covisionmedia.ai/en/scanners>. 4
- [2] Kosmas Alexopoulos, Nikolaos Nikolakis, and George Chrysolouris. Digital twin-driven supervised machine learning for the development of artificial intelligence applications in manufacturing. *International Journal of Computer Integrated Manufacturing*, 33(5):429–439, 2020. 2
- [3] Moayad Aloqaily, Ouns Bouachir, Fakhri Karray, Ismaeel Al Ridhawi, and Abdulmotaleb El Saddik. Integrating digital twin and advanced intelligent technologies to realize the metaverse. *IEEE Consumer Electronics Magazine*, 12(6):47–55, 2022. 2
- [4] Marcin Andrychowicz, Filip Wolski, Alex Ray, Jonas Schneider, Rachel Fong, Peter Welinder, Bob McGrew, Josh Tobin, OpenAI Pieter Abbeel, and Wojciech Zaremba. Hindsight experience replay. In *NeurIPS*, 2017. 8
- [5] Jonathan T Barron and Jitendra Malik. Shape, illumination, and reflectance from shading. *IEEE TPAMI*, 37(8):1670–1687, 2014. 3
- [6] Sean Bell, Kavita Bala, and Noah Snavely. Intrinsic images in the wild. *ACM TOG*, 33(4):1–12, 2014. 3
- [7] Sai Bi, Zexiang Xu, Kalyan Sunkavalli, Miloš Hašan, Yannick Hold-Geoffroy, David Kriegman, and Ravi Ramamoorthi. Deep reflectance volumes: Relightable reconstructions from multi-view photometric images. In *ECCV*, 2020. 3
- [8] Sai Bi, Zexiang Xu, Kalyan Sunkavalli, David Kriegman, and Ravi Ramamoorthi. Deep 3D capture: Geometry and reflectance from sparse multi-view images. In *CVPR*, 2020. 4
- [9] Martin Bichsel and Alex P Pentland. A simple algorithm for shape from shading. In *CVPR*, 1992. 3
- [10] Mark Boss, Varun Jampani, Kihwan Kim, Hendrik P.A. Lensch, and Jan Kautz. Two-shot Spatially-varying BRDF and Shape Estimation. In *CVPR*, 2020. 4
- [11] Mark Boss, Raphael Braun, Varun Jampani, Jonathan T. Barron, Ce Liu, and Hendrik P.A. Lensch. NeRD: Neural reflectance decomposition from image collections. In *ICCV*, 2021. 3, 7
- [12] Mark Boss, Varun Jampani, Raphael Braun, Ce Liu, Jonathan Barron, and Hendrik Lensch. Neural-PIL: Neural pre-integrated lighting for reflectance decomposition. In *NeurIPS*, 2021. 7
- [13] Mark Boss, Andreas Engelhardt, Abhishek Kar, Yuanzhen Li, Deqing Sun, Jonathan Barron, Hendrik Lensch, and Varun Jampani. SAMURAI: Shape and material from unconstrained real-world arbitrary image collections. In *NeurIPS*, 2022. 3
- [14] Guangyan Cai, Kai Yan, Zhao Dong, Ioannis Gkioulekas, and Shuang Zhao. Physics-based inverse rendering using combined implicit and explicit geometries. *CGF*, 41(4):129–138, 2022. 3
- [15] Wenzheng Chen, Joey Litalien, Jun Gao, Zian Wang, Clement Fuji Tsang, Sameh Khamis, Or Litany, and Sanja Fidler. DIB-R++: learning to predict lighting and material with a hybrid differentiable renderer. In *NeurIPS*, 2021. 3
- [16] Jasmine Collins, Shubham Goel, Kenan Deng, Achleshwar Luthra, Leon Xu, Erhan Gundogdu, Xi Zhang, Tomas F Yago Vicente, Thomas Dideriksen, Himanshu Arora, et al. ABO: Dataset and benchmarks for real-world 3D object understanding. In *CVPR*, 2022. 2, 3
- [17] Erwin Coumans and Yunfei Bai. Pybullet, a python module for physics simulation for games, robotics and machine learning. <http://pybullet.org>, 2016–2019. 8
- [18] 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, 3
- [19] Matt Deitke, Ruoshi Liu, Matthew Wallingford, Huong Ngo, Oscar Michel, Aditya Kusupati, Alan Fan, Christian Laforte, Vikram Voleti, Samir Yitzhak Gadre, et al. Objaverse-XL: A universe of 10m+ 3D objects. *NeurIPS*, 2024. 3, 7, 8
- [20] Yiming Ding, Carlos Florensa, Pieter Abbeel, and Mariano Phielipp. Goal-conditioned imitation learning. In *NeurIPS*, 2019. 8
- [21] Laura Downs, Anthony Francis, Nate Koenig, Brandon Kinman, Ryan Hickman, Krista Reymann, Thomas B McHugh, and Vincent Vanhoucke. Google scanned objects: A high-quality dataset of 3D scanned household items. In *ICRA*, 2022. 2, 3, 10
- [22] Jakob Engel, Kiran Somasundaram, Michael Goesele, Albert Sun, Alexander Gamino, Andrew Turner, Arjang Talatof, Arnie Yuan, Bilal Souti, Brighid Meredith, Cheng Peng, Chris Sweeney, Cole Wilson, Dan Barnes, Daniel DeTone, David Caruso, Derek Valleroy, Dinesh Ginjupalli, Duncan Frost, Edward Miller, Elias Mueggler, Evgeniy Oleinik, Fan Zhang, Guruprasad Somasundaram, Gustavo Solaira, Harry Lanaras, Henry Howard-Jenkins, Huixuan Tang, Hyo Jin Kim, Jaime Rivera, Ji Luo, Jing Dong, Julian Straub, Kevin Bailey, Kevin Eickenhoff, Lingni Ma, Luis Pesqueira, Mark Schwesinger, Maurizio Monge, Nan Yang, Nick Charron, Nikhil Raina, Omkar Parkhi, Peter Borschowa, Pierre Moulon, Prince Gupta, Raul Mur-Artal, Robbie Pennington, Sachin Kulkarni, Sagar Miglani, Santosh Gondi, Saransh Solanki, Sean Diener, Shangyi Cheng, Simon Green, Steve Saarinen, Suvam Patra, Tassos Mourikis, Thomas Whelan, Tripti Singh, Vasileios Balntas, Vijay Baiyya, Wilson Dreewes, Xiaqing Pan, Yang Lou, Yipu Zhao, Yusuf Mansour, Yuyang Zou, Zhaoyang Lv, Zijian Wang, Mingfei Yan, Carl Ren, Renzo De Nardi, and Richard Newcombe. Project Aria: A new tool for egocentric multi-modal ai research, 2023. 5
- [23] Ruohan Gao, Yen-Yu Chang, Shivani Mall, Li Fei-Fei, and Jiajun Wu. ObjectFolder: A dataset of objects with implicit visual, auditory, and tactile representations. In *CoRL*, 2021. 2
- [24] Dibya Ghosh, Abhishek Gupta, Ashwin Reddy, Justin Fu, Coline Devin, Benjamin Eysenbach, and Sergey Levine. Learning to reach goals via iterated supervised learning. In *ICML*, 2019. 8
- [25] Roger Grosse, Micah K Johnson, Edward H Adelson, and

