# Common Objects in 3D: Large-Scale Learning and Evaluation of Real-life 3D Category Reconstruction

Jeremy Reizenstein[1] Roman Shapovalov[1] Philipp Henzler[2] Luca Sbordone[1]
Patrick Labatut[1] David Novotny[1]

{reizenstein,romansh,lsbordone,plabatut,dnovotny}@fb.com {p.henzler}@cs.ucl.ac.uk

[1]Facebook AI Research [2]University College London

https://github.com/facebookresearch/co3d


Figure 1: We introduce the **Common Objects in 3D (CO3D)** dataset comprising 1.5 million multi-view images of almost 19k objects from 50 MS-COCO categories annotated with accurate cameras and 3D point clouds (visualized above).

## Abstract

*Traditional approaches for learning 3D object categories have been predominantly trained and evaluated on synthetic datasets due to the unavailability of real 3D-annotated category-centric data. Our main goal is to facilitate advances in this field by collecting real-world data in a magnitude similar to the existing synthetic counterparts. The principal contribution of this work is thus a large-scale dataset, called* Common Objects in 3D*, with real multi-view images of object categories annotated with camera poses and ground truth 3D point clouds. The dataset contains a total of 1.5 million frames from nearly 19,000 videos capturing objects from 50 MS-COCO categories and, as such, it is significantly larger than alternatives both in terms of the number of categories and objects.*

*We exploit this new dataset to conduct one of the first large-scale "in-the-wild" evaluations of several new-view-synthesis and category-centric 3D reconstruction methods. Finally, we contribute* NerFormer *- a novel neural rendering method that leverages the powerful Transformer to reconstruct an object given a small number of its views.*

## 1. Introduction

Recently, the community witnessed numerous advances in deeply learning to reconstruct category-centric 3D models. While a large variety of technical approaches was proposed [56, 62, 13, 25, 26, 46, 8], they are predominantly trained and benchmarked either on synthetic data [6], or on real datasets of specific object categories such as birds [62] or chairs [37]. The latter is primarily a consequence of a lack of relevant real-world datasets with 3D ground truth.

Our main goal is therefore to collect a large-scale open real-life dataset of common objects in the wild annotated with 3D ground truth. While the latter can be collected with specialized hardware (turn-table 3D scanner, dome [29]), it is challenging to reach the scale of synthetic datasets [6] comprising thousands of instances of diverse categories.

Instead, we devise a photogrammetric approach only requiring object-centric multi-view RGB images. Such data can be effectively gathered in huge quantities by means of crowd-sourcing "turn-table" videos captured with smartphones, which are nowadays a commonly owned accessory. The mature Structure-from-Motion (SfM) framework then provides 3D annotations by tracking cameras and reconstructs a dense 3D point cloud capturing the object surface.

To this end, we collected almost 19,000 videos of 50 MS-COCO categories with 1.5 million frames, each annotated with camera pose, where 20% of the videos are annotated with a semi-manually verified high-resolution 3D point cloud. As such, the dataset exceeds alternatives [9, 1, 21] in terms of number of categories and objects.

Our work is an extension of the dataset from [26]. Here, we significantly increase the dataset size from less than 10 categories to 50 and, more importantly, conduct a human-in-the-loop check ensuring reliable accuracy of all cameras. Finally, the dataset from [26] did not contain any point cloud annotations, the examples of which are in fig. 1.

We also propose a novel *NerFormer* model that, given a small number of input source views, learns to reconstruct object categories in our dataset. NerFormer mates two of the main workhorses of machine learning and 3D computer vision: Transformers [65] and neural implicit rendering [43]. Specifically, given a set of 3D points along a rendering ray, features are sampled from known images and stacked into a tensor. The latter is in fact a ray-depth-ordered sequence of sets of sampled features which admits processing with a sequence-to-sequence Transformer. Therefore, by means of alternating feature pooling attention and ray-wise attention layers, NerFormer learns to jointly aggregate features from the source views and raymarch over them.

Importantly, NerFormer outperforms a total of 14 baselines which leverage the most common shape representations to date. As such, our paper conducts one of the first truly large-scale evaluations of learning 3D object categories in the wild.

## 2. Related Work

In this section we review current 3D datasets and related methods in the areas of single-image reconstruction, generative modelling and novel-view synthesis.

**3D object datasets** The main enabler of early works in 3D reconstruction was the synthetic ShapeNet [6] dataset. Pascal3D [72] introduced a real world dataset providing pose estimation for images, but only approximate 3D models. Choi et al. [9] provide a large set of realobject-centric RGB-D videos, however only a small subset is annotated with 3D models and cameras. Increasing the number of categories and objects, Objectron [1] contains object-centric videos, object/camera poses, point clouds, surface planes and 3D bounding boxes. Unfortunately, only a limited number of object-centric videos cover full 360 degrees. Our dataset further increases the number of categories by a factor of 5 and covers the full 360 degree range. Requiring 3D scanners, GSO [21] provides clean full 3D models including textures of real world objects. Due to the requirement of 3D scanning, it contains less objects. A detailed comparison of the aforementioned datasets is presented in tab. 1.

| Dataset | [6] | [72] | [9] | [1] | [21] | [26] | **Ours** |
|---|---|---|---|---|---|---|---|
| Real | × | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ |
| # Categories | 55 | 12 | 9 | 9 | NA | 7 | 50 |
| # Objects | 51k | 36k | 2k | 15k | 2k | 2k | 19k |
| 3D GT | full | approx. | depth | pcl, box | full | pcl | pcl |
| Multi-view | full | none | full | limited | full | full | full |

Table 1: Common Objects in 3D compared to alternatives. The "pcl" abbreviation stands for "point clouds".

**3D reconstruction** A vast amount of methods studied fully supervised 3D reconstruction from 2D images making use of several different representations: voxel grids [10, 18], meshes [19, 67], point clouds [14, 73], signed distance fields, [49, 2] or continuous occupancy fields [42, 8, 17, 16].

Methods overcoming the need for 3D supervision are based on differentiable rendering allowing for comparison of 2D images rather than 3D shapes [51, 63, 31]. Generating images from meshes was achieved via soft rasterization in [32, 40, 61, 30, 7, 36, 77, 20, 67]. Volumetric represenations are projected to 2D via differentiable raymarching [25, 15, 41, 43, 39, 55] or in a similar fashion via sphere tracing for signed distance fields [45, 74]. [28] introduce differentiable point clouds. Another line of work focuses on *neural rendering*, i.e. neural networks are trained to approximate the rendering function [44, 56]. [35, 34] map pixels to object-specific derformable template shapes. [47, 48] canonically align point clouds of object-categories in an unsupervised fashion, but do not reason about colour. Exploiting symmetries and reasoning about lighting [71], compose appearance, shape and lighting.

Similar to us, [52, 75, 26, 60] utilize per-pixel warp-conditioned embedding [26]. In contrast to our method, multi-view aggregation is handled by averaging over encodings which is prone to noise. Instead, we learn the aggregation by introducing the NerFormer module. Finally, a very recent IBRNet [68] learns to copy existing colors from known views in an IBR fashion [24, 5], whereas our method can hallucinate new colors, which is crucial since we leverage far fewer source views (at most 9).

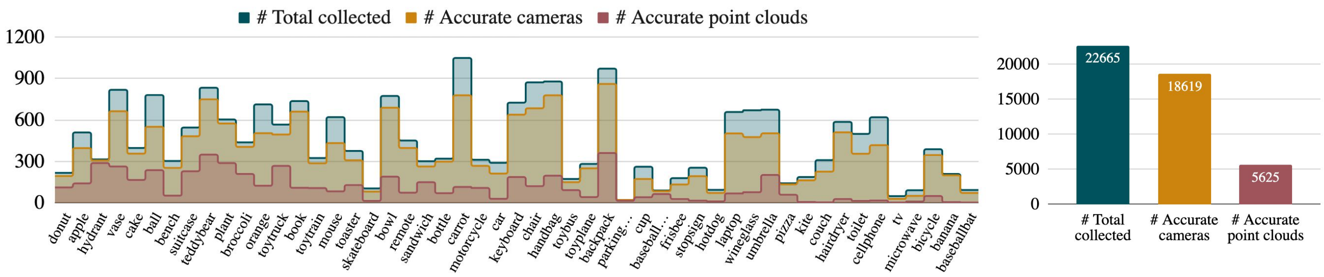


Figure 2: **Statistics of the Common Objects in 3D dataset** reporting the total (right) and per-category (left) numbers of collected videos, and the number of videos with accurate cameras and point clouds.

## 3. Common Objects in 3D

In this section we describe the dataset collection process.

