

Seeing Through the Glass: Neural 3D Reconstruction of Object Inside a Transparent Container

Jinguang Tong^{1,2}, Sundaram Muthu², Fahira Afzal Maken², Chuong Nguyen^{1,2}, Hongdong Li¹

¹The Australian National University ²Data61, CSIRO

{jinguang.tong, chuong.nguyen, hongdong.li}@anu.edu.au

{sundaram.muthu, fahira.afzalmaken, chuong.nguyen}@data61.csiro.au

Abstract

In this paper, we define a new problem of recovering the 3D geometry of an object confined in a transparent enclosure. We also propose a novel method for solving this challenging problem. Transparent enclosures pose challenges of multiple light reflections and refractions at the interface between different propagation media e.g. air or glass. These multiple reflections and refractions cause serious image distortions which invalidate the single viewpoint assumption. Hence the 3D geometry of such objects cannot be reliably reconstructed using existing methods, such as traditional structure from motion or modern neural reconstruction methods. We solve this problem by explicitly modeling the scene as two distinct sub-spaces, inside and outside the transparent enclosure. We use an existing neural reconstruction method (NeuS) that implicitly represents the geometry and appearance of the inner subspace. In order to account for complex light interactions, we develop a hybrid rendering strategy that combines volume rendering with ray tracing. We then recover the underlying geometry and appearance of the model by minimizing the difference between the real and rendered images. We evaluate our method on both synthetic and real data. Experiment results show that our method outperforms the state-of-the-art (SOTA) methods. Codes and data will be available at <https://github.com/hirotong/ReNeuS>

1. Introduction

Digitization of museum collections is an emerging topic of increasing interest, especially with the advancement of virtual reality and Metaverse technology. Many famous museums around the world have been building up their own digital collections^{1,2} for online exhibitions. Among these collections, there is a kind of special but important collec-

¹<https://collections.louvre.fr/en/>

²<https://www.metmuseum.org/>

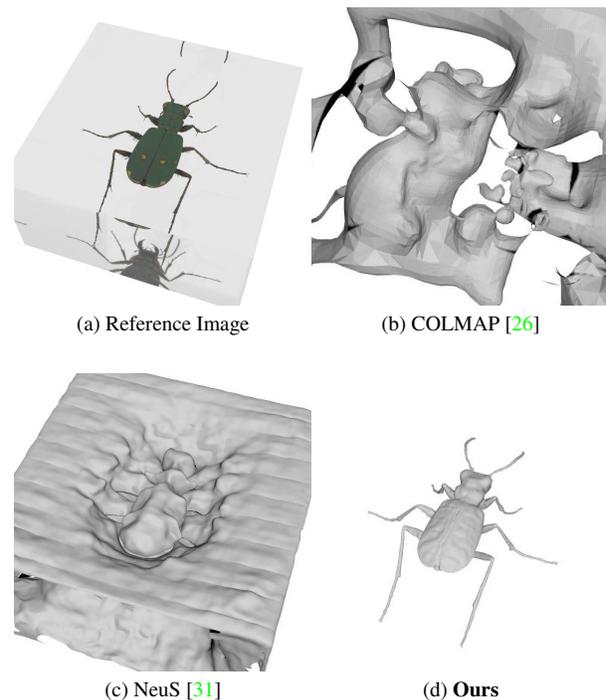


Figure 1. We propose a new research problem of recovering shape of objects inside a transparent container as shown in (a) and solve it by a physically-based neural reconstruction solution. Other methods (b) and (c) without consideration of the optical effects fail to reconstruct, (d) our proposed ReNeuS can retrieve high-quality mesh of the insect object.

tions such as insects [7, 23], human tissues [30], aquatic creatures [6] and other fragile specimens that need to be preserved inside some hard but transparent materials (e.g. resin, glass) as shown in Fig. 1a. In order to digitize such collections, we abstract a distinct research problem which is seeing through the glassy outside and recovering the geometry of the object inside. To be specific, the task is to reconstruct the 3D geometry of an object of interest from multiple 2D views of the scene. The object of interest is usually con-

tained in a transparent enclosure, typically a cuboid or some regular-shaped container.

A major challenge of this task is the severe image distortion caused by repeated light reflection and refraction at the interface of air and transparent blocks. Since it invalidates the single viewpoint assumption [9], most existing multi-view 3D reconstruction methods tend to fail.

One of the most relevant research problems is 3D reconstruction through refractive interfaces, known as Refractive Structure-from-Motion (RSfM) [1, 4]. In RSfM, one or more parallel interfaces between different propagation media, with different refractive indices, separate the optical system from the scene to be captured. Light passing through the optical system refracts only once at each interface causing image distortion. Unlike RSfM methods, we deal with intersecting interfaces where multiple reflections and refractions take place. Another related research topic is reconstruction of transparent objects [13, 17, 19]. In this problem, lights only get refracted twice while entering and exiting the target. In addition to the multiple refractions considered by these methods, our method also tackles the problem of multiple reflections within the transparent enclosure.

Recently, neural implicit representation methods [21, 22, 28, 29, 31] achieve state-of-the-art performance in the task of novel view synthesis and 3D reconstruction showing promising ability to encode the appearance and geometry. However, these methods do not consider reflections. Although NeRFReN [12] attempts to incorporate the reflection effect, it cannot handle more complicated scenarios with multiple refraction and reflection.