- William T Freeman. Ground truth dataset and baseline evaluations for intrinsic image algorithms. In *ICCV*, 2009. 3
- [26] Qiao Gu, Zhaoyang Lv, Duncan Frost, Simon Green, Julian Straub, and Chris Sweeney. Egolifter: Open-world 3d segmentation for egocentric perception. In *ECCV*, pages 382–400. Springer, 2024. 7
- [27] Agrim Gupta, Piotr Dollar, and Ross Girshick. LVIS: A dataset for large vocabulary instance segmentation. In *CVPR*, 2019. 4, 1
- [28] Wassim Hamidouche, Lina Bariah, and Mérouane Debbah. Immersive media and massive twinning: Advancing toward the metaverse. *IEEE Communications Magazine*, 62(7):20–32, 2024. 2
- [29] Jon Hasselgren, Nikolai Hofmann, and Jacob Munkberg. Shape, light & material decomposition from images using Monte Carlo rendering and denoising. In *NeurIPS*, 2022. 4, 7
- [30] Yicong Hong, Kai Zhang, Jiuxiang Gu, Sai Bi, Yang Zhou, Difan Liu, Feng Liu, Kalyan Sunkavalli, Trung Bui, and Hao Tan. LRM: Large reconstruction model for single image to 3D. In *The Twelfth International Conference on Learning Representations*, 2024. 2
- [31] Binbin Huang, Zehao Yu, Anpei Chen, Andreas Geiger, and Shenghua Gao. 2D gaussian splatting for geometrically accurate radiance fields. In *SIGGRAPH 2024 Conference Papers*, 2024. 3, 7, 11, 12
- [32] Wenzel Jakob, Sébastien Speierer, Nicolas Roussel, Merlin Nimier-David, Delio Vicini, Tizian Zeltner, Baptiste Nicolet, Miguel Crespo, Vincent Leroy, and Ziyi Zhang. Mitsuba 3 renderer, 2022. <https://mitsuba-renderer.org>. 5
- [33] Michael Janner, Jiajun Wu, Tejas Kulkarni, Ilker Yildirim, and Joshua B Tenenbaum. Self-supervised intrinsic image decomposition. In *NeurIPS*, 2017. 4
- [34] Rasmus Jensen, Anders Dahl, George Vogiatzis, Engin Tola, and Henrik Aanæs. Large scale multi-view stereopsis evaluation. In *CVPR*, 2014. 2, 3
- [35] Haian Jin, Isabella Liu, Peijia Xu, Xiaoshuai Zhang, Songfang Han, Sai Bi, Xiaowei Zhou, Zexiang Xu, and Hao Su. TensorIR: Tensorial inverse rendering. In *CVPR*, 2023. 2
- [36] James T Kajiya. The rendering equation. In *Proceedings of the 13th annual conference on Computer graphics and interactive techniques*, pages 143–150, 1986. 3
- [37] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3D gaussian splatting for real-time radiance field rendering. *ACM TOG*, 2023. 2, 3, 7, 11, 12
- [38] Zhengfei Kuang, Kyle Olszewski, Menglei Chai, Zeng Huang, Panos Achlioptas, and Sergey Tulyakov. NeROIC: neural rendering of objects from online image collections. *ACM TOG*, 41(4):1–12, 2022. 3
- [39] Zhengfei Kuang, Yunzhi Zhang, Hong-Xing Yu, Samir Agarwala, Elliott Wu, Jiajun Wu, et al. Stanford-orb: a real-world 3D object inverse rendering benchmark. In *Advances in Neural Information Processing Systems Datasets and Benchmarks Track*, 2023. 3, 5, 7, 8, 9
- [40] Samuli Laine, Janne Hellsten, Tero Karras, Yeongho Seol, Jaakko Lehtinen, and Timo Aila. Modular primitives for high-performance differentiable rendering. *ACM Transactions on Graphics*, 39(6), 2020. 5
- [41] Jiahao Li, Hao Tan, Kai Zhang, Zexiang Xu, Fujun Luan, Yinghao Xu, Yicong Hong, Kalyan Sunkavalli, Greg Shakhnarovich, and Sai Bi. Instant3D: Fast text-to-3D with sparse-view generation and large reconstruction model. In *ICLR*, 2024. 2
- [42] Min Li, Zhenglong Zhou, Zhe Wu, Boxin Shi, Changyu Diao, and Ping Tan. Multi-view photometric stereo: A robust solution and benchmark dataset for spatially varying isotropic materials. *IEEE TIP*, 29:4159–4173, 2020. 3
- [43] Xuanlin Li, Kyle Hsu, Jiayuan Gu, Karl Pertsch, Oier Mees, Homer Rich Walke, Chuyuan Fu, Ishikaa Lunawat, Isabel Sieh, Sean Kirmani, Sergey Levine, Jiajun Wu, Chelsea Finn, Hao Su, Quan Vuong, and Ted Xiao. Evaluating real-world robot manipulation policies in simulation. *arXiv preprint arXiv:2405.05941*, 2024. 2
- [44] Zhengqin Li, Kalyan Sunkavalli, and Manmohan Chandraker. Materials for masses: SVBRDF acquisition with a single mobile phone image. In *ECCV*, 2018. 3
- [45] Zhengqin Li, Zexiang Xu, Ravi Ramamoorthi, Kalyan Sunkavalli, and Manmohan Chandraker. Learning to reconstruct shape and spatially-varying reflectance from a single image. *ACM TOG*, 37(6):1–11, 2018. 3, 4
- [46] Isabella Liu, Linghao Chen, Ziyang Fu, Liwen Wu, Haian Jin, Zhong Li, Chin Ming Ryan Wong, Yi Xu, Ravi Ramamoorthi, Zexiang Xu, and Hao Su. OpenIllumination: A multi-illumination dataset for inverse rendering evaluation on real objects. *NeurIPS*, 2023. 2, 3
- [47] Stephen Lombardi and Ko Nishino. Reflectance and natural illumination from a single image. In *ECCV*, 2012. 3
- [48] Stephen Lombardi and Ko Nishino. Reflectance and illumination recovery in the wild. *IEEE TPAMI*, 2015. 3
- [49] Fujun Luan, Shuang Zhao, Kavita Bala, and Zhao Dong. Unified shape and SVBRDF recovery using differentiable Monte Carlo rendering. *CGF*, 40(4):101–113, 2021. 3
- [50] Stephen Robert Marschner. *Inverse Rendering for Computer Graphics*. Cornell University, 1998. 3
- [51] Nelson Max. Optical models for direct volume rendering. *IEEE TVCG*, 1(2):99–108, 1995. 3
- [52] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. NeRF: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, 2020. 2, 3, 7, 9
- [53] Jacob Munkberg, Jon Hasselgren, Tianchang Shen, Jun Gao, Wenzheng Chen, Alex Evans, Thomas Müller, and Sanja Fidler. Extracting triangular 3D models, materials, and lighting from images. In *CVPR*, 2022. 4, 7
- [54] Soroush Nasiriany, Abhiram Maddukuri, Lance Zhang, Adeet Parikh, Aaron Lo, Abhishek Joshi, Ajay Mandlekar, and Yuke Zhu. Robocasa: Large-scale simulation of everyday tasks for generalist robots. In *Robotics: Science and Systems (RSS)*, 2024. 7
- [55] Michael Niemeyer, Lars Mescheder, Michael Oechsle, and Andreas Geiger. Differentiable volumetric rendering: Learning implicit 3D representations without 3D supervision. In *CVPR*, 2020. 4
- [56] Geoffrey Oxholm and Ko Nishino. Multiview shape and reflectance from natural illumination. In *CVPR*, 2014. 3