**AMT video collection** In order to scale the collection of object-centric videos, we crowd-sourced it on Amazon Mechanical Turk (AMT). Each AMT task asks a worker to select an object of a given category, place it on a solid surface and take a video where they keep the whole object in view while moving full circle around it. We pre-selected 50 MS-COCO [38] categories (listed in fig. 2) comprising stationary objects that are typically large enough to be reconstructed. The workers were instructed to avoid actions that would hinder the ensuing reconstruction stage such as abrupt movements leading to motion blur. Each video was reviewed to ensure that it fulfills the requirements.

**Generating 3D ground-truth** As explained below, we use off-the-shelf software to produce object masks, camera tracking, and 3D reconstructions for each video. We then semi-automatically filter out poor reconstructions.

**1) Sparse reconstruction** Given the set of valid object-centric videos, we reconstruct the extrinsic (3D location and orientation) and intrinsic (calibration) properties of the cameras that captured the videos. To this end, each video is first converted into a time-ordered sequence of images $\mathcal{V} = (I^i \mid I^i \in \mathbb{R}^{3\times H\times W})_{i=1}^{N_I}$ by extracting $n_I = 100$ frames uniformly spaced in time. The frames are then fed to the COLMAP SfM pipeline [53] which annotates each image with camera projection matrices $\mathcal{P} = (P^i \mid P^i \in \mathbb{R}^{4\times 4})_{i=1}^{N_I}$. Fig. 3 shows example camera tracks together with estimated sparse scene geometries.

**2) Object segmentation** We segment the object in each image $I^i$ with PointRend [33], state-of-the-art instance segmentation method, resulting in a sequence of soft binary masks $\mathcal{M} = (M^i \mid M^i \in [0,1]^{H\times W})_{i=1}^{N_I}$ per video. Note that, while masks aid further dense reconstruction, we have not used them for camera tracking which is typically anchored to the background regions.

**3) Semantic dense reconstruction** Having obtained the camera motions and segmentations, we now describe the process of annotating the captured objects with a 3D surface. We first execute the multi-view stereo (MVS) algorithm of COLMAP [54] to generate per-frame dense depth maps $(D^i \mid D^i \in \mathbb{R}_+^{H\times W})_{i=1}^{N_I}$. We then run COLMAP's point cloud fusion algorithm, which back-projects the depth values masked with $\mathcal{M}$ and retains the points that are consistent across frames, to get a point cloud $\mathcal{P}(\mathcal{V}) = \{\mathbf{x}_i\}_{i=1}^{N_\mathcal{P}}$. Example dense point clouds are visualized in fig. 3.

**4) Labelling reconstruction quality with Human-in-the-loop.** Since our reconstruction pipeline is completely automated, any of the aforementioned steps can fail, resulting in unreliable 3D annotations. We thus incorporate a semi-manual check that filters inaccurate reconstructions.

To this end, we employ an active learning [11] pipeline which cycles between: a) manually labelling the point-cloud and camera-tracking quality; b) retraining a "quality" SVM classifier; and c) automatically estimating the shape/tracking quality of unlabelled videos. Details of this process are deferred to the supplementary.

**5) The dataset** The active-learned SVM aids the final dataset filtering step. As accurate camera annotations are crucial for the majority of recent 3D category reconstruction methods, the first filtering stage completely removes all scenes with camera tracking classified as "inaccurate" (18% of all videos). While all videos that pass the camera check are suitable for training, the scenes that pass both camera and point cloud checks (30% of the videos with accurate cameras) comprise the pool from which evaluation videos are selected. Note, a failure to pass the point cloud check does not entail that the corresponding scene is irreconstructible and should therefore be removed from training – instead this merely implies that the MVS method [54] failed, while other alternatives (see sec. 5) could succeed.

Fig. 2 summarizes the size of Common Objects in 3D. Reconstructing a single video took on average 1h 56 minutes with the majority of execution time spent on GPU-accelerated MVS and correspondence estimation. This amounts to the total of 43819 GPU-hours of reconstruction time distributed to a large GPU cluster.

## 4. Learning 3D categories in the wild

Here, we describe the problem set, give an overview of implicit shape representations, and explain our main techni-

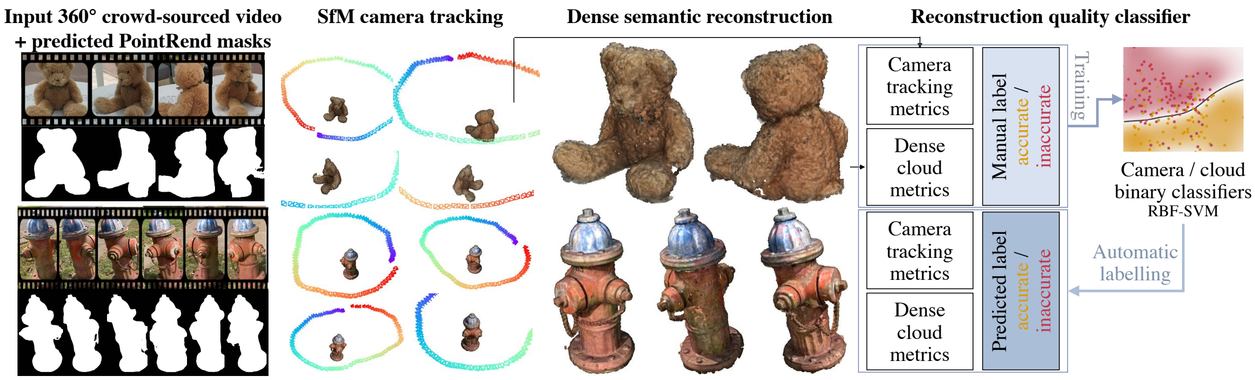


Figure 3: **3D ground truth for Common Objects in 3D** was generated with active learning. Videos are first annotated with cameras and dense point clouds with COLMAP [53]. Given several reconstruction metrics and manual binary annotations ("accurate"/"inaccurate") of a representative subset of reconstructions, we train an SVM that automatically labels all videos.

cal contribution: the NerFormer neural rendering model.

**Problem setting** We tackle the task of generating a representation of appearance and geometry of an object given a small number of its observed source RGB views $\{I_i^{\text{src}}\}_{i=1}^{N_{\text{src}}}$ and their cameras $\{P_i^{\text{src}}\}_{i=1}^{N_{\text{src}}}$ which, in our case, are samples from the set of frames $\mathcal{V}_v$ and cameras $\mathcal{P}_v$ of a video $v$.

In order to visualize the generated shape and appearance, we differentiably render it to a *target* view for which only the camera parameters $P^{\text{tgt}} \in \mathbb{R}^{4\times4}$ are known. This results in a render $\hat{I}^{\text{tgt}}$ which is, during training, incentivised to match the ground truth target view $I^{\text{tgt}}$.

Methods are trained on a dataset of category-centric videos $\{\mathcal{V}_v\}_{v=1}^{n_\mathcal{V}}$, which allows to exploit the regular structure of the category for better generalization to previously unseen objects. Note that, in our experiments (sec. 5), we additionally consider "overfitting" with $n_\mathcal{V} = 1$.

### 4.1. Representing common objects in 3D

In order to predict novel views, all methods need to reconstruct 3D shape in some form. The latter can be represented with a plethora of recent frameworks such as voxel grids [64, 41], implicit surfaces [43, 46, 45], point clouds [28, 69], or meshes [40, 50]. Among those, implicit representations have been successfully applied to reconstructing real scenes and object categories [43, 25, 75], therefore we choose to use them in our pipeline.

**Implicit surface** An implicit surface is defined as a level set $\mathcal{S}_f = \{\mathbf{x} \mid f(\mathbf{x}, \mathbf{z}) = C, C \in \mathbb{R}\}$ of a function $f : \mathbb{R}^3 \times \mathbb{D}_z \mapsto \mathbb{R}$ that accepts a 3D point $\mathbf{x} \in \mathbb{R}^3$ and $D_z$-dimensional latent code $\mathbf{z} \in \mathbb{R}^{D_z}$. In addition to $f$, which represents geometry, a second function $c : \mathbb{R}^3 \times \mathbb{S}^2 \times \mathbb{D}_z \mapsto \mathbb{R}^3$ assigns colors $c(\mathbf{x}, \mathbf{r}, \mathbf{z})$ to the input points $\mathbf{x}$. Note that, in line with recent work [43, 74], $c$ is further conditioned on the 3D direction vector $\mathbf{r} \in \mathbb{S}^2$ from which $\mathbf{x}$ is imaged in order to model viewpoint-dependent effects such as specularities. Finally, both functions $f$ and $c$ depend on the latent code $\mathbf{z}$ encoding geometry and appearance of the scene. Changes in $\mathbf{z}$ alter the level set of $f$ and colors $c$ allowing for representing different instances of object categories.