In order to handle such complicated cases and make the problem tractable, we make the following two reasonable simplifications:

i. **Known geometry and pose of the transparent block.**

It is not the focus of this paper to estimate the pose of the interface plane since it is either known [1, 4] or calculated as a separate research problem [14] as is typically assumed in RSfM.

ii. **Homogeneous background and ambient lighting.**

Since the appearance of transparent objects is highly affected by the background pattern, several transparent object reconstruction methods [13, 19] obtain ray correspondence with a coded background pattern. To concentrate on recovering the shape of the target object, we further assume monochromatic background with only homogeneous ambient lighting conditions similar to a photography studio.

To handle both **Reflection** and **Refraction** in 3D reconstruction, we propose **ReNeuS**, a novel extension of a neural implicit representation method NeuS [31]. In lieu of

dealing with the whole scene together, a divide-and-conquer strategy is employed by explicitly segmenting the scene into internal space, which is the transparent container containing the object, and an external space which is the surrounding air space. For the internal space, we use NeuS [31] to encode only the geometric and appearance information. For the external space, we assume a homogeneous background with fixed ambient lighting as described in simplification ii. In particular, we use a novel hybrid rendering strategy that combines ray tracing with volume rendering techniques to process the light interactions across the two sub-spaces. The main contributions of this paper are as follows:

- We introduce a new research problem of 3D reconstruction of an object confined in a transparent enclosure, a rectangular box in this paper.
- We develop a new 3D reconstruction method, called ReNeuS, that can handle multiple light reflections and refractions at intersecting interfaces. ReNeuS employs a novel hybrid rendering strategy to combine the light intensities from the segmented inner and outer sub-spaces.
- We evaluate our method on both synthetic as well as a newly captured dataset that contains 10 real insect specimens enclosed in a transparent glass container.

2. Related work

Refractive Structure from Motion The Refractive Structure-from-Motion (RSfM) problem involves recovering 3D scene geometry from underwater images where light refraction at the interface between different propagation media causes distortions which invalidates the assumption of a single viewpoint [10] (Pinhole camera assumption fails). The refractive structure from motion method developed in [15] accurately reconstructs underwater scenes that are captured by cameras confined in an underwater glass housing. This is achieved by incorporating multiple refractions in the camera model that takes place at interfaces of varying densities of glass, air, and water. [4] presents a robust RSfM framework to estimate camera motion and 3D reconstruction using cameras with thin glass interfaces that seal the optical system. To achieve this, [4] considers the refractive interface explicitly by proposing a refractive fundamental matrix. Unlike the methods described in [4, 15], where the camera is confined in front of a piece of glass, our method works with small objects, such as insects, that are embedded within a glassy box.

Transparent Object Reconstruction Acquiring the 3D geometry of transparent objects is challenging as general-purpose scanning and reconstruction techniques cannot handle specular light transport phenomena. 3D shape reconstruction of highly specular and transparent objects is

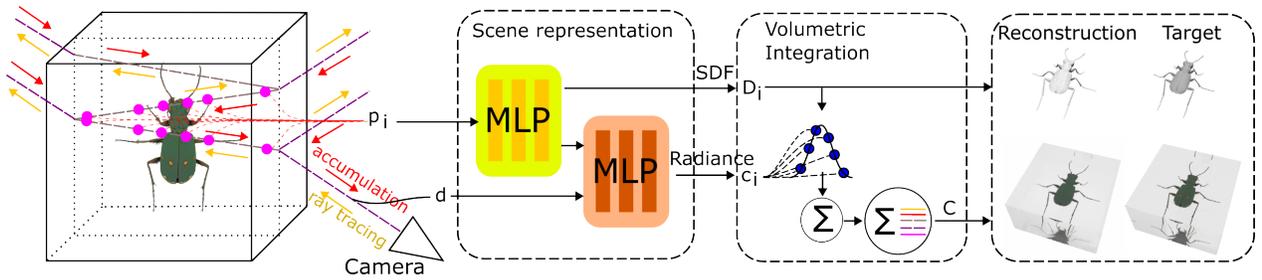


Figure 2. An overview of the ReNeuS framework. The scene is separated into two sub-spaces w.r.t. the interface. The internal scene is represented by two multi-layer perceptrons (MLPs) for both geometry and appearance. For neural rendering, we recursively trace a ray through the scene and collect a list of sub-rays. Our neural implicit representation makes the tracing process controllable. We show a ray tracing process with the depth of recursion $D_{re} = 3$ here. Color accumulation is conducted on irradiance of the sub-rays in the inverse direction of ray tracing. We optimized the network by the difference between rendered image and the ground truth image. Target mesh can be extracted from SDF by Marching Cubes [18]

studied under controlled settings in [32]. The method in [32] mounts the object on a turntable. The model is optimized using the refraction normal consistency, silhouette consistency, and surface smoothness. However, this method requires a known refractive index assumption of the transparent object. [13] presents a method for reconstruction of transparent objects with unknown refractive index. This method partially immerses the object in a liquid to alter the incident light path. Then the object surface is recovered through reconstructing and triangulating such incident light paths. [19] captures fine geometric details of the transparent objects by using correspondences between the camera view rays and the locations on the static grey-coded background. Recently, [17] proposes a neural network to perform 3D reconstruction of transparent objects using a small number of mobile camera images. This is achieved by modelling the reflections and refractions in a rendering layer and a cost volume refinement layer. The rendering layer optimizes the surface normal of the object to guide the point cloud reconstruction.