- [57] Geoffrey Oxholm and Ko Nishino. Shape and reflectance estimation in the wild. *IEEE TPAMI*, 2015. 3
- [58] Xiaqing Pan, Nicholas Charron, Yongqian Yang, Scott Peters, Thomas Whelan, Chen Kong, Omkar Parkhi, Richard Newcombe, and Yuheng (Carl) Ren. Aria digital twin: A new benchmark dataset for egocentric 3D machine perception. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 20133–20143, 2023. 3
- [59] Viktor Rudnev, Mohamed Elgharib, William Smith, Lingjie Liu, Vladislav Golyanik, and Christian Theobalt. NeRF for outdoor scene relighting. In *ECCV*, pages 615–631. Springer, 2022. 3
- [60] Tianchang Shen, Jun Gao, Kangxue Yin, Ming-Yu Liu, and Sanja Fidler. Deep marching tetrahedra: a hybrid representation for high-resolution 3D shape synthesis. In *NeurIPS*, 2021. 7
- [61] Jian Shi, Yue Dong, Hao Su, and Stella X Yu. Learning non-Lambertian object intrinsics across ShapeNet categories. In *CVPR*, 2017. 2, 3, 4
- [62] Cheng Sun, Guangyan Cai, Zhengqin Li, Kai Yan, Cheng Zhang, Carl Marshall, Jia-Bin Huang, Shuang Zhao, and Zhao Dong. Neural-PBIR reconstruction of shape, material, and illumination. In *ICCV*, 2023. 2, 4, 5, 6
- [63] Stephen Tian, Yancheng Cai, Hong-Xing Yu, Sergey Zakharov, Katherine Liu, Adrien Gaidon, Yunzhu Li, and Jiajun Wu. Multi-object manipulation via object-centric neural scattering functions. In *CVPR*, 2023. 7
- [64] Marcel Torne, Anthony Simeonov, Zechu Li, April Chan, Tao Chen, Abhishek Gupta, and Pulkit Agrawal. Reconciling reality through simulation: A real-to-sim-to-real approach for robust manipulation. *Arxiv*, 2024. 7
- [65] Marco Toschi, Riccardo De Matteo, Riccardo Spezialetti, Daniele De Gregorio, Luigi Di Stefano, and Samuele Salti. Relight my NeRF: A dataset for novel view synthesis and relighting of real world objects. In *CVPR*, 2023. 3
- [66] Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. NeuS: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. In *NeurIPS*, 2021. 4
- [67] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE TIP*, 13(4):600–612, 2004. 9
- [68] Henrique Weber, Donald Prévost, and Jean-François Lalonde. Learning to estimate indoor lighting from 3D objects. In *3DV*, 2018. 3
- [69] Xinyue Wei, Minghua Liu, Zhan Ling, and Hao Su. Approximate convex decomposition for 3D meshes with collision-aware concavity and tree search. *ACM TOG*, 2022. 7, 13
- [70] Xinyue Wei, Kai Zhang, Sai Bi, Hao Tan, Fujun Luan, Valentin Deschaintre, Kalyan Sunkavalli, Hao Su, and Zexiang Xu. MeshLRM: Large reconstruction model for high-quality mesh. *arXiv preprint arXiv:2404.12385*, 2024. 2, 10, 11
- [71] Felix Wimbauer, Shangzhe Wu, and Christian Rupprecht. De-rendering 3D objects in the wild. In *CVPR*, 2022. 3, 4
- [72] Haoqian Wu, Zhipeng Hu, Lincheng Li, Yongqiang Zhang, Changjie Fan, and Xin Yu. NeFII: Inverse rendering for reflectance decomposition with near-field indirect illumination. In *CVPR*, 2023. 4, 7
- [73] Shangzhe Wu, Ameet Makadia, Jiajun Wu, Noah Snavely, Richard Tucker, and Angjoo Kanazawa. De-rendering the world’s revolutionary artefacts. In *CVPR*, 2021. 3, 4
- [74] Tong Wu, Jiarui Zhang, Xiao Fu, Yuxin Wang, Liang Pan, Jiawei Ren, Wayne Wu, Lei Yang, Jiaqi Wang, Chen Qian, Dahua Lin, and Ziwei Liu. OmniObject3D: Large-vocabulary 3D object dataset for realistic perception, reconstruction and generation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 2, 3
- [75] Kohei Yamashita, Yuto Enyo, Shohei Nobuhara, and Ko Nishino. nLMVS-Net: Deep non-lambertian multi-view stereo. In *WACV*, 2023. 3
- [76] Kai Yan and Shuang Zhao. PSDR-JIT: Path-space differentiable renderer, 2020. <https://github.com/andyankai/psdr-jit>. 5
- [77] Lior Yariv, Yoni Kasten, Dror Moran, Meirav Galun, Matan Atzmon, Basri Ronen, and Yaron Lipman. Multiview neural surface reconstruction by disentangling geometry and appearance. In *NeurIPS*, 2020. 7
- [78] Lior Yariv, Jiatao Gu, Yoni Kasten, and Yaron Lipman. Volume rendering of neural implicit surfaces. *Advances in Neural Information Processing Systems*, 34:4805–4815, 2021. 11
- [79] Vickie Ye, Ruilong Li, Justin Kerr, Matias Turkulainen, Brent Yi, Zhuoyang Pan, Otto Seiskari, Jianbo Ye, Jeffrey Hu, Matthew Tancik, and Angjoo Kanazawa. gsplat: An open-source library for Gaussian splatting. *arXiv preprint arXiv:2409.06765*, 2024. 7
- [80] Hong-Xing Yu, Samir Agarwala, Charles Herrmann, Richard Szeliski, Noah Snavely, Jiajun Wu, and Deqing Sun. Accidental light probes. In *CVPR*, 2023. 3, 5
- [81] Hong-Xing Yu, Michelle Guo, Alireza Fathi, Yen-Yu Chang, Eric Ryan Chan, Ruohan Gao, Thomas Funkhouser, and Jiajun Wu. Learning object-centric neural scattering functions for free-viewpoint relighting and scene composition. *TMLR*, 2023. 3
- [82] Kai Zhang, Fujun Luan, Qianqian Wang, Kavita Bala, and Noah Snavely. PhySG: Inverse rendering with spherical gaussians for physics-based material editing and relighting. In *CVPR*, 2021. 4, 7, 9
- [83] Kai Zhang, Fujun Luan, Zhengqi Li, and Noah Snavely. IRON: Inverse rendering by optimizing neural SDFs and materials from photometric images. In *CVPR*, 2022. 4
- [84] Kai Zhang, Sai Bi, Hao Tan, Yuanbo Xiangli, Nanxuan Zhao, Kalyan Sunkavalli, and Zexiang Xu. GS-LRM: Large reconstruction model for 3D gaussian splatting. *arXiv preprint arXiv:2404.19702*, 2024. 2
- [85] Ruo Zhang, Ping-Sing Tsai, James Edwin Cryer, and Mubarak Shah. Shape-from-shading: a survey. *IEEE TPAMI*, 21(8):690–706, 1999. 3
- [86] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018. 9
- [87] Xiuming Zhang, Pratul P Srinivasan, Boyang Deng, Paul Debevec, William T Freeman, and Jonathan T Barron. NeRFactor: Neural factorization of shape and reflectance under an unknown illumination. *ACM TOG*, 40(6):1–18, 2021. 3

- [88] Yuanqing Zhang, Jiaming Sun, Xingyi He, Huan Fu, Rongfei Jia, and Xiaowei Zhou. Modeling indirect illumination for inverse rendering. In *CVPR*, 2022. [4](#), [7](#)
- [89] Yunzhi Zhang, Shangzhe Wu, Noah Snavely, and Jiajun Wu. Seeing a rose in five thousand ways. In *CVPR*, 2023. [4](#)

Digital Twin Catalog: A Large-Scale Photorealistic 3D Object Digital Twin Dataset

Supplementary Material

We provide the following content in the supplementary materials:

- We provide a comprehensive dataset statistics of DTC dataset in Sec. A, including the full category of DTC objects, the list of objects used in DSLR and egocentric recording respectively. We also provided visualizations for examples of DTC objects, DSLR and egocentric data.
- We include complementary details of benchmark in Sec. B. First, as indicated by the main paper, we include a sparse view setting of benchmark. We further provide the baseline comparisons and analysis for DSLR benchmark and egocentric benchmark. We use the same baselines and dataset split as described in the main paper.
- We provide details of the simulation and experiments for our robotics experiments in Sec. C.

