

Dex-NeRF: Using a Neural Radiance Field to Grasp Transparent Objects

Yahav Avigal



Electrical Engineering and Computer Sciences
University of California, Berkeley

Technical Report No. UCB/EECS-2022-266

<http://www2.eecs.berkeley.edu/Pubs/TechRpts/2022/EECS-2022-266.html>

December 16, 2022

Copyright © 2022, by the author(s).
All rights reserved.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission.

Acknowledgement

The topic and contents of this work are based on Dex-NeRF: Using a Neural Radiance Field to Grasp Transparent Objects, by Jeffrey Ichnowski*, Yahav Avigal*, Justin Kerr, and Ken Goldberg, in the Conference on Robot Learning (CoRL) 2021. The research was performed at the AUTOLAB at UC Berkeley in affiliation with the Berkeley Artificial Intelligence Research (BAIR) Lab, Berkeley Deep Drive (BDD), the Real-Time Intelligent Secure Execution (RISE) Lab, and the CITRIS "People and Robots" (CPAR) Initiative. I am indebted to the many individuals who have supported and guided me throughout this experience. For giving me the opportunity to do research at the AUTOLAB, supporting my work, and providing invaluable advice, I thank my advisor at UC Berkeley Prof Ken Goldberg.

**Dex-NeRF: Using a Neural Radiance Field to Grasp Transparent
Objects**

by Yahav Avigal

Research Project

Submitted to the Department of Electrical Engineering and Computer Sciences,
University of California at Berkeley, in partial satisfaction of the requirements for the
degree of **Master of Science, Plan II**.

Approval for the Report and Comprehensive Examination:

Committee:

Ken Goldberg

Professor Ken Goldberg
Research Advisor

2022-12-15

* * * * *

Angjoo Kanazawa

Professor Angjoo Kanazawa
Second Reader

2022-12-15

Abstract

Dex-NeRF: Using a Neural Radiance Field to Grasp Transparent Objects

by

Yahav Avigal

Masters of Science in Electrical Engineering and Computer Sciences in

University of California, Berkeley

Ken Goldberg

The ability to grasp and manipulate transparent objects is a major challenge for robots. Existing depth cameras have difficulty detecting, localizing, and inferring the geometry of such objects. We propose using neural radiance fields (NeRF) to detect, localize, and infer the geometry of transparent objects with sufficient accuracy to find and grasp them securely. We leverage NeRF’s view-independent learned density, place lights to increase specular reflections, and perform a transparency-aware depth-rendering that we feed into the Dex-Net grasp planner. We show how additional lights create specular reflections that improve the quality of the depth map, and test a setup for a robot workcell equipped with an array of cameras to perform transparent object manipulation. We also create synthetic and real datasets of transparent objects in real-world settings, including singulated objects, cluttered tables, and the top rack of a dishwasher. In each setting we show that NeRF and Dex-Net are able to reliably compute robust grasps on transparent objects, achieving 90 % and 100 % grasp-success rates in physical experiments on an ABB YuMi, on objects where baseline methods fail.

Contents

Contents	i
List of Figures	ii
List of Tables	iv
1 Introduction	1
2 Related Work	4
3 Problem Statement	6
4 Method	7
5 Experiments	11
6 Limitations	19
7 Conclusion	20
Bibliography	21

List of Figures

1.1	Using NeRF to grasp transparent objects Given a scene with transparent objects (left column), we the pipeline on the right to compute grasps (middle column). The top row shows Dex-NeRF working in a simulated scene while the bottom row shows it working in a physical scene.	2
4.1	Comparison to RealSense Depth Camera. We compare the results of the proposed pipeline in a real-world setting against the depth map produced by an Intel RealSense camera. In the left image is the real-world scene, the middle shows the depth image from the RealSense, and the right shows the result of our pipeline. The color scheme in the RealSense image is provided by the RealSense SDK, while the color scheme in the right column is from Matplotlib. We observe that the RealSense depth camera is unable to recover depth from a large portion of the scene, shown in black. On the other hand, the proposed pipeline, while having a few holes, can recover depth for most of the scene.	8
4.2	Using NeRF to render depth for grasping transparent objects. Dex-NeRF uses a transparency-aware depth rendering to render depth maps that can be used for grasp planning. In contrast, Vanilla-NeRF’s depth maps are filled with holes and result in poor grasp predictions.	9
5.1	Synthetic singulated objects used in simulation experiments. Top row: image of the object in the training data. Bottom row: computed depth map and candidate grasp.	12
5.2	Grasp-success rate vs training epochs. As opposed to view-synthesis, which requires over 200k epochs, we observe high grasp success rates after 50k to 60k epochs.	13
5.3	Physical grasps objects. The objects are placed on a rack, and in the background is the base of the YuMi robot.	14