Neural 3D reconstruction The use of neural networks has been widely used in recent years to predict new views and geometric representations of 3D scenes and objects based on a limited set of images captured with known camera poses. Several works exist that represent 3D structures using deep implicit representations. The main advantage over using conventional voxel grids or meshes is the fact that the implicit representations are continuous and memory efficient to model shape parts of the objects. [21] uses implicit representations of curves and surfaces of any arbitrary topology using level sets. [22] presents NeRF that synthesizes novel views of complex scenes by optimizing a continuous neural radiance field representation of a scene from a set of input images with known camera poses. [12] presents NeRFReN, which is based on NeRF and allows us to model scenes with reflections. In particular, [12] splits the scene into transmitted and reflected components and models the two components with separate neural radiance

fields. Michael *et al.* propose UNISURF [24], a unified framework that couples the benefits of neural implicit surface models (for surface rendering) with those of neural radiance fields (for volume rendering) to reconstruct accurate surfaces from multi-view images. However, this method is limited to representing solid and non-transparent surfaces. The method presented in [33] relaxes the requirement of multi-view camera calibration by describing an end-to-end neural system that is capable of learning 3D geometry, appearance, and cameras from masked 2D images and noisy camera initialization. NeuS [31] performs high-quality reconstruction of objects with severe occlusions and complex geometries by combining radiance fields and occupancy representations. This method integrates both neural volumetric representation of NeRF and neural implicit representations using SDFs to represent the 3D surfaces. The implicit representation greatly improves the NeRF. However, NeuS cannot handle the severe image distortion created due to multiple reflections and refractions of the object placed inside a glass. Hence, our method ReNeuS modifies the NeuS [31] to overcome the problem by using a hybrid rendering strategy.

3. Method

Given a set of N posed images and a known transparent rectangular box geometry, we aim to recover the geometry of the internal object. As shown in Fig. 2, we proposed a neural reconstruction method called ReNeuS which extends NeuS to work on scenes with complex light refractions and reflections. To decompose such a complicated scene, we first divide it into two sub-spaces (inside and outside). Following NeuS [31], the internal space is represented by an implicit neural network which contains an implicit surface representation with signed distance function (SDF) and a volumetric appearance representation. To render a color image of the scene, we explicitly trace its ray path through the scene. For all the sub-rays in the ray path, we retrieve the

radiance of sub-rays according to their location (inner or outer) and then physically accumulate them to get the color of a ray. We call this a hybrid rendering method. All of the network parameters are then optimized by minimizing the difference between the rendered image and the ground truth image.

3.1. Neural Implicit Surface Revisited

Scene representation. Neural Implicit Surface (NeuS) [31] represents scenes with neural networks which contain a neural surface representation of signed distance function (SDF) and a volumetric radiance representation. Two independent multi-layer perceptrons (MLP) are used to encode the scene’s geometry and appearance. To be specific, the geometric MLP: $g_\theta : (\mathbf{x}) \rightarrow (s)$ with θ as parameters, represent the scene as an SDF field that it takes over a point coordinate \mathbf{x} and outputs its signed distance value s . From $g(\mathbf{x})$, the surface of the object \mathcal{S} can be extracted from a zero-level set of its SDF as:

$$\mathcal{S} = \{\mathbf{x} \in \mathbb{R}^3 \mid g_\theta(\mathbf{x}) = 0\} \quad (1)$$

Since the appearance is view-dependent, another appearance MLP: $f_\phi : (\mathbf{x}, \mathbf{v}, \mathbf{n}) \rightarrow (\mathbf{c})$ encodes radiance information of a point \mathbf{c} with its position \mathbf{x} , view direction \mathbf{v} and the normal vector \mathbf{n} .

Volume Rendering Let’s first recall the volume rendering equation. Given a light ray $\ell(t) = \mathbf{o} + t\mathbf{d}$, the color $\mathbf{C}(\ell)$ is calculated by the weighted integration of radiance along the ray by:

$$\hat{\mathbf{C}}(\ell) = \int_{t_n}^{t_f} w(t)c(\mathbf{x}(t), \mathbf{d})dt \quad (2)$$

where $w(t)$ is the weight at point $\mathbf{x}(t)$ and $c(\mathbf{x}(t), \mathbf{d})$ is the radiance emitted from point $\mathbf{x}(t)$ along the direction \mathbf{d} . In neural radiance field based representation [22], the weight is a composition of volume density $\sigma(\mathbf{x}(t))$ and accumulated transmittance $T(t)$:

$$w(t) = T(t)\sigma(\mathbf{x}(t)), \quad (3)$$

$$\text{where } T(t) = \exp\left(-\int_{t_n}^t \sigma(\mathbf{x}(s))ds\right) \quad (4)$$

In order to apply volume rendering on the SDF network, Peng *et al.* [31] derives an *unbiased* and *occlusion-aware* weight function as:

$$w(t) = T(t)\rho(t) \quad (5)$$

with the aid of an opaque density function:

$$\rho(t) = \max\left(\frac{-\frac{d\Phi_s}{dt}(g(\mathbf{x}(t)))}{\Phi_s(g(\mathbf{x}(t)))}, 0\right) \quad (6)$$

where $\Phi_s(x) = (1 + e^{-sx})^{-1}$ is the Sigmoid function. Finally, the integration in Eq. (2) is estimated by using numerical quadrature [20] with hierarchical sampling strategy. For more detail, refer to the original papers [22, 31].

3.2. ReNeuS Scene Representation

Since the problem originates from the heterogeneity of the scene, our solution is to separate the scene by the interface of different mediums (*e.g.* air/glass in our case), and obtain two sub-spaces *internal* \mathbb{S}_{in} and *external* \mathbb{S}_{out} .