A. DTC Dataset Statistics & Details

In this section, we present the category list of DTC 3D objects (Sec. A.1) and showcase additional examples of DTC digital-twin-quality 3D objects (Fig. 9). We also provide the list of selected objects used for creating the DSLR evaluation dataset (Sec. A.2) and the egocentric evaluation dataset (Sec. A.3). Furthermore, we include additional examples of real-world captures and recordings to better illustrate the DSLR (Sec. A.4) and Egocentric (Sec. A.5) evaluation datasets.

A.1. Categories of DTC Scanned Objects

In DTC dataset, we selected and scanned 2,000 physical-world objects across 40 carefully curated categories from the taxonomy of LVIS [27]. These categories were chosen to ensure a diverse representation of common daily objects while remaining compatible with the scanner’s capabilities. Table 5 provides an overview of these categories, including the number of models, average vertex count, minimum/maximum vertex numbers for each category and corresponding category label in LVIS taxonomy.

For each scanned model, we generated high-quality 4K-resolution PBR material maps—including albedo, roughness, metallic, and normal maps—to achieve a photorealistic appearance. Fig. 9 showcases additional examples of DTC scanned digital-twin-quality 3D models, each accompanied by a full set of PBR material maps (Fig. 10, 11, 12, 13, 14).

A.2. List of Models for DSLR Evaluation Data

The full list of object models selected for DSLR evaluation data is as follows, and each model is captured under 2 dif-

Category	# Models	avg # vert / std	min / max	LVIS Category
airplane	15	129,153 / 15,535	80,719 / 144,587	airplane
axe	3	133,588 / 1,922	132,096 / 136,301	ax
basketball	55	132,260 / 3,013	128,652 / 140,985	basketball
birdhouse	102	147,577 / 11,057	98,405 / 182,594	birdhouse
bowl	106	127,456 / 1,019	125,883 / 130,964	bowl
building blocks	129	96,592 / 46,377	28,181 / 301,475	toy
calculator	38	131,244 / 3,638	126,379 / 139,970	calculator
candle	31	108,221 / 58,406	15,897 / 309,601	candle
candle holder	1	132,285 / 0	132,285 / 132,285	candle_holder
cast iron	31	165,954 / 127,424	128,074 / 711,536	pan
cup	58	137,154 / 43,605	126,346 / 417,343	cup
cutting board	16	133,796 / 11,244	127,543 / 173,141	chopping_board
dino	103	82,403 / 29,477	13,796 / 151,148	animal
dish	51	129,732 / 4,130	127,014 / 147,451	dish
dumbbell	39	156,915 / 165,650	126,593 / 1,177,907	dumbbell
fake food can	79	117,829 / 22,772	78,646 / 270,271	can
fakefruit	96	124,037 / 9,724	100,833 / 151,870	fruit
figurine	77	57,152 / 47,753	15,482 / 140,812	figurine
football	48	132,323 / 2,953	126,719 / 138,459	football
gargoyle	50	137,921 / 5,019	130,206 / 151,325	gargoyle
gravestone	24	85,245 / 55,760	10,750 / 150,966	gravestone
hammer	33	133,413 / 36,272	46,808 / 320,845	hammer
hardcover_book	17	174,884 / 102,168	130,267 / 500,735	hardback_book
key	2	60,456 / 41,100	19,357 / 101,556	key
keyboard	25	146,684 / 12,414	129,360 / 174,488	computer_keyboard
knife	10	126,359 / 8,118	102,569 / 133,250	knife
mallard (fake duck)	48	99,728 / 45,926	9,000 / 131,058	mallard
marker	54	93,394 / 39,460	34,796 / 306,272	marker
miscellaneous	44	149,799 / 168,694	27,772 / 1,208,696	NA
mouse	52	144,184 / 46,099	121,879 / 304,533	mouse
pistol	2	133,587 / 91	133,496 / 133,678	pistol
pottery	46	143,284 / 162,783	12,354 / 1,206,223	pottery
remote	2	42,993 / 17,068	25,925 / 60,061	remote_control
shampoo	45	143,639 / 75,261	126,285 / 604,491	shampoo
shaver	20	142,231 / 57,736	38,246 / 306,105	shaver
shoes	121	139,363 / 4,171	130,604 / 148,465	shoe
speaker	40	201,013 / 152,410	127,154 / 852,714	speaker
spoon	34	107,649 / 56,459	29,712 / 302,920	spoon
teapot	99	145,149 / 100,792	85,289 / 926,273	teapot
vase	101	142,927 / 142,946	61,488 / 1,568,710	vase
volleyball	52	133,921 / 4,056	129,653 / 144,770	volleyball

Table 5. Categories of DTC Scanned Objects. We include the number of models per object (#Models), the average number of vertices per object categories (avg # vert) and its standard deviation (std), the minimum and maximum number of vertices within the object category (min/max), and the label name in LVIS category taxonomy.

ferent environment lighting conditions. For the 15 object recordings that are used in our benchmark evaluations, we highlight them in **bold**.

- **Airplane_B097C7SHJH.WhiteBlue**
- **Airplane_B0B2DC5QBP.BlueGray**
- **BirdHouse**
- **BirdHouse_B0B8F27TFK.BrownRoofYellowWalls**
- **BirdHouseRedRoofYellowWindows**
- **BirdHouseWoodenRoofGreenWall**
- **Bowl_B0BQR77WRW.LightGrey-1.TU**
- **Box_ADITR.DecorativeBoxHexLarge.Green**
- **Car_38330969.Toy**
- **CaramicBowlBluewithBrown**
- **CeramicBowlBigWhite**
- **Cup_B08TWHJ33Q.Tan**
- **Cup_B0B3JKZW76.Brown**




Figure 9. Examples of DTC 3D models. The PBR Materials of each object are presented in the following figures from Fig. 10 to Fig. 14.

- Cup_B0CQXPND8L.Stripes
- Dutch.Oven_B0B916N11D.Black
- **Figurine_B08FYFNYP4.LionKing**
- Figurine_B0983CQ2HH.Angel
- Figurine_B0CR3Y5T3K.Gnome
- Gargoyle_B005KDPAFW.BatWings
- **Gargoyle_B08SQMBDXY.HandsOnKnees**
- **Gargoyle_B0C2PNF2C1.Meditating**
- Gravestone_B08TBJQ5XP.LightGrayKitty
- Hammer_B000FK3VZ6.Wood
- Home_ADTIR_A6116F6.DraganBoxSmall.Wood
- Home_ADTIR_L041OV8.Rinnig.PlateHolder
- Keyboard_B07P6K5GMY.Black
- Keyboard_B0CL8S2DW9.Pink
- Kitchen.Spoon_B008H2JLP8.LargeWooden
- **Mallard_B082D168CK.MintGreen**



Figure 10. PBR Materials of the example DTC objects (the list of objects in Fig.9 Row 1), From left to right: albedo map, roughness map, metallic map, normal map, and PBR rendering.