Following recent successes of [55, 26], we model category-specific 3D shapes with opacity functions $f_o$. Specifically, $f_o$ assigns $f_o(\mathbf{x}, \mathbf{z}) = 0$ to the unoccupied points $\mathbf{x} \notin \mathcal{S}_f$, and $f(\mathbf{x}', \mathbf{z}) > 0$ to surface points $\mathbf{x}' \in \mathcal{S}_f$.

**Neural implicit surfaces** Recent methods implement functions $f_o$ and $c$ as multi-layer perceptrons (MLP) $f_{\text{MLP}}$ and $c_{\text{MLP}}$ [43, 45, 22, 2]. They typically learn shallow specialized networks on top of the shared deep feature extractor $f'_{\text{MLP}}$: $f_{\text{MLP}} = f_{\text{HEAD}} \circ f'_{\text{MLP}}$ and $c_{\text{MLP}} = c_{\text{HEAD}} \circ f'_{\text{MLP}}$.

We depart from representing occupancies with plain MLPs as they process input 3D points $\mathbf{x}$ independently, without any form of spatial reasoning, which is crucial in our case where input source views provide only a partial information about the reconstructed shape.

**Positional embedding** Following [43, 66], avoiding loss of detail, we pre-process the raw 3D coordinates $\mathbf{x}$ with a positional embedding (PE) $\gamma(\mathbf{x}) = [\sin(\mathbf{x}), \cos(\mathbf{x}), ..., \sin(2^{N_f}\mathbf{x}), \cos(2^{N_f}\mathbf{x})] \in \mathbb{R}^{2N_f}$ before feeding to $f_{\text{MLP}}(\mathbf{x}, \mathbf{z})$. While [43] was the first to demonstrate benefits of PE, in sec. 5 we also combine PE with other pre-NeRF methods, such as SRN [56] or DVR [45].

**Rendering an implicit surface** In order to admit image-supervised learning, implicit surfaces are converted into an explicit representation of appearance and geometry with a rendering function $r$. Formally, given a target camera $P^{\text{tgt}}$, the goal is to generate the target image $I^{\text{tgt}} = r(f, c, P^{\text{tgt}})$ which depicts the scene from $P^{\text{tgt}}$'s viewpoint.

We render opacity fields with the Emission-Absorption model (EA). EA renders the RGB value $I_\mathbf{u}^{\text{tgt}} \in \mathbb{R}^3$ at a pixel $\mathbf{u} \in \{1, \dots, W\} \times \{1, ..., H\}$, by evaluating the opacity function $f_o(\mathbf{x}, \mathbf{z})$ for an ordered point-tuple $(\mathbf{x}_i^{\mathbf{r_u}})_{i=1}^{N_S}$ sampled along $\mathbf{u}$'s projection ray $\mathbf{r_u}$ at approximately equidistant intervals $\Delta$. The color $I_\mathbf{u}^{\text{tgt}}(\mathbf{r_u}, \mathbf{z}) = \sum_{i=1}^{N_S} w_i(\mathbf{x}_i^{\mathbf{r_u}}, \mathbf{z}) c(\mathbf{x}_i^{\mathbf{r_u}}, \mathbf{r_u}, \mathbf{z})$ is a sum of per-point colors

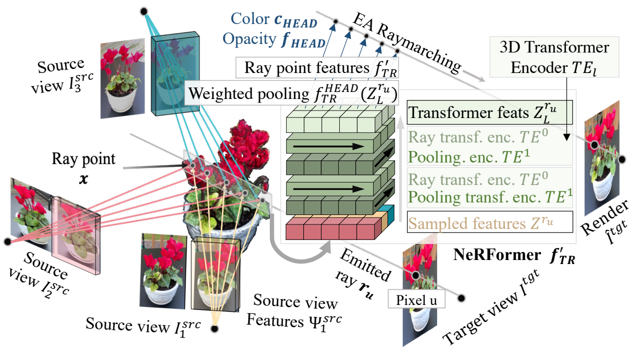


Figure 4: **We propose NerFormer** which jointly learns to pool features from source views and to raymarch by means of a series of transformers alternating between attention along the ray and pooling dimensions.

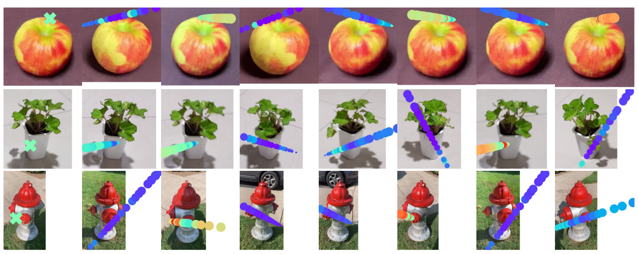

Figure 5: **NerFormer learns to attend to features from source images**. A ray is emitted from a target-image pixel (1st column), and its points are projected to the source views (columns 2-8) from which features are sampled. For each source feature, NerFormer predicts attention (red=high, blue=low) for aggregating to a single source embedding per point. Note how the attention model implicitly learns to pool features from nearby source views.

$c(\mathbf{x}_i^{\mathbf{r_u}}, \mathbf{r_u}, \mathbf{z})$ weighted by the emission-absorption product $w_i = \left(\prod_{j=0}^{i-1} T_j\right)(1 - T_i)$ with $T_i = \exp(-\Delta f_o(\mathbf{x}_i^{\mathbf{r}_u}, \mathbf{z}))$.

## 4.2. Latent shape encoding $\mathbf{z}$

A crucial part of a category-centric reconstructor is the latent embedding $\mathbf{z}$. Early methods [18, 70, 58, 61, 40] predicted a global scene encoding $\mathbf{z}_{\text{global}} = \Phi_{\text{CNN}}(I^{\text{src}})$ with a deep convolutional network $\Phi_{\text{CNN}}$ that solely analyzed the colors of source image pixels. While this approach was successful for the synthetic ShapeNet dataset where shapes are rigidly aligned, it has been recently shown in [26, 75] that such approach is infeasible in real world settings where objects are arbitrarily placed in the 3D space. This is because, unlike in the former case where color-based shape inference is possible since similar images of aligned scenes generate similar 3D shapes, in the latter case, similarly looking images of unaligned scenes can generate vastly different 3D.

**Warp-conditioned embedding** To fix the latter, [26] proposed Warp-Conditioned Embedding (WCE): given a world coordinate point $\mathbf{x}$ and a source view $I^{\text{src}}$ with camera $P^{\text{src}}$, the warp-conditioned embedding $\mathbf{z}_{\text{WCE}} \in \mathbb{R}^{D_z}$

$$\mathbf{z}_{\text{WCE}}(\mathbf{x}, I^{\text{src}}, P^{\text{src}}) = \Psi_{\text{CNN}}(I^{\text{src}})[\pi_{P^{\text{src}}}(\mathbf{x})],$$

is formed by sampling a tensor of source image descriptors $\Psi_{\text{CNN}}(I^{\text{src}}) \in \mathbb{R}^{D_z \times H \times W}$ at a 2D location $\pi_{P^{\text{src}}}(\mathbf{x})$. Here, $\pi_{P^{\text{src}}}(\mathbf{x}) = P^{\text{src}}[\mathbf{x}; 1] = d_{\mathbf{u}}[\mathbf{u}; 1]$ expresses perspective camera projection of a 3D point $\mathbf{x}$ to a pixel $\mathbf{u}$ with depth $d_{\mathbf{u}} \in \mathbb{R}$. Intuitively, since $\mathbf{z}_{\text{WCE}}$ is a function of the world coordinate $\mathbf{x}$, the ensuing implicit $f$ can perceive the specific 3D location of the sampled appearance element in the world coordinates, which in turn enables $f$ to learn invariance to rigid scene misalignment.

In the common case where multiple source views are given, the aggregate WCE $\mathbf{z}^*_{\text{WCE}}(\mathbf{x}, \{I_i^{\text{src}}\}, \{P_i^{\text{src}}\})$ is defined as a concatenation of the mean and standard deviation of the set of view-specific source embeddings $\{\mathbf{z}_{\text{WCE}}(\mathbf{x}, I_i^{\text{src}}, P^{\text{src}})\}_{i=1}^{N_{\text{src}}}$.

**Boosting baselines with WCE** The vast majority of existing methods for learning 3D categories leverage global shape embeddings $\mathbf{z}_{\text{global}}$, which renders them inapplicable to our real dataset. As WCE has been designed to alleviate this critical flaw, and because of its generic nature, in this paper we endow state-of-the-art category-centric 3D reconstruction methods with WCE in order to enable them for learning category-specific models on our real dataset.