Internal space \mathbb{S}_{in} . We adapt NeuS to represent the local internal space \mathbb{S}_{in} which contains the target object. As described in Sec. 3.1, the appearance of \mathbb{S}_{in} is represented by a volumetric radiance field. From the principle of volumetric radiance field [20], light interaction (including absorbing, reflecting *etc.*) is explained by the probability of interaction and has been embedded in the volume rendering equation. In our case, a light ray can go straightly through \mathbb{S}_{in} and its color can be retrieved by volume rendering as Eq. (2).

External space \mathbb{S}_{out} . As our simplification *ii*, the external space is assumed to be empty with homogeneous ambient lighting $\mathbf{C}^{out} = \mathbf{C}^{ambient}$.

3.3. ReNeuS Renderer

To render a color image of the scene, we utilize a novel rendering strategy that combines ray tracing [25] and volume rendering [20] techniques.

Ray tracing. In computer graphics, ray tracing is the process of modelling light transport for rendering algorithms. We slightly modified the ordinary ray tracing method. In ReNeuS, the appearance of internal space \mathbb{S}_{in} is represented as a volumetric radiance field. When tracing a ray through \mathbb{S}_{in} , we do not try to explicitly model the interaction with the target object but let the ray go straightly through the space. This allows us to get much better control of the scene and we only need to consider the light interactions (reflection and refraction) happening on the interfaces of different sub-spaces.

To be specific, for a ray ℓ_i , we calculate the intersection of ℓ_i and the transparent box. If there is no hit, the tracing process will be terminated with only the original ray ℓ_i recorded. If ℓ_i hits the box, we explicitly trace the reflected ray ℓ_i^r and refracted ray ℓ_i^t according to the Law of Reflection and Snell’s Law [3], so that we have:

$$\ell_i^r, \ell_i^t = \mathcal{RE}(\ell_i, \mathbf{n}_i) \quad (7)$$

where \mathbf{n}_i is the surface normal at the intersection and \mathcal{RE} represents the operation of **reflection** and **refraction**. Besides, the reflectance R_i is calculated according to the Fresnel equations under “natural light” assumption [3] and also the transmittance $T_i^{re} = 1 - R_i$. Note the transmittance T_i^{re} is different from $T(t)$ in Eq. (4) that T_i^{re} is the transmittance

of refraction light while $T(t)$ represents the transmittance of the space along ray direction. The color of reflected and refracted rays can be expressed as:

$$\mathbf{C}(\ell_i^r) = R_i \cdot \hat{\mathbf{C}}(\ell_i^r) \quad (8)$$

$$\mathbf{C}(\ell_i^t) = T_i^{re} \cdot \hat{\mathbf{C}}(\ell_i^t) \quad (9)$$

The tracing process is performed recursively on ℓ_i^r and ℓ_i^t and constrained by the depth of recursion D_{re} which is a hyper-parameter in our method. After ray tracing, each original camera ray ℓ_i is extended to \mathbb{L}_i which contains a set of traced rays.

Color Accumulation. The color of a ray emitted from the camera ℓ_i is restored by accumulating that among all the sub-rays in \mathbb{L}_i .

$$\mathbf{C}(\ell_i) = \mathcal{ACC}_{\mathbf{m}_i \in \mathbb{L}_i}(\mathbf{C}(\mathbf{m}_i)) \quad (10)$$

where \mathcal{ACC} means physically-based color accumulation. To do this, we first retrieve the color of each sub-ray $\mathbf{m}_i \in \mathbb{L}_i$. For internal rays, the color is obtained by volume rendering. And for external rays, a pre-defined background color \mathbf{C}^{out} is directly assigned. Next, the ray color is accumulated progressively in the reverse order of ray tracing. For a ray ℓ , if meets the interface, the color $\mathbf{C}(\ell)$ is modified according to Eq. (8) and (9) and when passing through S^{in} , we further multiply the color $\mathbf{C}(\ell)$ with the accumulated transmittance along the ray, computed as:

$$T_\ell = \exp\left(-\int_{t_s}^{t_e} \rho(s) ds\right) \quad (11)$$

where t_s and t_e indicate the start and end of the ray.

Working with Irradiance. Since we use a physical-based renderer, it should work with scene irradiance rather than image intensity. The camera response function (CRF) describes the relationship between image intensity and scene irradiance. To compensate for the nonlinear CRF, we assume that the appearance MLP predicts colors in linear space and adopts a gamma correction function to the final output:

$$I = \mathcal{G}(\mathbf{C}(\ell)), \quad (12)$$

where $\mathcal{G}(\mathbf{C}) = \mathbf{C}^{\frac{1}{2.2}}$ is the gamma correction function.

3.4. Loss Function

Following previous works [22, 31, 33], the parameters of ReReUS are merely optimized and supervised by multi-view images with known poses. A binary mask M_{in} of the transparent enclosure is also used. To be specific, we randomly sample a batch of pixels. For each pixel p in the batch, a camera ray ℓ_p is generated with the corresponding camera pose. After that, we trace the camera ray ℓ_p and get a ray

set \mathbb{L}_p . For each ray in \mathbb{L}_p , if it belongs to S_{in} , we consistently sample n points along the ray and do volume rendering. The color of the pixel \hat{I}_p is retrieved by accumulation and gamma correction as described in Sec. 3.3.