- **Mallard_B09LV16HD5_LightBrown2**
- **Mallard_B0BPY18VHR.White**
- **Mallard_B0C6MQWM21.BlackWhite**
- **Mouse_B0CHNVBBLF_Honeycomb_1**
- **Pan_B0CFQWYJZ8.BlackWoodHandle**
- **Pan_B0CHW1KK8Z.Black**
- **Planter_B0C4G81ZPF_Cat**
- **Pottery_B097S319TR.Woman**
- **Pottery_B0CJJ59SLH.BlueHairFairy**
- **Shoe_B000ZP6MIY_Navy7L_TU**
- **Spoon_B08M3XNKYR.Slotted**
- **TeaPot_B074ZQYRP7.BrownDragonShaped**
- **TeaPot_B07GL8MH3X.PinkFlamingo**
- **TeaPot_B07QP5MFQ1.BlackCastIron**
- **TeaPot_B084G3K8TD.YellowBlackSunflowers**
- **TeaPot_B08HSDHBM4.BlackGoldLeaves**
- **TeaPot_B00ESU7PFG.WhiteRoseFlowers**
- **TeaPot_B01KFCZB2Y.WhiteWoodHandle**



Figure 11. PBR Materials of the example DTC objects (the list of objects in Fig.9 Row 2). From left to right: albedo map, roughness map, metallic map, normal map, and PBR rendering.

- Vase_B09ZGXSVTT.White_TU
- Vase_B0BV44B4R4.BlueBirdsYellowBirds
- Vase_Corrected

A.3. List of Models for Egocentric Evaluation Data

The full list of models selected that contains pairs of egocentric data is as follows, and each model is captured with both *active* and *passive* trajectories. For the 15 object recordings that are used in our benchmark evaluations, we highlight

them in **bold**.

- Airplane_B097C7SHJH.WhiteBlue
- Airplane_B09WN2RN15.Black-1
- BasketPlasticRectangular
- BirdHouseToy
- **BirdHouse_B0B8F27TFK.BrownRoofYellowWalls**
- BirdHouse_B004HJE8AS.WhiteWallsTwoPorches_2
- BirdHouseRedRoofYellowWindows
- BirdHouseWoodenRoofGreenWall



Figure 12. PBR Materials of the example DTC objects (the list of objects in Fig.9 Row 3). From left to right: albedo map, roughness map, metallic map, normal map, and PBR rendering.

- BirdHouse_B08GYBKJ8N_RedBarn
- Birdhouse_B09FJYJYDQ_BoatHouse
- BirdHouseMetalRoofYellowWall_1.TU
- BirdHouse_A79823645_BearWithLogStump
- BirdHouseWoodenRoofRedWall
- BlackCeramicDishLarge
- BlackCeramicMug
- Block_B007GE75HY_RedBlue
- Bowl_B0BQR77WRW_LightGrey_1.TU
- Bowl_B07ZNJ5RQV_Orange_TU
- CandleDishSmall
- CrateBarrelBowlRed
- **CeramicBowlBigWhite**
- CelebrateBowlPink
- Car_38330969_Toy
- CeramicBowlBluewithBrown
- Candle_B0B2JQWNNQ.White
- Candle_B0B764F39X_Turquoise



Figure 13. PBR Materials of the example DTC objects (the list of objects in Fig.9 Row 4). From left to right: albedo map, roughness map, metallic map, normal map, and PBR rendering.

- Candle_B09MP8NDML.White
- Calculator_B0C7GP2D5C.Purple
- **Cup_B08TWHJ33Q.Tan**
- Cup_B0B3JKZW76.Brown
- Cup_B0CQXPND8L.Stripes
- Cup_B08TWHJ33Q.Gray
- Cup_B09QCYR1SL.Pitcher_1
- Cup_B01LYONYPB.SkyBlue
- Dumbbell_B00PY62Y90.Black
- Dumbbell_B0CN56CQPS.PurpleCoolGray
- Dutch.Oven_B0B916N11D.Black
- FakeFruit_B076H96CS1.Banana
- FakeFruit_B09992T572.WatermelonSlice
- FakeFruit_B0815W7RKC.StarFruit
- Figurine_B08FYFNYP4.LionKing
- **Figurine_B0983CQ2HH.Angel**
- Figurine_B0CR3Y5T3K.Gnome
- **Flask**




Figure 14. PBR Materials of the example DTC objects ((the list of objects in Fig.9 Row 5)), From left to right: albedo map, roughness map, metallic map, normal map, and PBR rendering.

- Gargoyle_B07GHVQ3C4_BatCat
- **Gargoyle_B08SQMBDXH_HandsOnKnees**
- Gargoyle_B0C3RQ5254_Bronze
- Gargoyle_B0BYTQT173_Dragon
- Gargoyle_B0C2PNF2C1_Meditating
- Gargoyle_B00N08IU24_Dog
- Gravestone_B08TBJQ5XP_LightGrayKitty
- Gravestone_B07WDGH3NR_GrayAngel
- Hammer_B000FK3VZ6_Wood
- Hammer_B0BN6FXDQ7_BlueHandle
- Hammer_B01N63ONKY_DrillingSledge
- Kitchen_Cup_B09G2WNN61_DarkBlue
- Kitchen_Spoon_B008H2JLP8_LargeWooden
- Kitchen_Spoon_D146567C_Green_1
- LargeLightGreyDish
- Mallard_B082D168CK_MintGreen
- **Mallard_B09LV16HD5_LightBrown2**
- **Mallard_B0BPY18VHR.White**

- Mallard_B0C6MQWM21.BlackWhite
- Mallard_B00RTSJU7K.Red
- Mouse_B0CHNVBBLF.Honeycomb_1
- Pan_B0CFQWYJZ8.BlackWoodHandle
- Pan_B0CHW1KK8Z.Black
- Planter_B0C4G81ZPF.Cat
- PlasticBowlGreen_1.TU
- **Pottery_B097S319TR.Woman**
- Pottery_B0CJJ59SLH.BlueHairFairy
- Pottery_B07TGC6TGL.White
- Pottery_B075SX9GVK.White
- Pottery_B0CF8FW987.Man
- **Shoe_B000ZP6MIY.Navy7L.TU**
- Shoe_B094ZCQK75.Red6HL.TU
- **Spoon_B08M3XNKYR.Slotted**
- TeaPot_B074ZQYRP7.BrownDragonShaped
- TeaPot_B07GL8MH3X.PinkFlamingo
- TeaPot_B07QP5MFQ1.BlackCastIron
- TeaPot_B084G3K8TD.YellowBlackSunflowers.TU
- **TeaPot_B08HSDHBM4.BlackGoldLeaves**
- **TeaPot_B00ESU7PFG.WhiteRoseFlowers**
- **TeaPot_B01KFCZB2Y.WhiteWoodHandle**
- TeaPot_B07RT7BYXL.WhiteSnowOwl.TU
- **Vase_B09ZGXSVTT.White.TU**
- Vase_B0BV44B4R4.BlueBirdsYellowBirds
- Vase_Corrected
- Vase_B0BNX2CWW8.YellowTall
- Vase_B00858OOXI.Green_1
- Vase_B09FYBCM1R.BeeSunflowers
- Vase_B0BY8PYLLC.PinkPineapple
- Vase_B0BNX2CWW8.YellowShort
- WhiteContainerBox
- WhiteClip_1
- WoodBlocks_B00FIX22YQ.BlueArch
- WoodBlocks_B098PHYN3P.SmallOctagon
- WoodenBoxSmall
- WoodenFork

A.4. An Example of DSLR Data

As described in Sec. 3.3, we used three DSLR cameras mounted on a rig to capture 360-degree inward-facing photos of each object. As shown in Fig. 15, these DSLR cameras were positioned at different elevation angles and captured images circularly around the object with a 9-degree interval between shots. Each camera captured 40 photos, resulting in a total of 120 photos per object. For each photo, we provide the corresponding camera pose and an accurate foreground object mask (2nd row overlaid on the original image in Fig. 15) to facilitate straightforward and reliable evaluation. Additionally, to support the evaluation of reliable reconstruction, we utilized a mirror ball to capture and reconstruct an HDR environment map (bottom row in Fig. 15).