To this end, we complement SRN [56], NeuralVolumes [41], and the Implicit Point Cloud (discussed later) with WCE. These extensions are detailed in the supplementary.

## 4.3. Attention is all you nerf

**Limitations of WCE** While [26] has demonstrated that the combination of NeRF and WCE (termed NeRF-WCE) leads to performance improvements, our experiments indicated that a major shortcoming of NeRF-WCE is its inability to deal with cases where parts of the 3D domain are labelled with noisy WCE. This is because NeRF's MLP $f'_{\text{MLP}}$ independently processes each 3D point and, as such, cannot detect failures and recover from them via *spatial reasoning*. The 3D deconvolutions of Neural Volumes [41] are a potential solution, but we found that the method ultimately produces blurry renders due to the limited resolution of the voxel grid. The LSTM marcher of SRN [56] is capable of spatial reasoning, which is however somewhat limited due to the low-capacity of the LSTM cell. Last but not least, a fundamental flaw is that simple averaging of the source-view WCE embeddings can suppress important features.

Our main technical contribution aims to alleviate these issues and follows a combination of two main design guidelines: 1) We replace the $f'_{\text{MLP}}$ with a more powerful architecture *capable of spatial reasoning*. 2) Instead of engineering the WCE aggregation function, we propose to *learn* it.

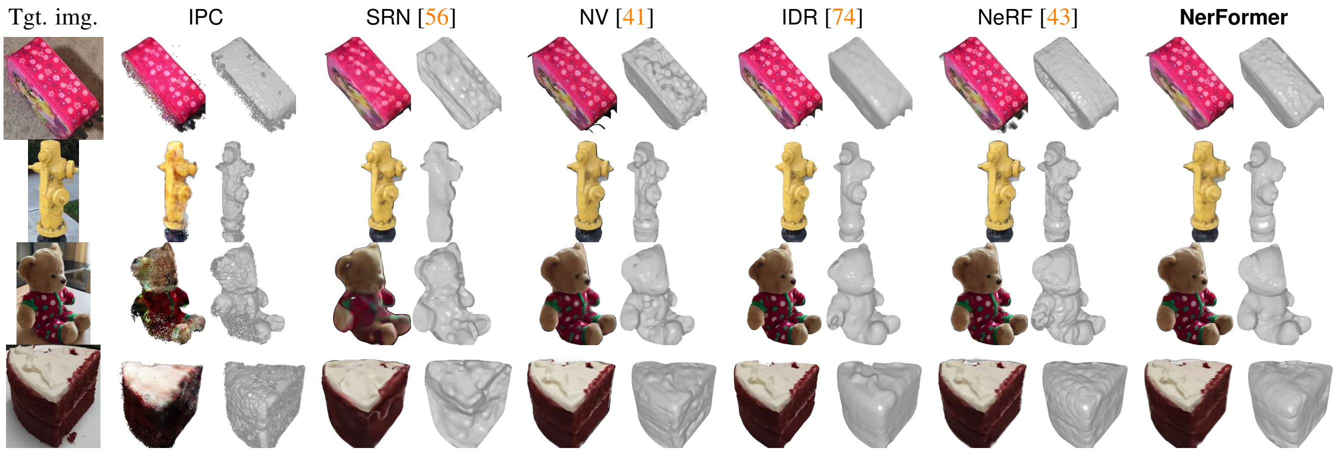


Figure 6: **Single-scene new-view synthesis on Common Objects in 3D** depicting a target image from the training video (left), and corresponding synthesized view generated by IPC, SRN [56], NV [41], IDR [74], NeRF [43], and our NerFormer .

**NerFormer** As a solution, we propose to leverage the popular Transformer architecture [65]. Employing a sequence-to-sequence model is intuitive since the set of all WCE embeddings along a projection ray is in fact a depth-ordered descriptor sequence along the ray dimension, and an unordered sequence along the feature pooling dimension.

Formally, given a ray $\mathbf{r_u}$, we define $Z^{\mathbf{r_u}} \in \mathbb{R}^{N_S \times N_{\text{src}} \times D_{\mathbf{z}}}$ as a stacking of un-aggregated WCEs of all ray-points $\mathbf{x}_j^{\mathbf{r}_u}$:

$$Z^{\mathbf{r_u}} = \left( \left\{ \mathbf{z}_{\text{WCE}}(\mathbf{x}_j^{\mathbf{r}_u}, I_i^{\text{src}}, P_i^{\text{src}}) \right\}_{i=1}^{n_{\text{src}}} \right)_{j=1}^{N_S}. \quad (1)$$

The *NerFormer* module $f'_{\text{TR}}(Z^{\mathbf{r_u}}) = f_{\text{TR}}^{\text{HEAD}} \circ \text{TE}_L \circ \cdots \circ \text{TE}_1(Z^{\mathbf{r_u}})$ replaces the feature backbone $f'_{\text{MLP}}$ (sec. 4.1) with a series of $L$ 3D transformer modules $\text{TE}_l$ terminated by a weighted pooling head $f_{\text{TR}}^{\text{HEAD}}$ (fig. 4). Here, each 3D transformer module $\text{TE}_l$ is a pairing of Transformer Encoder [65] layers $\text{TE}^0(Z)$ and $\text{TE}^1(Z)$:

$$\text{MHA}_l^d(Z_l) = Z'_l = \text{LN}(\text{MHA}_l(Z_l, \text{dim}{=}d) + Z_l) \quad (2)$$

$$\text{TE}_l^d(Z_l) = \text{LN}(\text{MLP}_l(Z'_l) + Z'_l) \quad (3)$$

$$\text{TE}_l(Z_l) = \text{TE}_l^0(\text{TE}_l^1(Z_l)) = Z_{l+1}, \quad (4)$$

$\text{MHA}(Z, \text{dim}{=}d)$ is a multi-head attention layer [65] whose attention vectors span the $d$-th dimension of the input tensor $Z$, MLP is a two-layer MLP with ReLU activation, and LN is Layer Normalization [3]. Intuitively, the alternation between ray and pooling attention of $\text{TE}^0(Z)$ and $\text{TE}^1(Z)$ facilitates learning to jointly aggregate WCE features from the source views and ray-march over them respectively.

Finally, $f'_{\text{TR}}$ is terminated by a weighted pooling head $f_{\text{TR}}^{\text{HEAD}}(Z_L)$ that aggregates the second dimension of $Z_L$ output by the final $L$-th 3D transformer module $\text{TE}_L$ : $f_{\text{TR}}^{\text{HEAD}}(Z_L) = \sum_{i=1}^{n_{\text{src}}} \omega_i(Z_L) Z_L[:, i, :] \in \mathbb{R}^{N_S \times D_{\mathbf{z}}}$, where the weights $\omega_i \in [0, 1], \sum_i \omega_i = 1$ are output by a linear layer with softmax activation. We show $\omega_i$ in fig. 5.

| method | PSNR | LPIPS | $\ell_1^{\text{depth}}$ | IoU | method | PSNR | LPIPS | $\ell_1^{\text{depth}}$ | IoU |
|---|---|---|---|---|---|---|---|---|---|
| **NerFormer** | 23.3 | 0.17 | 0.40 | **0.96** | SRN[56] | 20.4 | 0.21 | 0.60 | 0.93 |
| NeRF+WCE[26] | 21.0 | 0.19 | 0.74 | 0.91 | SRN+WCE+$\gamma$ | 16.9 | 0.30 | 0.60 | 0.75 |
| NeRF[43] | **23.6** | 0.17 | **0.38** | 0.95 | SRN+WCE | 15.8 | 0.26 | 0.64 | 0.80 |
| NV[41] | 22.2 | 0.20 | 0.91 | 0.91 | SRN+$\gamma$ | 16.9 | 0.30 | 0.59 | 0.75 |
| NV+WCE | 18.7 | 0.25 | 0.85 | 0.90 | DVR[45] | 15.0 | 0.33 | 0.89 | 0.68 |
| IDR[74] | 18.5 | **0.15** | 0.81 | 0.92 | DVR+$\gamma$ | 15.8 | 0.38 | 0.74 | 0.65 |
| IPC+WCE | 13.9 | 0.25 | 1.59 | 0.83 | P3DMesh[50] | 17.3 | 0.22 | 0.69 | 0.91 |
| IPC | 13.8 | 0.25 | 1.58 | 0.83 | | | | | |

Table 2: **Single-scene new-view synthesis results on Common Objects in 3D** comparing the baseline approaches [43, 41, 74, 56, 45, 67], IPC, their variants with Warp-conditioned Embeddings (+**WCE**) or Positional Embedding (+$\gamma$), and our NerFormer (the **best** / 2nd best result).