The overall loss function is given below, and described in detail afterwards:

$$\mathcal{L} = \mathcal{L}_{color} + \lambda_1 \mathcal{L}_{trans} + \lambda_2 \mathcal{L}_{reg} \quad (13)$$

Photometric loss. We empirically choose L1 loss between rendered pixel color and the ground truth as:

$$\mathcal{L}_{color} = \frac{1}{|M_{in}|} \sum_{\mathbf{p} \in M_{in}} \|\hat{I}_p - I_p\|_1 \quad (14)$$

Sparsity prior. Based on the crystalline appearance of the box, we exploit a sparsity prior that a considerable part of S_{in} (except the area occupied by the object) should be somehow transparent to the light. We apply this prior by regularizing the transmittance of S_{in} as:

$$\mathcal{L}_{trans} = \frac{1}{|M_{in}|} \sum_{p \in M_{in}} \sum_{\ell \in \mathbb{L}_p \cap S_{in}} \|1 - T_\ell\| \quad (15)$$

This prior is valid in our case since the opaque object usually takes up a limited area of the transparent box. Experiment results show that \mathcal{L}_{trans} helps to get cleaner geometry.

Regularization. The same as NeuS, we apply the Eikonal regularization from [11] on the sampled points to regularize the geometric MLP as:

$$\mathcal{L}_{reg} = \mathbb{E}_{\mathbf{x}} (|\nabla_{\mathbf{x}} g_\theta(\mathbf{x})| - 1)^2, \quad (16)$$

where \mathbf{x} represents all the sampled points on the rays inside S_{in} within a batch.

4. Experiments

4.1. Dataset.

We introduce two (one synthetic and another real) new datasets to evaluate our method and baseline methods since there is no existing dataset dedicated to our problem setting.

Synthetic Dataset. We choose 13 objects with complicated appearances and/or geometry to generate synthetic data. Blender [5] with the physical-based render engine `Cycles` is used to render photorealistic images. We put the object at the center of a box, manually adjust the size of the box to fit the object and then scale it within a unit sphere. The object generally has a Lambertian surface. The material of the box is adjusted by `Principled BSDF` to appear as glass with a refractive index of 1.45. The camera pose is randomly sampled on a full sphere of radius 5. For each camera pose, we not only render an image of the complete scene but also render another image without the transparent box. We render 60 viewpoints with a resolution

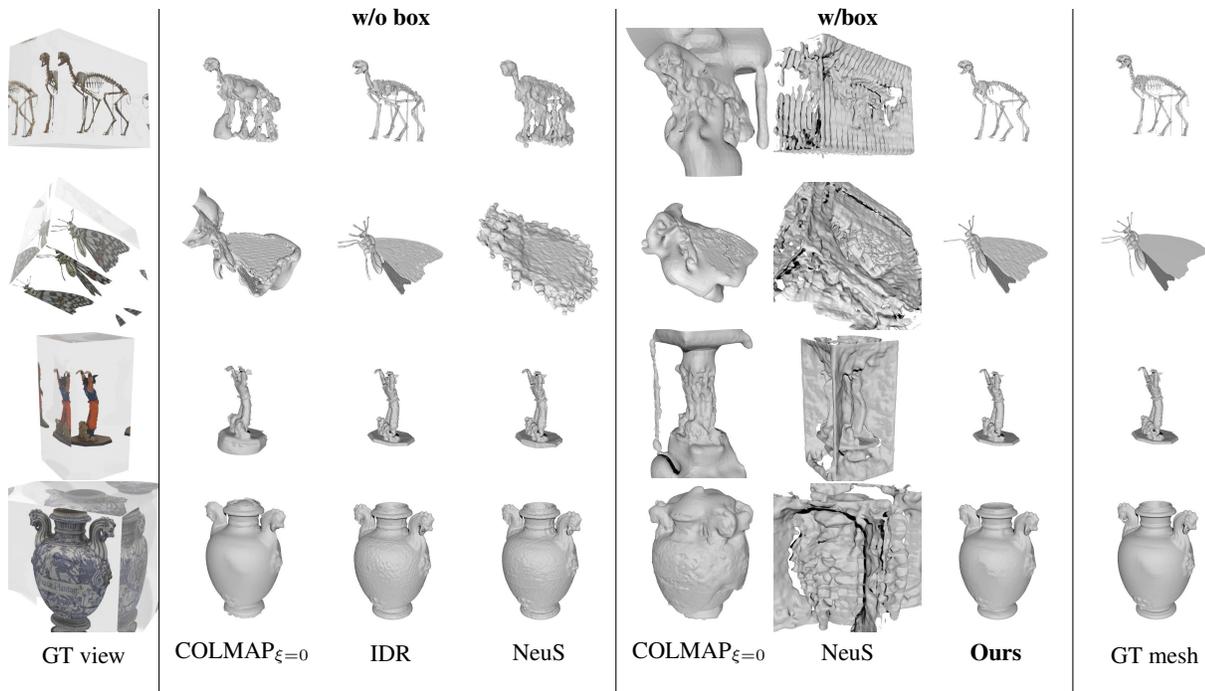


Figure 3. Reconstruction of 3D shapes on the synthetic dataset. We qualitatively compare ReNeuS with baseline methods. Results show that our method successfully achieves high-quality reconstruction with transparent containers where other methods fail.

of 800×800 for each scene, getting two subsets named w/ box and w/o box datasets.

Real Dataset. Another 10 real scenes of insect specimens data are captured. We capture the dataset in a photography light tent where the specimen is placed on a sheet of ChArUco board³ [8] over a rotating display stand. The images are captured on a hemisphere with different azimuths and altitudes. We use ChArUco board to calibrate the intrinsics as well as the pose of the camera. As for the transparent box, we manually label it by edges in the images and calculate its dimensions and pose according to the known camera. The transparent box mask is then generated by projecting the box to 2D images. Finally, we center-crop all images to be 1728×1152 pixels and randomly sample 60 images per scene. You can find more details about the dataset in the supplementary material.