- 5.4 More lights mean more specular reflections, and result in better NeRF depth estimation of transparent surfaces. In (a) and (b), we show a scene lit by a single overhead high-intensity light. In (c) and (d) we show the same scene lit by an overhead 5x5 array of lights. The combined light wattage is equal in both scenes. Images (a) and (c) are views of the scene, and (b) and (d) are the corresponding depth images obtained from the pipeline. Two glasses on their sides are missing top surfaces (outlined in dashed red) in (b), while the effect is reduced in (d) due to the additional light sources. 17
- 5.5 **depth rendering using NeRF with different thresholds** Here we show the effect of the threshold value on the depth rendering of an isolated deer figurine. Values too low result in excess noise, while values too high cause parts of the scene to disappear. 18
- 5.6 **Depth rendering using a grid of overhead cameras.** Using increasing amounts of overhead cameras improves the quality of the depth map and its utility in grasping, however, beyond a certain number of cameras there is a diminishing return. 18

List of Tables

- 5.1 **Physical grasp success rate.** For each object, we compute a depth map using a PhoXi camera, unmodified Vanilla NeRF, and Dex-NeRF for grasping transparent objects. From the depth map, Dex-Net computes a 10 different grasps, and an ABB YuMi attempts the grasp. Successful grasps lift the object. 15

Acknowledgments

The topic and contents of this work are based on Dex-NeRF: Using a Neural Radiance Field to Grasp Transparent Objects, by Jeffrey Ichnowski*, Yahav Avigal*, Justin Kerr, and Ken Goldberg, in the Conference on Robot Learning (CoRL) 2021. The research was performed at the AUTOLAB at UC Berkeley in affiliation with the Berkeley Artificial Intelligence Research (BAIR) Lab, Berkeley Deep Drive (BDD), the Real-Time Intelligent Secure Execution (RISE) Lab, and the CITRIS "People and Robots" (CPAR) Initiative.

I am indebted to the many individuals who have supported and guided me throughout this experience. For giving me the opportunity to do research at the AUTOLAB, supporting my work, and providing invaluable advice, I thank my advisor at UC Berkeley Prof Ken Goldberg. I was introduced to him by Prof Dan Halperin, who helped me find my way to robotics during my time at Tel Aviv University. I would also like to express my deepest appreciation to my friend and co-author Dr Jeffrey Ichnowski, a brilliant roboticist and critical thinker who was always willing to tackle the hardest and most challenging problems together.

Chapter 1

Introduction

Transparent objects are common in homes, restaurants, retail packaging, labs, gift shops, hospitals, and industrial warehouses. Effectively automating robotic manipulation of transparent objects could have a broad impact, from helping in everyday tasks and performing tasks in hazardous environments. Existing depth cameras assume that surfaces of observed objects reflect light uniformly in all directions, but this assumption does not hold for transparent objects as their appearance varies significantly under different view directions and illumination conditions due to reflection and refraction properties of transparent materials. To address this, we propose and demonstrate *Dex-NeRF*, a new method to sense the geometry of transparent objects and allow for robots to interact with them.

Dex-NeRF uses a Neural Radiance Fields (NeRF) as part of a pipeline (Fig. 1.1, right) to compute and execute robot grasps on transparent objects. While NeRF was originally proposed as an alternative for explicit volumetric representations and shown to render novel views of complex scenes realistically [24], it can also reconstruct the scene geometry. In particular, due to the view-dependent nature of the NeRF model, it can learn to represent the geometry associated with transparency accurately. The only input requirement to train a NeRF model is a set of images taken from a camera with known intrinsics (e.g., focal length, distortion) and extrinsics (position and orientation in the world). While the intrinsics can be determined from calibration techniques or from the camera itself, determining the extrinsics is often a challenge [34, 35]. However, robots operating in a fixed workcell or with a camera mounted on the manipulator arm, can readily determine camera intrinsics. This makes NeRF a particularly good match for robot manipulators.