Figure 15. Example DSLR evaluation data: captured images with 3 DSLR cameras from different elevation angles (first row), the object mask overlaid on each image (second row) and the recovered HDR envmap (last row).

A.5. An Example of Egocentric Data

Fig. 16 provides an example visualization of the egocentric recording observed from its RGB camera and the rendered ground truth using the 3D aligned digital twin model using the method described in Sec. 3.4. The raw RGB input from Project Aria is a wide angle fisheye camera, which limits the amount of existing off-the-shelf baselines using pin-hole camera models. To resolve this limitation, we rectify all the raw images using a linear camera model, and render the corresponding ground truth of the linear camera model as well. All the evaluations on egocentric recordings are performed using the linear rectified images as input and output. We use the focal length 1200 and 2400x2400 resolution which retains the original pixel resolution and field of view from the raw image input.

It is worth noting the image input from Project Aria is gravity aligned. Both the raw image and rectified image will appear as 90 degree counter-clock wise rotated according to the gravity direction without explicitly updating the camera calibration information. For the reconstruction and evaluation, we do not rotate the image or update the calibration. We only rotate the rendered images 90 degree clockwise for the benefit of visualization purpose (e.g. Fig. 22)

We provided two different type of recordings to benchmark reconstruction from egocentric views. Fig. 17 shows a



Figure 16. An example of visualization for the egocentric recordings. From left to right, we present the original egocentric view from the fisheye RGB camera input, the rectified image view, object mask, rendered depth, and rendered normal. The egocentric recording are gravity aligned respect to the global trajectory coordinate. We use the rectified image for all evaluations.




Figure 17. An snapshot visualization of the active (above) and passive (below) of egocentric recordings in the 3d space. The left of image shows the trajectory (visualized in Violet color) of the camera movement in each scenario. The yellow points are the scene point cloud.

snapshot of the egocentric recording. We demonstrate each type of camera trajectory overlaid in the 3D space using the scene point cloud. The active trajectory represents a common object-centric 360 view of the objects while the passive trajectory represents a more casual walk around the object from different distances.

B. Complementary Benchmarks

In this section, we provide complementary results on the different applications of benchmarking in Sec. 4. First, we include the qualitative comparisons of baselines in the inverse rendering application using the DTC dataset. We use the same baselines described in Sec. 4.1. Second, as indicated in the main paper, we include a new baseline evaluation

for sparse-view reconstruction in Sec. B.2 for DSLR data and Sec. B.3 for egocentric data. Third, we present the qualitative comparisons of egocentric active reconstruction baselines described in Sec. 4.2.

Explanations on Evaluation Metrics. In Sec. 4.1, we evaluate the quality of *geometry reconstruction* by comparing the rendered depth maps, normal maps, and predicted 3D meshes to the ground truth. For depth maps, we use the Scale-Invariant Mean Squared Error (**SI-MSE**) as the evaluation metric. Normal maps are evaluated using **Cosine Distance**, and for the 3D meshes, we compute the **Bi-directional Chamfer Distance** between sampled surface points. For *novel view synthesis* and *novel scene relighting*, we utilize widely adopted metrics: Peak Signal-to-Noise Ratio for both HDR (**PSNR-H**) and LDR (**PSNR-L**), Structural Similarity Index Measure (**SSIM**) [67], and Learned Perceptual Image Patch Similarity (**LPIPS**) [86]. We follow PhySG [82] and adapt all these metrics to be scale-invariant, addressing the ambiguity in material and lighting decomposition. We use the same metrics for novel-view synthesis in the egocentric novel-view synthesis applications as well. The same evaluation metrics are used in StanfordORB [39].

B.1. Inverse Rendering Baseline Comparisons

We provide qualitative comparisons as a supplement to the inverse rendering baseline comparisons in Sec. 4.1. Our DTC dataset contains objects with varying geometric complexity, ranging from simpler ones, *e.g.*, cups (Fig. 18 Right), to complex ones, *e.g.*, toy birdhouses (Fig. 18 Left). The surface material ranges from diffuse ones, *e.g.*, toy houses, to highly reflective ones, *e.g.*, cups. For the two examples in Fig. 18, we observed that methods using surface-based representations, *i.e.*, PhySG and InvRender, and hybrid representations, *i.e.*, NVDiffRec and NVDiffRecMC, tend to produce smoother geometry compared to methods using NeRF [52]-based representations, *i.e.*, Neural-PIL and NeRD.




Figure 18. **Qualitative Comparisons of Baseline Methods.** We show qualitative comparisons of baseline methods, used in Table 2, on two different DTC models BirdHouseToy (Left) and Cup_B08TWHJ33Q_Tan (Right).

Table 6. Quantitative evaluation of trained LRM-VolSDF baseline for view synthesis on GSO dataset with different number of images. It achieves the state-of-the-art performance compared to prior work.

LRM-VolSDF	PSNR (\uparrow)	SSIM (\uparrow)	LPIPS (\downarrow)
MeshLRM [70]	28.13	0.923	0.093
4 images	28.72	0.940	0.070
8 images	30.19	0.947	0.061

Table 7. Quantitative results on novel view synthesis for sparse-view reconstruction using DTC DSLR data.

	PSNR-H (\uparrow)	PSNR-L (\uparrow)	SSIM (\uparrow)	LPIPS (\downarrow)
LRM-VolSDF	20.53	21.16	0.993	0.006




Figure 19. LRM view synthesis results from DSLR data.

B.2. Sparse-view Reconstruction Application for DSLR data

We provide a sparse-view reconstruction baseline on our DSLR data using learning based reconstruction method. To build the baseline, we train an LRM model similar [70] that achieves state-of-the-art mesh reconstruction results as demonstrated on public synthetic dataset on GSO[21] dataset

in Table 6.

LRM baseline Our LRM baseline has the same transformer architecture and tri-plane representation as [70]. It consists of 24 self-attention blocks with feature dimension 1024 and 16 heads of attention. The tri-plane token number is 32×32 . Each token is decoded to a 8×8 feature patch through a



Figure 20. Example input data of egocentric sparse-view reconstruction. We used the masked pixels from the rectified egocentric passive data as input. Some objects can appear very small due to the challenges in free-view movement in egocentric passive trajectory.

Table 8. Quantitative results on novel view synthesis for sparse-view reconstruction using egocentric *passive* data.

	PSNR-H (\uparrow)	PSNR-L (\uparrow)	SSIM (\uparrow)	LPIPS (\downarrow)
LRM-VolSDF	20.03	21.72	0.971	0.065

linear layer, which leads us to final tri-plane resolution of 256×256 . During training, our LRM directly outputs SDF value instead of density and we use the volume ray tracing method proposed in [78] to render images for supervision. Compared to the NeRF representation used in [70], our LRM baseline is robust to train and can achieve accurate geometry reconstruction. We evaluate our LRM-VolSDF baseline on synthetic GSO dataset. Results are summarized in Table 6. With 8 input views, our baseline achieves reconstruction accuracy comparable to the state-of-the-arts.

Experiments and evaluation We randomly sample 16 images from the 120 training views as inputs to our LRM-VolSDF network. The testing view is the same as the dense-view inverse rendering experiments. The quantitative numbers for view synthesis are summarized in Table 7. We observe that all three metrics are much worse compared to the results from synthetic GSO dataset, indicating the existence of a domain gap between synthetic and real data. Fig. 19 visualize some of our reconstruction results. We observe that for relative simple shape, our model generalizes well and achieves highly detailed reconstruction that closely match the ground-truths. However, it struggles at reconstructing more complicated object, such as the Birdhouse. Further investigation is required to understand the gap on real world dataset.