4.2. Baseline methods.

Since there is no specific method aiming at the same problem, we choose several state-of-the-art 3D reconstruction methods as baselines. To make our results more convincing, apart from the w/ box data, we also compare with the baselines on the w/o box dataset.

COLMAP [26, 27] is a popular MVS method. For a fair comparison, we don't run structure-from-motion with COLMAP, but directly use the ground truth camera pose

³https://docs.opencv.org/4.x/df/d4a/tutorial_charuco_detection.html

for reconstruction. We obtain a mesh by applying screened Poisson Surface Reconstruction [16] on the output point cloud of COLMAP.

IDR [33] is a neural rendering-based 3D reconstruction method that achieves results with high fidelity. We run IDR with object mask of w/o box dataset only.

NeuS. Our method is actually an extension of NeuS to an optically complex scene. We compare with NeuS to verify our credit.

4.3. Implementation Details.

We implement ReNeuS on the top of NeuS [31]. We adopt the same network structure for geometric MLP g_θ and appearance MLP f_ϕ as NeuS. The geometric MLP is initialized as [2] to be an ellipsoid-like SDF field. Different from NeuS, we replace the activation function (Softplus) with a periodic activation function SIREN [28] for its stronger capability for neural representation. The input parameters position \mathbf{x} and view direction \mathbf{v} are prompted with positional encoding [22]. The geometry of \mathbb{S}_{out} is represented as a mesh of a regular box with known dimensions. For the real data, we manually calibrate it with the edges of the box. Please refer to the supplementary material for more details. For ReNeuS rendering, we set the recursion depth $D_{re} = 2$ as a trade-off between accuracy and efficiency. The ambient lighting C^{out} is set to $[0.8, 0.8, 0.8]$ (before gamma correction) for synthetic data, and we calibrate that by measuring the average background for real data. For loss function,

Items	w/o box				w/ box			
	COLMAP $_{\xi=0}$	COLMAP $_{\xi=7}$	IDR	NeuS	COLMAP $_{\xi=0}$	COLMAP $_{\xi=7}$	NeuS	ReNeuS
beetle	1.10	3.67	0.76	2.42	15.48	17.55	27.40	1.40
box	4.09	5.28	0.51	3.76	35.38	14.15	15.43	2.29
butterfly	6.10	1.15	0.78	9.98	12.44	18.82	19.85	1.15
coral	4.93	11.23	1.68	2.11	142	17.32	21.18	8.78
coral2	1.81	2.75	1.66	1.42	21.38	29.25	20.75	2.44
dinosaur	2.22	1.18	0.82	2.80	28.15	36.33	15.46	1.47
goku	2.21	2.13	0.81	1.37	22.48	11.09	17.23	1.48
insect	2.56	0.96	0.62	5.89	45.78	35.88	21.71	1.10
insect2	1.34	1.15	0.76	4.07	15.34	18.65	22.11	1.30
lobster	13.27	12.20	17.17	12.97	40.52	26.72	14.61	11.08
shiba	2.33	17.73	2.24	2.22	24.30	13.92	20.30	3.38
statuette	0.85	0.92	0.87	0.79	61.74	22.16	25.44	2.26
vase	1.37	1.30	1.31	1.35	15.02	15.06	24.68	3.68
mean	3.22	4.47	2.28	3.72	35.86	20.83	20.59	3.19

Table 1. Qualitative evaluation on the synthetic dataset. We highlight the best result in terms of **w/o box** and **w/ box** data respectively.

Equation (13), we empirically set $\lambda_1 = 0.1$ and $\lambda_2 = 0.1$ for all experiments. We sample 1024 rays per batch during training, with hierarchical sample [22], totally 128 points (64 for coarse and 64 for fine) are sampled along a ray. No point outside is additionally sampled. We train all the models for totally 200k iterations on a single NVIDIA Tesla P100 GPU. During evaluation time, we extract a mesh from the zero-level set of the geometric MLP g_θ with Marching Cubes [18]. Automatic post-processing is applied by clustering and retaining the mesh with most vertices.

4.4. Evaluation on the synthetic dataset.

We evaluate ReNeuS and other baseline methods on the synthetic dataset. We only run ReNeuS on the w/ box dataset while all the others are evaluated on both w/ box and w/o box datasets. For quantitative comparison, we report the Chamfer- L_1 distance between ground truth and the reconstruction. Due to the unbalanced number of points in the ground truth mesh, we uniformly sample $N = \max(10000, 0.2num_v)$ points from the reconstructed mesh as input, where num_v is the number of vertices of the ground truth mesh. Note that for all results reported in the paper, we magnify it by 100x for better comparison.

We report the experiment results in Tab. 1 and Fig. 3. For the w/ box dataset, it’s clear that ReNeuS outperforms other baselines by a large margin both quantitatively and qualitatively. Comparing with the baselines on the w/o box dataset further proves the power of our method. Even though working on much more challenging data, our method achieves comparable performance with the baselines methods, only falling behind IDR [33] which needs a ground truth mask

	NeuS ⁺	Ours (w/o Trans Loss)	Ours
Chamfer- L_1 ↓	22.07	19.08	17.90
Failure case	3	2	0

Table 2. Results of ablation study.

for training. It’s worth noting that ReNeuS get superior performance than NeuS, even on data without a container box. We assume this to be the contribution of reflection. Every time when a ray is reflected, it’s like we have a new view of the scene which helps with reconstruction.