The final opacity and coloring functions of NerFormer are thus $f_o = f_{\text{HEAD}} \circ f'_{\text{TR}}$, $c_o = c_{\text{HEAD}} \circ f'_{\text{TR}}$ respectively, which are rendered with the EA raymarcher (sec. 4.1).

**Technical details** Training minimizes, with Adam (learning rate $5 \cdot 10^{-4}$), a sum of the RGB MSE error $\|I^{\text{tgt}} - \hat{I}^{\text{tgt}}\|^2$ and the binary cross entropy between the rendered alpha mask $\hat{M}^{\text{tgt}}$ and the ground truth mask $M^{\text{tgt}}$. We iterate over batches comprising a randomly sampled target view and 1 to 9 source views of a training video until convergence.

## 5. Experiments

**Datasets** In order to evaluate and train reconstruction methods on our dataset, we split the 18,619 collected CO3D videos into 4 different sets as follows. For each of the 50 categories, we split its videos to a `train` and a `test` set in 9:1 ratio. For each video, we further define a set of frames that are removed from the training set by randomly dividing each `train` video in a 8:2 ratio to 80 `train-known` training and 20 `train-unseen` holdout frames. `test` video frames are split according to the same protocol resulting in `test-known` and `test-unseen` sets. As all base-

| Method | (a) Average statistics: train-unseen PSNR | train-unseen LPIPS | train-unseen $\ell_1^{\text{depth}}$ | train-unseen IoU | test-unseen PSNR | test-unseen LPIPS | test-unseen $\ell_1^{\text{depth}}$ | test-unseen IoU | (b) PSNR @ # source views: train-unseen 9 | train-unseen 7 | train-unseen 5 | train-unseen 3 | train-unseen 1 | test-unseen 9 | test-unseen 7 | test-unseen 5 | test-unseen 3 | test-unseen 1 | (c) PSNR @ target view difficulty: train-unseen easy | train-unseen med. | train-unseen hard | test-unseen easy | test-unseen med. | test-unseen hard |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| **NerFormer** | $\mathbf{\underline{17.9}}$ | 0.26 | 0.87 | 0.82 | $\mathbf{\underline{17.6}}$ | $\underline{0.27}$ | 0.91 | $\underline{0.81}$ | $\mathbf{\underline{19.3}}$ | $\mathbf{\underline{19.0}}$ | $\mathbf{\underline{18.3}}$ | $\underline{17.4}$ | 15.6 | $\mathbf{\underline{18.9}}$ | $\mathbf{\underline{18.6}}$ | $\mathbf{\underline{18.1}}$ | $\mathbf{\underline{17.1}}$ | $\mathbf{\underline{15.1}}$ | $\mathbf{\underline{18.9}}$ | 15.5 | $\underline{14.6}$ | $\mathbf{\underline{18.6}}$ | $\mathbf{\underline{14.9}}$ | 14.7 |
| SRN+WCE+$\gamma$ | $\underline{17.6}$ | $\underline{0.24}$ | $\mathbf{\underline{0.28}}$ | $\underline{0.89}$ | 14.4 | $\underline{0.27}$ | $\underline{0.40}$ | $\underline{0.81}$ | $\underline{18.0}$ | $\underline{18.0}$ | $\underline{17.8}$ | $\mathbf{\underline{17.6}}$ | $\mathbf{\underline{16.8}}$ | 14.6 | 14.5 | 14.6 | 14.5 | $\underline{13.9}$ | $\underline{18.0}$ | $\mathbf{\underline{16.8}}$ | $\mathbf{\underline{16.0}}$ | 14.7 | 13.6 | $\mathbf{\underline{15.1}}$ |
| SRN+WCE | 16.6 | 0.26 | $\underline{0.31}$ | 0.87 | $\underline{14.6}$ | $\underline{0.27}$ | $\mathbf{\underline{0.36}}$ | $\mathbf{\underline{0.82}}$ | 17.0 | 17.0 | 16.7 | 16.4 | 15.8 | $\underline{14.9}$ | $\underline{14.8}$ | $\underline{14.8}$ | 14.6 | $\underline{13.9}$ | 16.9 | $\underline{15.7}$ | 14.5 | $\underline{14.9}$ | $\underline{13.7}$ | $\underline{14.8}$ |
| NeRF+WCE[26] | 14.3 | 0.27 | 2.14 | 0.72 | 13.8 | $\underline{0.27}$ | 2.23 | 0.70 | 14.3 | 15.0 | 14.9 | 14.7 | 14.2 | 12.6 | 14.5 | 14.4 | 14.2 | 13.8 | 12.1 | 13.6 | 13.6 | 14.4 | 13.0 | 13.0 |
| IPC+WCE | 14.1 | 0.36 | 2.12 | 0.70 | 13.5 | 0.37 | 2.24 | 0.69 | 14.4 | 14.4 | 14.2 | 14.1 | 13.4 | 13.8 | 13.8 | 13.7 | 13.6 | 12.6 | 14.4 | 13.7 | 13.4 | 13.8 | 12.8 | 12.2 |
| P3DMesh | 17.2 | $\mathbf{\underline{0.23}}$ | 0.50 | $\mathbf{\underline{0.91}}$ | 12.4 | $\mathbf{\underline{0.26}}$ | 2.49 | 0.69 | 17.6 | 17.5 | 17.4 | 17.1 | $\underline{16.2}$ | 12.6 | 12.5 | 12.5 | 12.5 | 12.1 | 17.5 | $\mathbf{\underline{16.8}}$ | $\mathbf{\underline{16.0}}$ | 12.6 | 11.8 | 13.2 |
| NV+WCE | 12.3 | 0.34 | 2.87 | 0.54 | 11.6 | 0.35 | 3.01 | 0.53 | 12.5 | 12.5 | 12.3 | 12.2 | 12.0 | 11.7 | 11.6 | 11.6 | 11.6 | 11.3 | 12.4 | 12.0 | 13.6 | 11.7 | 11.2 | 12.0 |
| SRN+$\gamma$+AD | $\mathbf{\underline{21.7}}$ | $\mathbf{\underline{0.21}}$ | $\underline{0.31}$ | 0.89 | - | - | - | - | $\mathbf{\underline{21.7}}$ | $\mathbf{\underline{21.7}}$ | $\mathbf{\underline{21.7}}$ | $\mathbf{\underline{21.7}}$ | $\mathbf{\underline{21.8}}$ | - | - | - | - | - | $\mathbf{\underline{21.7}}$ | $\mathbf{\underline{21.3}}$ | $\mathbf{\underline{19.7}}$ | - | - | - |
| SRN[56]+AD | $\underline{21.2}$ | $\mathbf{\underline{0.21}}$ | $\mathbf{\underline{0.23}}$ | $\mathbf{\underline{0.94}}$ | - | - | - | - | $\underline{21.2}$ | $\underline{21.2}$ | $\underline{21.2}$ | $\underline{21.2}$ | $\underline{21.3}$ | - | - | - | - | - | $\underline{21.2}$ | $\underline{20.8}$ | $\underline{19.2}$ | - | - | - |
| NV[41]+AD | 19.7 | $\underline{0.23}$ | 0.41 | $\underline{0.93}$ | - | - | - | - | 19.7 | 19.7 | 19.6 | 19.6 | 19.7 | - | - | - | - | - | 19.7 | 19.3 | 17.7 | - | - | - |
| NeRF+AD | 17.1 | 0.25 | 0.55 | 0.92 | - | - | - | - | 17.1 | 17.0 | 17.0 | 17.0 | 17.1 | - | - | - | - | - | 17.1 | 16.9 | 15.8 | - | - | - |
| P3DMesh[50]+AD | 16.1 | 0.24 | 0.74 | 0.89 | - | - | - | - | 16.1 | 16.1 | 16.1 | 16.1 | 16.2 | - | - | - | - | - | 16.1 | 16.0 | 15.1 | - | - | - |
| IPC+AD | 14.0 | 0.37 | 2.20 | 0.69 | - | - | - | - | 14.4 | 14.3 | 14.1 | 14.0 | 13.3 | - | - | - | - | - | 14.3 | 13.6 | 13.2 | - | - | - |
| IDR+AD | 13.7 | 0.25 | 1.74 | 0.88 | - | - | - | - | 13.7 | 13.7 | 13.7 | 13.7 | 13.8 | - | - | - | - | - | 13.7 | 13.7 | 14.0 | - | - | - |
| DVR+$\gamma$+AD | 7.6 | 0.29 | 3.24 | 0.09 | - | - | - | - | 7.6 | 7.6 | 7.5 | 7.5 | 7.7 | - | - | - | - | - | 7.6 | 7.7 | 9.0 | - | - | - |
| DVR[45]+AD | 7.6 | 0.29 | 3.01 | 0.06 | - | - | - | - | 7.6 | 7.6 | 7.5 | 7.5 | 7.7 | - | - | - | - | - | 7.6 | 7.7 | 9.0 | - | - | - |

Table 3: **Category-centric new-view synthesis results on Common Objects in 3D** comparing mean performance over 10 object categories. Reported are: (a) the main metrics averaged over the `train-unseen`/`test-unseen` set; (b) PSNR averaged over test samples with a specific number of source views and; (c) the samples within one of three viewpoint difficulty bins. Methods with +AD are autodecoders and thus not applicable to the `test-unseen` set. (**$\underline{\text{best}}$** / $\underline{\text{2nd best}}$)

lines require a bounded scene coordinate frame, we zero-center each point cloud and normalize the scale by dividing with an average over per-axis standard deviations.