In experiments, we show qualitatively and quantitatively that *Dex-NeRF* can compute high accuracy depth images from photo-realistic synthetic and real scenes, and achieve 90% or better grasp-success rates on real objects.

In this thesis, my main contributions were the implementation of the NeRF transparency-aware depth rendering method, the integration of NeRF and *Dex-Net* to compute grasps on the depth renderings, training the NeRF model, and the benchmarks on in simulation and on a physical system. Dr Jeff Ichnowski generated synthetic data for training the NeRF model, and executed the experiments on the number of lights and the workcell setup. Justin Kerr

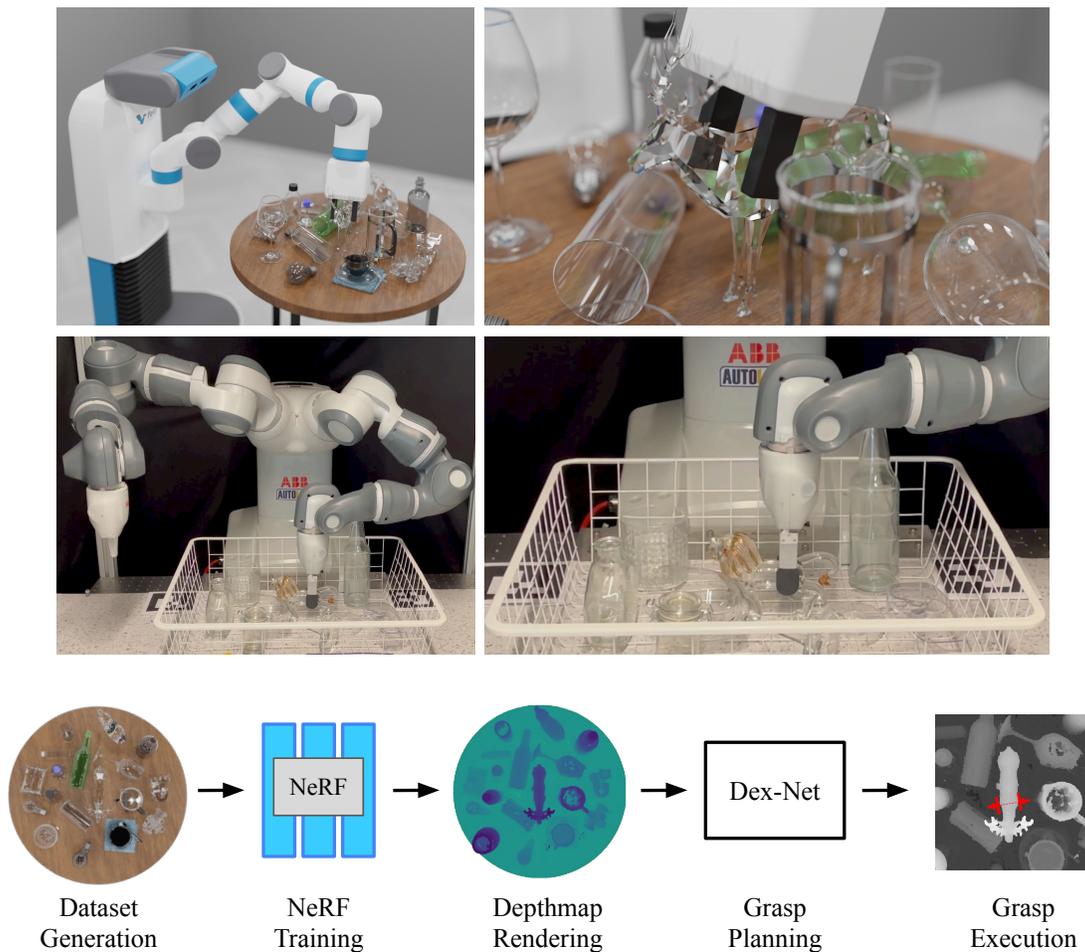


Figure 1.1: **Using NeRF to grasp transparent objects** Given a scene with transparent objects (left column), we use the pipeline on the right to compute grasps (middle column). The top row shows Dex-NeRF working in a simulated scene while the bottom row shows it working in a physical scene.

helped to set up the physical system. Myself, Dr Jeffrey Ichnowski and Prof Ken Goldberg all participated in discussions of all aspects of the project, including preliminary results and the overall research direction. Apart from this project, I have explored different approaches to increasing the efficiency and reliability of robotic manipulation in polyculture farming [5, 4], object transport [12, 13, 3], transparent object manipulation [16], and garment folding [6], which will be presented in my dissertation.