4.5. Visualization on real dataset.

We run ReNeuS and other baseline methods on the real dataset for quantitative comparison. For IDR, we run it with the mask of the transparent box since we don’t have an object mask. IDR failed to reconstruct anything but a convex hull of the masks. Experiment results are shown in Fig. 5 as well as the supplementary material. From the visualization results, we argue that NeReuS performs considerably well on the proposed task. Our method can even reconstruct the wings and antennae of small insects which is indeed a difficult problem in normal 3D reconstruction.

4.6. Ablation study

We conduct ablation studies to verify the effectiveness of our ReNeuS framework.

Rendering strategy. We believe that our novel hybrid rendering strategy is the key to solving the problem. To verify that, we propose a naive solution based on NeuS [31] that we only refract the rays once on the space interface. We refer to it as NeuS⁺. As the result shows in Fig. 4, the

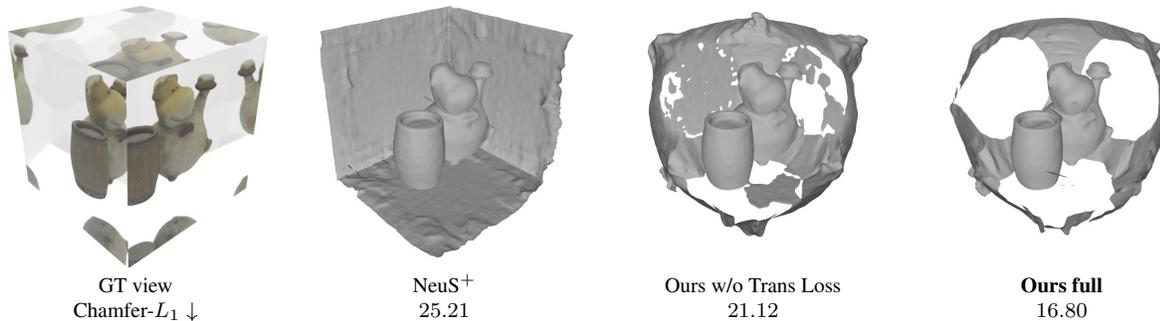


Figure 4. Visualization on ablation study.



Figure 5. Visualization of 3D reconstruction on real dataset.

result of NeuS⁺ is actually fully surrounded by another box-shaped mesh. The reason is that the NeuS⁺ method doesn't model the scene with light interaction exactly. While training with photometric loss, the network tries to render an image of the statuette on each side, thus producing these fake boundaries. As our method takes good care of the light interaction, the fake boundaries are greatly relieved. Based on this observation, we abandon the post-processing in the ablation study to better evaluate the influence of such errors. The results are shown in Tab. 2 where *Failure case* means that we fail to retrieve a mesh from the geometric MLP, after

trying 3 times. Apart from accuracy, our method performs consistently in different cases. Especially for the *dinosaur* case which is quite challenging due to its thin structure and texture-less appearance, all the two methods failed except the full model.

Transmittance loss. We also do an ablation study for our sparsity prior item as shown in Fig. 4. Those dummy facets are dramatically suppressed with the regularization.

5. Limitations and Conclusion

Limitation. The major limitation of our method is the assumption of known geometry and homogeneous background lighting. The known geometry assumption may be violated due to irregularly shaped transparent containers with rounded edges and corners and the lighting assumption may impede the data collection process. We notice from the results on real dataset that the violation of the assumptions can lead to a performance drop. We believe tackling these limitations by jointly reconstructing the object and container as well as modeling the lighting conditions can lead to promising future research directions.

Conclusion. In this paper, we have defined a new research problem of 3D reconstruction of specimens enclosed in transparent containers raised from the digitization of fragile museum collections. The image distortion caused by complex multi-path light interactions across interfaces, that separate multiple mediums, poses challenges to recovering accurate 3D geometry of the specimen. To tackle the problem, we propose a neural reconstruction method ReNeuS in which we use a novel hybrid rendering strategy that takes into account multiple light reflections and refractions across glass/air interfaces to accurately reconstruct 3D models. Experiments on both synthetic and our newly captured real dataset demonstrate the efficacy of our method. It is capable of producing high-quality 3D reconstructions when existing methods fail.

Acknowledgements. The project is supported by CSIRO Julius Career Award, CSIRO Early Research Career (CERC) Fellowship and CSIRO Future Digital Manufacturing Fund (FDMF) Postdoc Fellowship. Jinguang is supported by CSIRO Data61 PhD Scholarship.