B.3. Sparse-view Reconstruction Application for Egocentric data

We use the egocentric *passive* recordings in this evaluation, which is a challenging sparse-view reconstruction setting where we only have casual observation of an object captured by an egocentric RGB camera. Compared to the sparse-view reconstruction for DSLR data where the object always



Figure 21. LRM view synthesis results from egocentric data.

appear in the image center and occupies a large portion of the image, the egocentric recording has natural human moving trajectory and object can appear small in many views.

Fig. 20 shows several examples of the masked objects that we use as input to this evaluation. We can see the camera poses are much more diverse and object only occupies a very small portion of the image. This setting will defy any classical optimization-based 3D reconstruction method. We build a baseline using the same LRM-VolSDF model as mentioned in Sec. B.2. We tested the model on the challenging data without any fine-tuning.

Experiments and Evaluation. For each sequence, we randomly sample 24 views from the middle 1/3 of the whole sequence as this is when our Aria glasses are relatively close to the object. We use 16 of the 24 views as inputs and 8 views as ground-truth for testing. Similar to the DSLR experiment, we use centralized cropping to place the object in the center of the image and modify the plucker rays representation accordingly. Qualitative and quantitative results are summarized in Fig. 21 and Tab. 8 respectively. We observe that despite that the camera poses distribution of input views is very different from our training data, the LRM-VolSDF generalizes to the new type of input data. Our PSNR numbers for the masked region are much lower than those of synthetic data but still achieve reasonable performance. We calculate the LPIPS and SSIM based on the full image which results in higher numbers. This evaluation expose the domain gap in learning based reconstruction map for egocentric data input. We hope our current experimental results can be a promising starting point.

B.4. Dense-view Reconstruction for Egocentric data

As a complement to Sec. 4.2, we provided qualitative comparisons of novel view synthesis in Fig. 22. We rendered the depth from both baselines in a normalized depth range. For 3D-GS, we acquire the normal from its point cloud ray casted from each pixels. For 2D-GS, we directly use the predicted normal from the rasterizer.

Results analysis: Both 3D-GS[37] and 2D-GS [31] can

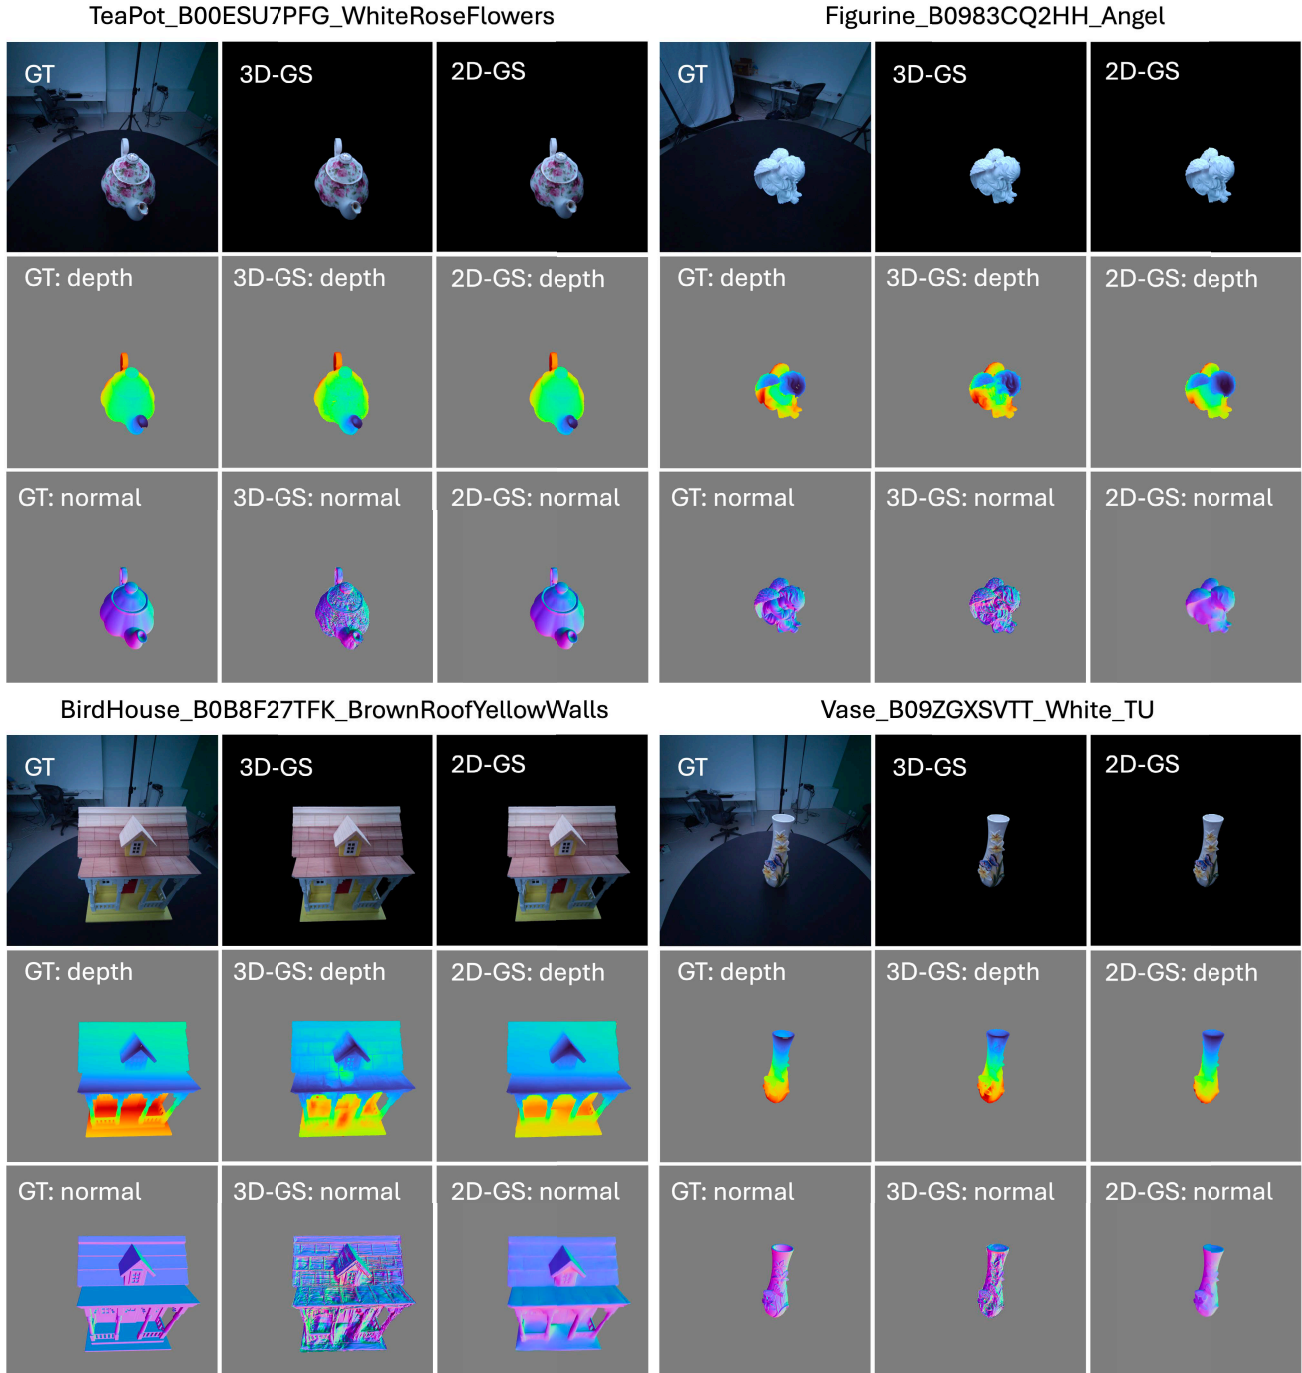


Figure 22. Qualitative comparisons of baselines reconstructions on egocentric recordings. We compare 3D-GS[37] and 2D-GS[31] to the ground truth using the modalities of rendered images, depth and normal. Both baselines can provide near photoreal view synthesis on the held-out validation view compared to ground truth. However, the comparison in their geometry (depth and normal) indicate the existing methods still fall short to recover the details in digital twin reconstruction. For the benefit for visualization, we rotate all the images 90 degree clockwise.