The contributions of this thesis are:

1. integration of NeRF with robot grasp planning,

2. a transparency-aware depth rendering method for NeRF,
3. experiments on synthetic and real images showing NeRF with Dex-Net generates high-quality grasps,
4. synthetic and real image datasets with transparent objects for training NeRF models.

Chapter 2

Related Work

Detecting Transparent Objects For robots to interact with transparent objects, they must first be able to detect them. The most recent approaches detecting and recognizing transparent objects are data-driven. Lai et al. [19] and Khaing et al. [17] propose using a Convolutional Neural Network (CNN) to detect transparent objects in RGB images. Recently, Xie et al. [41] developed a transformer-based pipeline [39] for transparent object segmentation. Other methods rely on deep-learning models to predict the object pose. Phillips et al. [29] trained a random forest to detect the contours of transparent objects for pose estimation and shape recovery. Xu et al. [42] proposed a two-stage method for estimating the 6-degrees-of-freedom (DOF) pose of a transparent object from a single RGBD image by replacing the noisy depth values with estimated values and training a DenseFusion-like network structure [40]. Sajjan et al. [33] extend this and incorporate a neural network trained for 3D pose estimation of transparent objects in a robotic picking pipeline. Zhou et al. [45, 44] train a grasp planner directly on raw images from a light-field camera. Zhu et al. [46] used an implicit function to complete missing depth given noisy RGBD observation of transparent objects. However, these data-driven methods rely on large annotated datasets that are hard to curate, whereas Dex-NeRF does not require any prior dataset.

Neural Radiance Fields Recently, implicit neural representations have led to significant progress in 3D object shape representation [23, 27, 9] and encoding the geometry and appearance of 3D scenes [36, 24]. Mildenhall et al. [24] presented Neural Radiance Fields (NeRF), a neural network whose input is a 3D coordinate with an associated view direction, and output is the volume density and view-dependent emitted radiance. Due to its view-dependent prediction, NeRF can represent non-Lambertian effects such as specularities and reflections, and therefore capture the geometry of transparent objects. However, NeRF is slow to train and has low data efficiency. Yu et al. [43] proposed *Plenotrees*, mapping coordinates to spherical harmonic coefficients, shifting the view-dependency from the input to the output. In addition, *Plenotrees* pre-samples the model into a sparse octree structure, achieving a significant speedup in training over NeRF. Deng et al. [10] proposed JaxNeRF, an efficient JAX implementation of NeRF reduces the training time of a NeRF model from

over a day to several hours. Deng et al. [11] add depth supervision to train NeRF 2 to 6× faster given fewer training views. Adamkiewicz et al. [1] proposed an algorithm that uses a NeRF model for robot navigation. In this work, we propose to use NeRF to recover the geometry of transparent objects for the purpose of robotic manipulation.

Robotic Grasping Traditional robot grasping methods analyze the object shape to identify successful grasp poses [18, 7, 26]. Data-driven approaches learn a prior using labeled data [15, 30] or through self-supervision over many trials in a simulated or physical environment [14, 28] and generalize to grasping novel objects with unknown geometry. These approaches rely on RGB and depth sensors to generate an accurate observation of the target object. Additionally, different methods use different inputs, such as depth maps [21, 20, 32], point clouds [25, 31, 38, 42], octrees [2], or a truncated signed distance function (TSDF) [8, 37]. In contrast, in this paper we propose a method to render a high-quality depth map from a NeRF model to then pass to Dex-Net [21] to compute a grasp. While standard depth cameras have gaps in their depth information that needs to be processed out with hole-filling techniques, the depth map rendering from NeRF is directly usable. It is possible that other grasp-planning techniques may be able to plan grasps from NeRF models.

Chapter 3

Problem Statement

We assume an environment with an array of cameras at known fixed locations or that the robot can manipulate a camera (e.g., wrist-mounted) to capture multiple images of the scene. Given the environment with rigid transparent objects, Dex-NeRF computes a frame for a robot gripper that will result in a stable grasp of a transparent object.

Chapter 4

Method

This section provides a brief background on NeRF, then describes recovering geometry of transparent objects, integrating with grasp analysis, and improving performance with additional lights.