**Evaluation protocol** All evaluations focus on new RGB+D-view synthesis where a method, given a known target camera $P^{\text{tgt}}$ and a set of source views $\{I_i^{\text{src}}\}_{i=1}^{n_{\text{src}}}$, renders new target RGB view $\hat{I}^{\text{tgt}}$ and depth $\hat{D}^{\text{tgt}}$. Four metrics are reported: The peak-signal-to-noise ratio (PSNR) and the LPIPS distance [76] between $\hat{I}^{\text{tgt}}$ and the ground truth target RGB frame $I^{\text{tgt}}$, the Jaccard index between the rendered/g.t. object mask $\hat{M}^{\text{tgt}}/M^{\text{tgt}}$ and the $\ell_1$ distance $\ell_1^{\text{depth}}$ between the rendered depth $\hat{D}^{\text{tgt}}$ and the ground truth point cloud depth render $D^{\text{tgt}}$. Note that PSNR and $\ell_1^{\text{depth}}$ are only aggregated over the foreground pixels $\{\mathbf{u}_{\text{fg}} | M^{\text{tgt}}[\mathbf{u}_{\text{fg}}] = 1\}$. Metrics are evaluated at an image resolution of 800x800 and 400x400 pixels for single-scene (sec. 5.2) and category-specific reconstruction (sec. 5.3) respectively.

### 5.1. Evaluated methods

We implemented a total of 15 approaches that leverage the majority of common shape representation types.

**Implicit surfaces** Among implicit surface estimators we selected the seminal opacity-based NeRF; and IDR [74] and DVR [45] expressing shapes as a SDF $f_d$ and render with sphere-tracing. We further benchmark SRN [56] which implements an implicit learned LSTM renderer. Importantly, following secs. 4.1 and 4.2, we evaluate the modifications SRN-$\gamma$, DVR-$\gamma$ that endow SRN and DVR respectively with positional embedding $\gamma$. Furthermore, SRN-$\gamma$-WCE, SRN-WCE, NeRF-WCE [26] complement SRN and NeRF with the Warp-conditioned Embedding [26]. Finally, we also evaluate our NerFormer method.

**Voxel grids** Neural Volumes (NV) [41] represents the State of the Art among voxel grid predictors. Similar to implicit methods, we also combine NV with WCE.

**Point clouds** In order to compare with a point cloud-based method, we devised an Implicit Point Cloud (IPC) baseline which represents shapes with a colored set of 3D points, converts the set into an implicit surface and then renders it with the EA raymarcher. We note that IPC is strongly inspired by SynSin [50, 69] (see supplementary).

**Meshes** We benchmark P3DMesh - the best-performing variant of PyTorch3D's soft mesh rasterizer from [50] inspired by Pixel2Mesh [67], which deforms an initial spherical mesh template with a fixed topology with a series of convolutions on the mesh graph.

### 5.2. Single-scene reconstruction

We first task the approaches to independently reconstruct individual object videos. More specifically, given a `test` video, every baseline is trained to reproduce the video's `test-known` frames (by minimizing method-specific loss functions such as $\ell_2$ RGB error), and evaluated by comparing the renders from the given `test-known` camera viewpoints to the corresponding ground truth images/depths/masks. Since training on all ~2k `test` videos is prohibitively expensive (each baseline trains at least for 24 hours on a single GPU), we test on 40 randomly selected `test` videos (two from each of 20 random classes). Quantitative / qualitative results are in fig. 6 / Tab. 2.

NerFormer is either the best or the second best across all metrics. NeRF+WCE is beaten by vanilla NeRF which suggests that the noisy WCE embeddings can hurt performance without NerFormer 's spatial reasoning. Interestingly, IDR's realistic, but less detailed renders win in terms of LPIPS, but are inferior in PSNR. Furthermore, we observed a large

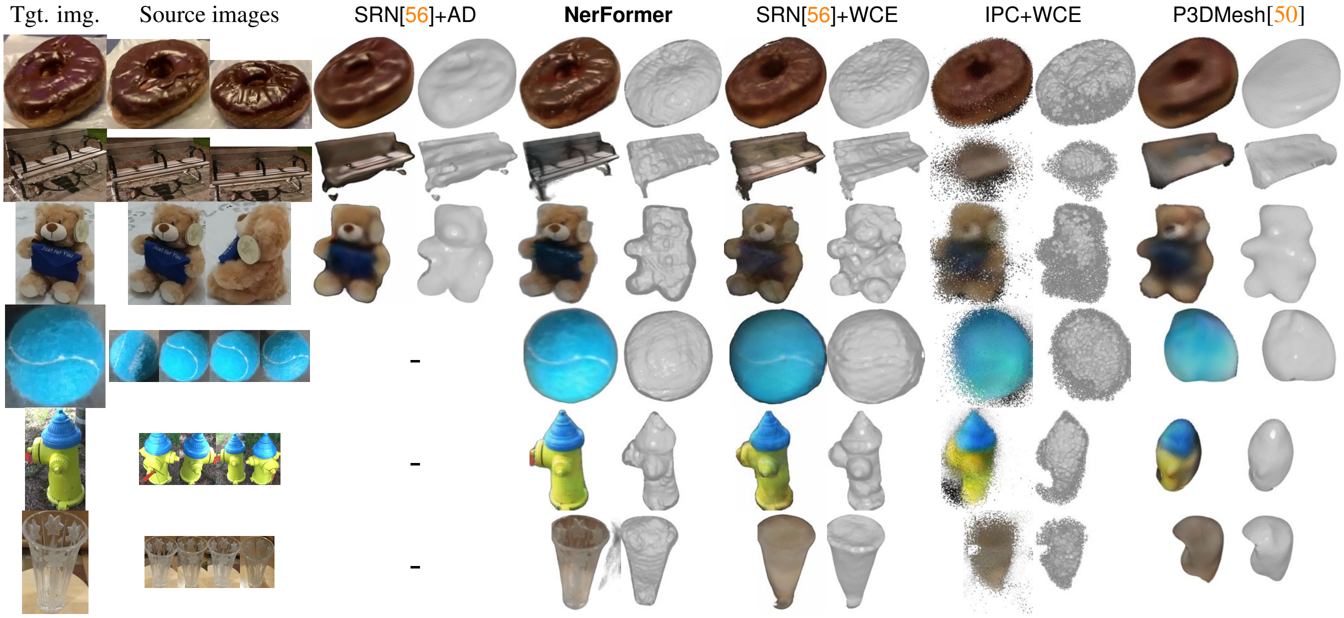


Figure 7: **Category-centric 3D reconstruction on Common Objects in 3D** depicting a target image, known source images $\{I_i^{\text{src}}\}$, and a synthesized new view. The first/last 3 rows are from the `train`/`test` set (SRN+AD is not applicable to `test`).

discrepancy between the train and test PSNR of $\gamma$/WCE-endowed SRN. This shows that, for the single-scene setting, increasing the model expressivity with WCE or $\gamma$ can lead to overfitting to the training views.

### 5.3. Learning 3D Object Categories

Our main task is learning category-centric 3D models. In more detail, a single model is trained on all `train-known` frames of an object category and evaluated on 1000 randomly generated test samples from `train-unseen` and `test-unseen` sets of the category. Each train or test sample is composed of a single target image $I^{\text{tgt}}$ and randomly selected source views $\{I^{\text{src}}\}_i^{n_{\text{src}}}$, where $n_{\text{src}}$ is randomly picked from $\{1, 3, 5, 7, 9\}$. Since training a model for each of the 50 object categories is prohibitively expensive (each method takes at least 7 days to train), we chose a subset of 10 categories for evaluation.