provide near photo-real view synthesis of the objects with high quality view-dependent effect. 3D-GS perform slightly better in PSNR metric consistent across all objects. In terms

of shape reconstruction, 2D-GS performs better in particular to recover surface normals. Compared to the ground truth rendering, the estimated depth from 3D-GS is significantly

noiser which also leads to visible artifacts in its normal map, while the 2D-GS tends to predict smoother depth and normal and can ignore certain details. For objects with simple shape (e.g., `Teapot_B00ESU7PFG.WhiteRoseFlowers`), both methods perform well. However, for objects with complex shapes, both methods fail to recover the complex geometric details in the object shape (e.g., `BirdHouse_B0B8F27TFK.BrownRoofYellowWalls`). The challenge in shape reconstruction indicates the direction for future research work in this area.

C. Experimental Details For Robotic Experiments

C.1. Experimental Objects

The specific objects that we use from the DTC and Objaverse-XL datasets ² are as follows:

DTC dataset:

- `Cup_B01LYONYPB.SkyBlue`
- `Cup_B0BR43SPKJ.Blue`
- `Cup_B0CJBZT7N5.Black`
- `Cup_B0CMPB8FNY.MountainBluebell`
- `Cup_B0CYL5PSR3.Gray`
- `Cup_B08PTSRWF8.Green`
- `Cup_B0CMD4LX4D.DarkBlue`
- `Cup_B08TWHJ33Q.Tan`
- `Cup_B0CNJP2KZF.GreenOrange`
- `Cup_B094NQH2YM.BlackGold`
- `Cup_B0BXB21T7Q.Blue`
- `Cup_B0CPX832ZP.Floral`
- `Cup.White`
- `Cup_B09L8DS2ZB.DarkBrown`
- `Cup_B0C3X3WY2Q.SageGreen`
- `Cup_B0CQTF6GF1.RedWhite`
- `Cup_B09QCYR1SL.Pitcher_1`
- `Cup_B0C81BCSXQ.WhiteRainbow`
- `Cup_B0CQXPND8L.Stripes`
- `Cup_B09YDPVRM7.RedTeaCup`
- `Cup_B0CDWXPDK1.Zebra`
- `Cup_B0CR45H24G.Blue`
- `Cup_B0B3JKZW76.Brown`
- `Cup_B0CHJYN6F1.Black`

For Objaverse-XL, we collect a subset of 24 objects by filtering objects by the STL file format and by the word `cup` appearing in the provided metadata tags, and then manually filtering the results by inspecting the object files to only include objects that can be considered “cups” by a human judge.

²The version of Objaverse-XL used in this work excludes all 3D models sourced from Sketchfab. Further, no Polycam assets were obtained from the Polycam source site

C.2. Simulation and Data Collection

When importing objects into the PyBullet simulator, we obtain separate collision meshes by performing convex decomposition of each object with CoACD [69] with a concavity threshold of 0.01, and rescale each object to have a maximum bounding box side length of 0.137 m (to match that of the test object). Since not all Objaverse-XL objects come with textures, we randomly color Objaverse-XL objects uniformly in RGB space.

For data collection, we uniformly randomly select one object from the considered object set to be placed into the scene. We uniformly randomly place the object within a 0.4×0.4 m box centered at the scene origin at a z height of 0.10m, and allow it to fall to the floor workspace. If the object rolls away (further than a 0.8×0.8 m bound), we re-sample the initial position and re-attempt the object placement.

Pushing: For pushing, we generate pseudo-random trajectories by executing a scripted policy in two stages: First, the robot samples a starting position from the push by randomly sampling an angle θ and radius r . It then takes steps to move its end-effector to the x-y location $(x_{obj} + r \sin(\theta), y_{obj} + r \cos(\theta))$, where (x_{obj}, y_{obj}) is the initial object position. After reaching within 2cm of this location, the robot begins to push the object, by moving from its current end-effector position toward the object location. All steps are normalized to have a magnitude of 0.03m, and the action at each step is affected by uniform random noise drawn from $\mathcal{U}(-0.05, 0.05)$ in each axis. The trajectory ends after 35 steps.

Grasping: For grasping, we generate successful grasps by sampling grasping candidates and filtering only the successful ones. After adding an object to the scene, we uniformly randomly sample a grasping position (x, y, z) where $x = x_{obj} + \mathcal{U}(-0.05, 0.05)$, $y = y_{obj} + \mathcal{U}(-0.05, 0.05)$, $z = 0.2 + \mathcal{U}(-0.1, 0.1)$ where (x_{obj}, y_{obj}) is the initial object position. The robot then performs four steps: (1) It moves in the air to the target grasp position x, y coordinates, (2) it moves vertically downwards to a height of z , (3) it closes the gripper, and (4) it lifts to a height of 0.5m. The grasp is considered successful if the object is at least 0.1m above the floor after step (4).

C.3. Policy Training

We train convolutional neural network policies to regress actions representing xy position changes and the (x, y, z, θ) grasp position for pushing and grasping respectively. For pushing, we train goal-image-conditioned policies, relabeling goals using a hindsight sampling mechanism. We sample goals uniformly randomly from future timesteps in the trajectory of a sampled initial state, and use the first action after the initial state as the label.

We use similar architectures for pushing and grasping.

The architecture consists of an encoder and an MLP head for action prediction. Below we describe the architecture for goal-conditioned pushing, with differences for the grasping policy noted. The input image size for both initial and goal images is $256 \times 256 \times 3$.

For pushing, we normalize action labels by multiplying each dimension by 100, and for grasping, we normalize the last dimension (grasp angle) by dividing by 10 such that it has similar magnitudes to the other action label dimensions.

Encoder:

- Conv2d(3, 32, 3, stride=2, padding=1)
- ReLU
- Conv2d(32, 64, 3, stride=2, padding=1)
- ReLU
- Conv2d(64, 64, 3, stride=2, padding=1)
- ReLU
- Flatten
- Linear($(\text{image_size}/8)^2 \times 64$, 512)
- ReLU

Note that for the pushing policy, the observation and goal image share encoder weights.

Action Prediction Head:

- Linear(512 \times 2, 256) (the input feature dimension is 512 for grasping, which has a single image input as opposed to 1024 and two image inputs for pushing, which have their features concatenated together).
- ReLU
- Linear(256, 256)
- ReLU
- Linear(256, 2) (for grasping, the final output size is 4)

We train policies using a mean-squared-error loss using the Adam optimizer with a learning rate of $3e - 4$ and a batch size of 32, holding out 5% of the data for validation. We train for 200 epochs and take the checkpoint with lowest validation loss for evaluation. Training uses a single NVIDIA Titan RTX GPU.

C.4. Evaluation

For evaluation, we generate test sets of 100 initial conditions for each task using the “cup_scene003” object from the StanfordORB dataset. For pushing, we generate trajectories in the same way as during data collection and use the initial state and final state as the initial condition and target goal for the policy. We roll out the policy for 35 steps and compute the final position error of the object in the x, y dimensions. For grasping, we generate initial object positions in the same way as in data collection and the policy performs a single grasp attempt, which is determined successful if the object ends at least 0.1m above the working surface.

In Figure 8, we present detailed results for the pushing experiments, plotting the success rates for each policy with respect to varying success thresholds based on final object

position error.