References

- [1] Amit Agrawal, Srikumar Ramalingam, Yuichi Taguchi, and Visesh Chari. A theory of multi-layer flat refractive geometry. In *CVPR*, pages 3346–3353, 2012. 2
- [2] Matan Atzmon and Yaron Lipman. Sal: Sign agnostic learning of shapes from raw data. In *CVPR*, pages 2565–2574, 2020. 6
- [3] Max Born and Emil Wolf. *Principles of optics: electromagnetic theory of propagation, interference and diffraction of light*. Elsevier, 2013. 4
- [4] Francois Chadebecq, Francisco Vasconcelos, George Dwyer, Rene Lacher, Sebastien Ourselin, Tom Vercauteren, and Danail Stoyanov. Refractive Structure-From-Motion Through a Flat Refractive Interface. In *ICCV*, 2017. 2
- [5] Blender Online Community. *Blender - a 3D modelling and rendering package*. Blender Foundation, Stichting Blender Foundation, Amsterdam, 2018. 5
- [6] CSIRO. Australian National Fish Collection. <https://www.csiro.au/en/about/facilities-collections/collections/anfc>. 1
- [7] CSIRO. Australian National Insect Collection. <https://www.csiro.au/en/about/facilities-collections/collections/anic>. 1
- [8] S. Garrido-Jurado, R. Muñoz-Salinas, F. J. Madrid-Cuevas, and M. J. Marín-Jiménez. Automatic generation and detection of highly reliable fiducial markers under occlusion. *Pattern Recognition*, 47(6):2280–2292, 2014. 6
- [9] Georg Glaeser and Hans-Peter Schröcker. Reflections on refractions. *Journal for Geometry and Graphics*, 4(1):1–18, 2000. 2
- [10] Georg Glaeser and Hans-Peter Schröcker. Reflections on refractions. *Journal for Geometry and Graphics*, 4(1):1–18, 2000. 2
- [11] Amos Gropp, Lior Yariv, Niv Haim, Matan Atzmon, and Yaron Lipman. Implicit geometric regularization for learning shapes. In *Proceedings of the 37th International Conference on Machine Learning*, pages 3789–3799, 2020. 5
- [12] Yuan-Chen Guo, Di Kang, Linchao Bao, Yu He, and Song-Hai Zhang. NeRFReN: Neural Radiance Fields With Reflections. In *CVPR*, 2022. 2, 3
- [13] Kai Han, Kwan-Yee K. Wong, and Miaomiao Liu. Dense Reconstruction of Transparent Objects by Altering Incident Light Paths Through Refraction. *IJCV*, 126(5):460–475, 2018. 2, 3
- [14] Xiao Hu, François Lauze, Kim Steenstrup Pedersen, and Jean Mérou. Absolute and relative pose estimation in refractive multi view. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops*, pages 2569–2578, 2021. 2
- [15] Anne Jordt-Sedlazeck and Reinhard Koch. Refractive Structure-from-Motion on Underwater Images. In *ICCV*, pages 57–64, 2013. 2
- [16] Michael Kazhdan and Hugues Hoppe. Screened poisson surface reconstruction. *ACM TOG*, 32(3):1–13, 2013. 6
- [17] Zhengqin Li, Yu-Ying Yeh, and Manmohan Chandraker. Through the Looking Glass: Neural 3D Reconstruction of Transparent Shapes. In *CVPR*, pages 1262–1271, 2020. 2, 3
- [18] William E. Lorensen and Harvey E. Cline. Marching cubes: A high resolution 3D surface construction algorithm. *ACM SIGGRAPH Computer Graphics*, 21(4):163–169, 1987. 3, 7
- [19] Jiahui Lyu, Bojian Wu, Dani Lischinski, Daniel Cohen-Or, and Hui Huang. Differentiable refraction-tracing for mesh reconstruction of transparent objects. *ACM TOG*, 39(6):195:1–195:13, 2020. 2, 3
- [20] N. Max. Optical models for direct volume rendering. *IEEE TVCG*, 1(2):99–108, 1995. 4
- [21] Mateusz Michalkiewicz, Jhony K. Pontes, Dominic Jack, Mahsa Baktashmotlagh, and Anders Eriksson. Implicit Surface Representations As Layers in Neural Networks. In *ICCV*, pages 4743–4752, 2019. 2, 3
- [22] 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*, pages 405–421, 2020. 2, 3, 4, 5, 6, 7
- [23] The Natural History Museum. Digital Collections. <https://www.nhm.ac.uk/our-science/our-work/digital-collections.html>. 1
- [24] Michael Oechsle, Songyou Peng, and Andreas Geiger. Unisurf: Unifying neural implicit surfaces and radiance fields for multi-view reconstruction. In *ICCV*, pages 5589–5599, October 2021. 3
- [25] Steven Parker, William Martin, Peter-Pike J. Sloan, Peter Shirley, Brian Smits, and Charles Hansen. Interactive ray tracing. In *ACM SIGGRAPH 2005 Courses*, pages 12–es, 2005. 4
- [26] Johannes L. Schönberger and Jan-Michael Frahm. Structure-From-Motion Revisited. In *CVPR*, pages 4104–4113, 2016. 1, 6
- [27] Johannes L Schönberger, Enliang Zheng, Jan-Michael Frahm, and Marc Pollefeys. Pixelwise view selection for unstructured multi-view stereo. In *ECCV*, pages 501–518, 2016. 6
- [28] Vincent Sitzmann, Julien Martel, Alexander Bergman, David Lindell, and Gordon Wetzstein. Implicit Neural Representations with Periodic Activation Functions. In *NeurIPS*, volume 33, pages 7462–7473, 2020. 2, 6
- [29] Jiaming Sun, Xi Chen, Qianqian Wang, Zhengqi Li, Hadar Averbuch-Elor, Xiaowei Zhou, and Noah Snavely. Neural 3d reconstruction in the wild. In *ACM SIGGRAPH 2022 Conference Proceedings*, pages 1–9, 2022. 2
- [30] UNSW. Museum of Human Disease. <https://www.unsw.edu.au/medicine-health/disease-museum>. 1
- [31] 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*, volume 34, pages 27171–27183, 2021. 1, 2, 3, 4, 5, 6, 7
- [32] Bojian Wu, Yang Zhou, Yiming Qian, Minglun Cong, and Hui Huang. Full 3d reconstruction of transparent objects. *ACM TOG*, 37(4), 2018. 3
- [33] Lior Yariv, Yoni Kasten, Dror Moran, Meirav Galun, Matan Atzmon, Basri Ronen, and Yaron Lipman. Multiview neural surface reconstruction by disentangling geometry and ap-

pearance. In *NeurIPS*, volume 33, pages 2492–2502, 2020.
3, 5, 6, 7