**Baselines** For each method we evaluate the ability to represent the shape space of training object instances by turning it into an *autodecoder* [4] which, in an encoder-less manner, learns a separate latent embedding $\mathbf{z}_{\text{scene}}(\text{sequence}_{\text{ID}})$ for each `train` scene as a free training parameter. Since autodecoders (abbreviated with the **+AD** suffix) only represent the `train` set, we further compare to all WCE-based methods which can additionally reconstruct `test` videos. Note that, as remarked in [26] and sec. 4.2, the alternative encoding $\mathbf{z}_{\text{global}} = \Phi_{\text{CNN}}(I^{\text{src}})$ is deemed to fail due to the world coordinate ambiguity of the training SfM reconstructions.

**Results** Quantitative and qualitative comparisons are shown in tab. 3 and fig. 7 respectively. Besides average metrics over both test sets, we also analyze the dependence on the number of available source views, and on the difficulty of the target view. For the latter, test frames are annotated with a measure of their distance to the set of available source views, and then split into 3 different difficulty bins. Average per-bin PSNR is reported. The supplementary details the distance metric and difficulty bins.

While SRN+AD has the best performance across all metrics on `test-known`; on the `test-unseen` set, where SRN+AD is inapplicable, our NerFormer is the best for most color metrics. Among implicit methods, the SDF-based DVR and IDR are outperformed by the opacity-based NeRF, with both DVR+WCE and IDR+WCE failing to converge. This is likely because regressing SDFs is more challenging than classifying 3D space with binary labels opaque/transparent. Finally, we observed poor performance of P3DMesh, probably due to the inability of meshes to represent complicated real geometries and textures.

## 6. Conclusion

We have introduced Common Objects in 3D (CO3D), a dataset of in-the-wild object-centric videos capturing 50 object categories with camera and point cloud annotations.

We further contributed NerFormer which is a marriage between Transformer and neural implicit rendering that can reconstruct 3D object categories from CO3D with better accuracy than a total of 14 other tested baselines.

The CO3D collection effort still continues at a steady pace of $\sim$500 videos per week which we plan to release in the near future.

## References


[1] Adel Ahmadyan, Liangkai Zhang, Jianing Wei, Artsiom Ablavatski, and Matthias Grundmann. Objectron: A large scale dataset of object-centric videos in the wild with pose annotations. *arXiv preprint arXiv:2012.09988*, 2020. 2

[2] Matan Atzmon and Yaron Lipman. Sal: Sign agnostic learning of shapes from raw data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2565–2574, 2020. 2, 4

[3] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv*, 2016. 6

[4] Piotr Bojanowski, Armand Joulin, David Lopez-Paz, and Arthur Szlam. Optimizing the latent space of generative networks. *arXiv*, 2017. 8

[5] Chris Buehler, Michael Bosse, Leonard McMillan, Steven Gortler, and Michael Cohen. Unstructured lumigraph rendering. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 425–432, 2001. 2

[6] Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*, 2015. 1, 2

[7] Wenzheng Chen, Jun Gao, Huan Ling, Edward J. Smith, Jaakko Lehtinen, Alec Jacobson, and Sanja Fidler. Learning to Predict 3D Objects with an Interpolation-based Differentiable Renderer. In *Proc. NeurIPS*, 2019. 2

[8] Wenzheng Chen, Huan Ling, Jun Gao, Edward Smith, Jaakko Lehtinen, Alec Jacobson, and Sanja Fidler. Learning to predict 3d objects with an interpolation-based differentiable renderer. In *Advances in Neural Information Processing Systems*, pages 9609–9619, 2019. 1, 2

[9] Sungjoon Choi, Qian-Yi Zhou, Stephen Miller, and Vladlen Koltun. A large dataset of object scans. *arXiv preprint arXiv:1602.02481*, 2016. 2

[10] Christopher B Choy, Danfei Xu, JunYoung Gwak, Kevin Chen, and Silvio Savarese. 3d-r2n2: A unified approach for single and multi-view 3d object reconstruction. In *European conference on computer vision*, pages 628–644. Springer, 2016. 2

[11] David Cohn, Les Atlas, and Richard Ladner. Improving generalization with active learning. *Machine learning*, 15(2):201–221, 1994. 3

[12] Corinna Cortes and Vladimir Vapnik. Support-vector networks. *Machine Learning*, 20(3), 1995. 17

[13] Alexey Dosovitskiy and Thomas Brox. Generating images with perceptual similarity metrics based on deep networks. In *Proc. NeurIPS*, 2016. 1

[14] Haoqiang Fan, Hao Su, and Leonidas J Guibas. A point set generation network for 3d object reconstruction from a single image. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 605–613, 2017. 2

[15] Matheus Gadelha, Subhransu Maji, and Rui Wang. 3d shape induction from 2d views of multiple objects. In *2017 International Conference on 3D Vision (3DV)*, pages 402–411. IEEE, 2017. 2

[16] Kyle Genova, Forrester Cole, Avneesh Sud, Aaron Sarna, and Thomas Funkhouser. Local deep implicit functions for 3d shape. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4857–4866, 2020. 2

[17] Kyle Genova, Forrester Cole, Daniel Vlasic, Aaron Sarna, William T Freeman, and Thomas Funkhouser. Learning shape templates with structured implicit functions. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 7154–7164, 2019. 2

[18] Rohit Girdhar, David F Fouhey, Mikel Rodriguez, and Abhinav Gupta. Learning a predictable and generative vector representation for objects. In *European Conference on Computer Vision*, pages 484–499. Springer, 2016. 2, 5

[19] Georgia Gkioxari, Justin Johnson, and Jitendra Malik. Mesh R-CNN. In *Proc. ICCV*, 2019. 2

[20] Shubham Goel, Angjoo Kanazawa, and Jitendra Malik. Shape and viewpoint without keypoints. *arXiv preprint arXiv:2007.10982*, 2020. 2

[21] GoogleResearch. Google scanned objects, September. 2

[22] David Ha, Andrew M. Dai, and Quoc V. Le. HyperNetworks. 2016. 4

[23] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *arXiv preprint arXiv:1512.03385*, 2015. 14

[24] Peter Hedman, Julien Philip, True Price, Jan-Michael Frahm, George Drettakis, and Gabriel Brostow. Deep blending for free-viewpoint image-based rendering. *ACM Transactions on Graphics (SIGGRAPH Asia Conference Proceedings)*, 37(6), November 2018. 2

[25] Philipp Henzler, Niloy Mitra, and Tobias Ritschel. Escaping plato's cave using adversarial training: 3d shape from unstructured 2d image collections. In *Proc. ICCV*, 2019. 1, 2, 4

[26] Philipp Henzler, Jeremy Reizenstein, Patrick Labatut, Roman Shapovalov, Tobias Ritschel, Andrea Vedaldi, and David Novotny. Unsupervised learning of 3d object categories from videos in the wild. *arXiv*, 2021. 1, 2, 4, 5, 6, 7, 8, 13

[27] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Computation*, 9(8), 1997. 15

[28] Eldar Insafutdinov and Alexey Dosovitskiy. Unsupervised learning of shape and pose with differentiable point clouds. In *Advances in neural information processing systems*, pages 2802–2812, 2018. 2, 4

[29] Hanbyul Joo, Tomas Simon, Xulong Li, Hao Liu, Lei Tan, Lin Gui, Sean Banerjee, Timothy Godisart, Bart C. Nabbe, Iain A. Matthews, Takeo Kanade, Shohei Nobuhara, and Yaser Sheikh. Panoptic studio: A massively multiview system for social interaction capture. *PAMI*, 41(1), 2019. 2

[30] Angjoo Kanazawa, Shubham Tulsiani, Alexei A. Efros, and Jitendra Malik. Learning category-specific mesh reconstruction from image collections. In *Proc. ECCV*, 2018. 2

[31] Abhishek Kar, Christian Häne, and Jitendra Malik. Learning a multi-view stereo machine. In *Advances in neural information processing systems*, pages 365–376, 2017. 2, 16

[32] Hiroharu Kato, Yoshitaka Ushiku, and Tatsuya Harada. Neural 3d mesh renderer. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3907–3916, 2018. 2

[33] Alexander Kirillov, Yuxin Wu, Kaiming He, and Ross Girshick. Pointrend: Image segmentation as rendering. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020. 3

[34] Nilesh Kulkarni, Abhinav Gupta, David F. Fouhey, and Shubham Tulsiani. Articulation-aware canonical surface mapping. In *Proc. CVPR*, pages 449–458, 2020. 2

[35] Nilesh Kulkarni, Abhinav Gupta, and Shubham Tulsiani. Canonical surface mapping via geometric cycle consistency. In *Proc. ICCV*, 2019. 2

[36] Xueting 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. *Proc. ECCV*, 2020. 2, 17

[37] Joseph J. Lim, Hamed Pirsiavash, and Antonio Torralba. Parsing IKEA Objects: Fine Pose Estimation. *Proc. ICCV*, 2013. 1

[38] Tsung-Yi Lin, Michael Maire, Serge J. Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft COCO: common objects in context. In *Proc. ECCV*, 2014. 3

[39] Lingjie Liu, Jiatao Gu, Kyaw Zaw Lin, Tat-Seng Chua, and Christian Theobalt. Neural sparse voxel fields. *arXiv preprint arXiv:2007.11571*, 2020. 2

[40] Shichen Liu, Tianye Li, Weikai Chen, and Hao Li. Soft rasterizer: A differentiable renderer for image-based 3D reasoning. *arXiv.cs*, abs/1904.01786, 2019. 2, 4, 5, 15

[41] Stephen Lombardi, Tomas Simon, Jason Saragih, Gabriel Schwartz, Andreas Lehrmann, and Yaser Sheikh. Neural volumes: Learning dynamic renderable volumes from images. *ACM Trans. Graph.*, 38(4):65:1–65:14, July 2019. 2, 4, 5, 6, 7, 13, 14, 15, 16

[42] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4460–4470, 2019. 2

[43] 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. *arXiv preprint arXiv:2003.08934*, 2020. 2, 4, 6, 13, 14

[44] Thu Nguyen-Phuoc, Chuan Li, Lucas Theis, Christian Richardt, and Yong-Liang Yang. HoloGAN: Unsupervised learning of 3D representations from natural images. *arXiv.cs*, abs/1904.01326, 2019. 2

[45] Michael Niemeyer, Lars Mescheder, Michael Oechsle, and Andreas Geiger. Differentiable volumetric rendering: Learning implicit 3d representations without 3d supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3504–3515, 2020. 2, 4, 6, 7, 13, 15

[46] Michael Niemeyer, Lars M. Mescheder, Michael Oechsle, and Andreas Geiger. Occupancy flow: 4d reconstruction by learning particle dynamics. In *Proc. ICCV*, 2019. 1, 4

[47] David Novotny, Diane Larlus, and Andrea Vedaldi. Learning 3d object categories by looking around them. In *Proc. ICCV*, 2017. 2

[48] David Novotný, Diane Larlus, and Andrea Vedaldi. Capturing the geometry of object categories from video supervision. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2018. 2

[49] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. Deepsdf: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 165–174, 2019. 2

[50] Nikhila Ravi, Jeremy Reizenstein, David Novotny, Taylor Gordon, Wan-Yen Lo, Justin Johnson, and Georgia Gkioxari. Accelerating 3d deep learning with pytorch3d. *arXiv:2007.08501*, 2020. 4, 6, 7, 8, 13, 14, 15, 16

[51] Danilo Jimenez Rezende, SM Ali Eslami, Shakir Mohamed, Peter Battaglia, Max Jaderberg, and Nicolas Heess. Unsupervised learning of 3d structure from images. In *Advances in neural information processing systems*, pages 4996–5004, 2016. 2

[52] Shunsuke Saito, , Zeng Huang, Ryota Natsume, Shigeo Morishima, Angjoo Kanazawa, and Hao Li. Pifu: Pixel-aligned implicit function for high-resolution clothed human digitization. *arXiv preprint arXiv:1905.05172*, 2019. 2

[53] Johannes Lutz Schönberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 3, 4

[54] Johannes Lutz Schönberger, Enliang Zheng, Marc Pollefeys, and Jan-Michael Frahm. Pixelwise view selection for unstructured multi-view stereo. In *European Conference on Computer Vision (ECCV)*, 2016. 3

[55] Katja Schwarz, Yiyi Liao, Michael Niemeyer, and Andreas Geiger. Graf: Generative radiance fields for 3d-aware image synthesis. *arXiv preprint arXiv:2007.02442*, 2020. 2, 4

[56] Vincent Sitzmann, Michael Zollhöfer, and Gordon Wetzstein. Scene representation networks: Continuous 3d-structure-aware neural scene representations. *CoRR*, abs/1906.01618, 2019. 1, 2, 4, 5, 6, 7, 8, 13, 14, 15

[57] Carole H Sudre, Wenqi Li, Tom Vercauteren, Sebastien Ourselin, and M Jorge Cardoso. Generalised dice overlap as a deep learning loss function for highly unbalanced segmentations. In *Deep learning in medical image analysis and multimodal learning for clinical decision support*. 2017. 17

[58] Maxim Tatarchenko, Alexey Dosovitskiy, and Thomas Brox. Multi-view 3d models from single images with a convolutional network. In *Proc. ECCV*, 2016. 5

[59] Simon Tong and Edward Y. Chang. Support vector machine active learning for image retrieval. In *ACM Multimedia*, 2001. 17

[60] Alex Trevithick and Bo Yang. Grf: Learning a general radiance field for 3d scene representation and rendering. *arXiv preprint arXiv:2010.04595*, 2020. 2

[61] Shubham Tulsiani, Alexei A Efros, and Jitendra Malik. Multi-view consistency as supervisory signal for learning shape and pose prediction. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2897–2905, 2018. 2, 5

[62] Shubham Tulsiani, Abhishek Kar, Joao Carreira, and Jitendra Malik. Learning category-specific deformable 3D models for object reconstruction. *PAMI*, 39(4):719–731, 2017. 1

[63] Shubham Tulsiani, Tinghui Zhou, Alexei A Efros, and Jitendra Malik. Multi-view supervision for single-view reconstruction via differentiable ray consistency. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2626–2634, 2017. 2

[64] Shubham Tulsiani, Tinghui Zhou, Alexei A. Efros, and Jitendra Malik. Multi-view supervision for single-view reconstruction via differentiable ray consistency. In *Proc. CVPR*, pages 209–217, 2017. 4

[65] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Proc. NeurIPS*, 2017. 2, 6

[66] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NIPS*, 2017. 4

[67] Nanyang Wang, Yinda Zhang, Zhuwen Li, Yanwei Fu, Wei Liu, and Yu-Gang Jiang. Pixel2mesh: Generating 3d mesh models from single rgb images. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 52–67, 2018. 2, 6, 7, 15

[68] Qianqian Wang, Zhicheng Wang, Kyle Genova, Pratul Srinivasan, Howard Zhou, Jonathan T Barron, Ricardo Martin-Brualla, Noah Snavely, and Thomas Funkhouser. Ibrnet: Learning multi-view image-based rendering. *arXiv*, 2021. 2

[69] Olivia Wiles, Georgia Gkioxari, Richard Szeliski, and Justin Johnson. Synsin: End-to-end view synthesis from a single image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7467–7477, 2020. 4, 7, 16

[70] Jiajun Wu, Chengkai Zhang, Tianfan Xue, Bill Freeman, and Josh Tenenbaum. Learning a probabilistic latent space of object shapes via 3D generative-adversarial modeling. In *Proc. NeurIPS*, 2016. 5

[71] Shangzhe Wu, Christian Rupprecht, and Andrea Vedaldi. Unsupervised learning of probably symmetric deformable 3d objects from images in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1–10, 2020. 2

[72] Yu Xiang, Roozbeh Mottaghi, and Silvio Savarese. Beyond PASCAL: A benchmark for 3D object detection in the wild. In *Proc. WACV*, 2014. 2

[73] Guandao Yang, Xun Huang, Zekun Hao, Ming-Yu Liu, Serge Belongie, and Bharath Hariharan. Pointflow: 3d point cloud generation with continuous normalizing flows. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4541–4550, 2019. 2

[74] Lior Yariv, Yoni Kasten, Dror Moran, Meirav Galun, Matan Atzmon, Basri Ronen, and Yaron Lipman. Multiview neural surface reconstruction by disentangling geometry and appearance. *Proc. NIPS*, 2020. 2, 4, 6, 7, 13, 15

[75] Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. pixelnerf: Neural radiance fields from one or few images. *arXiv*, 2020. 2, 4, 5

[76] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proc. CVPR*, 2018. 7

[77] Yuxuan Zhang, Wenzheng Chen, Huan Ling, Jun Gao, Yinan Zhang, Antonio Torralba, and Sanja Fidler. Image gans meet differentiable rendering for inverse graphics and interpretable 3d neural rendering. *arXiv preprint arXiv:2010.09125*, 2020. 2