Preliminary: Training NeRF

NeRF [24] learns a neural scene representation that maps a 5D coordinate containing a spatial location (x, y, z) and viewing direction (θ, ϕ) to the volume density σ and RGB color \mathbf{c} . Training NeRF’s multilayer perceptron (MLP) requires multiple RGB images of a static scene with their corresponding camera poses and intrinsic parameters. The expected color $C(\mathbf{r})$ of the camera ray $\mathbf{r} = \mathbf{o} + t\mathbf{d}$ between near and far scene bounds t_n and t_f is:

$$C(\mathbf{r}) = \int_{t_n}^{t_f} T(t)\sigma(\mathbf{r}(t))\mathbf{c}(\mathbf{r}(t), \mathbf{d})dt, \quad (4.1)$$

where $T(t) = \exp\left(-\int_{t_n}^t \sigma(\mathbf{r}(s))ds\right)$ is the probability that the camera ray travels from near bound t_n to point t without hitting any surface. NeRF approximates the expected color $\hat{C}(\mathbf{r})$ as:

$$\hat{C}(\mathbf{r}) = \sum_{i=1}^N T_i(1 - \exp(-\sigma_i\delta_i))\mathbf{c}_i, \quad (4.2)$$

where $T_i = \exp\left(-\sum_{j=1}^{i-1} \sigma_j\delta_j\right)$ and $\delta_i = t_{i+1} - t_i$ is the distance between consecutive samples on the ray \mathbf{r} . The training process minimizes the error between rendered and ground-truth colors.

Recovering Geometry of Transparent Objects

We observe that NeRF does not directly support transparent object effects—it casts a single ray per source image pixel without reflection, splitting, or bouncing. NeRF recovers non-Lambertian effects such as reflections from a specular surface by regressing on view direction

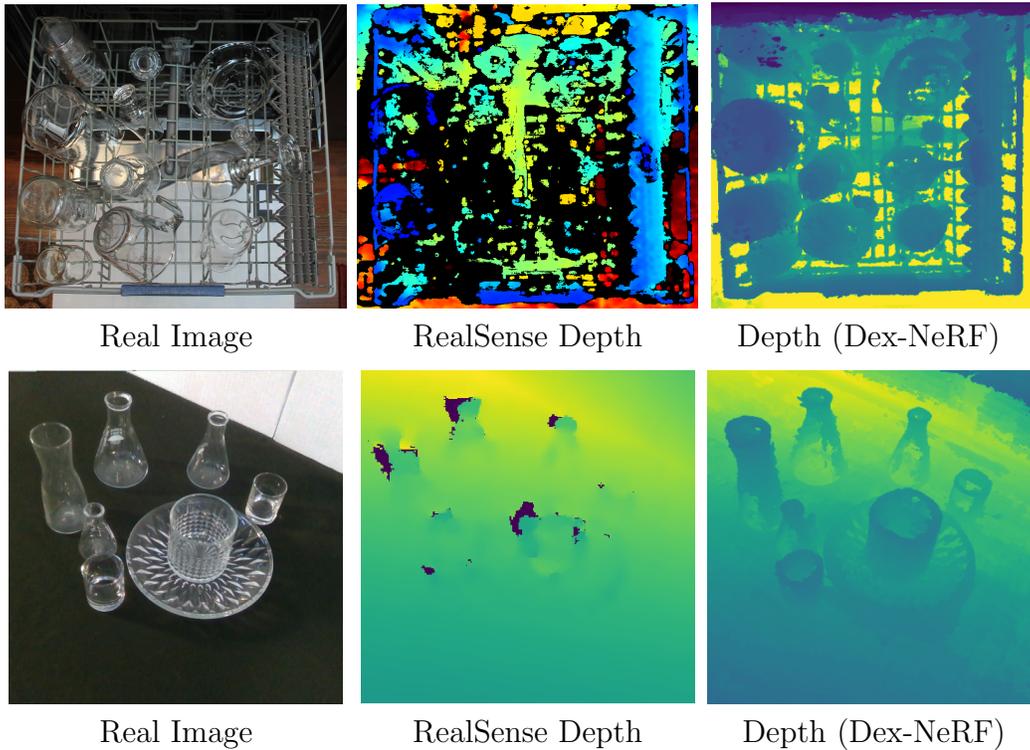


Figure 4.1: **Comparison to RealSense Depth Camera.** We compare the results of the proposed pipeline in a real-world setting against the depth map produced by an Intel RealSense camera. In the left image is the real-world scene, the middle shows the depth image from the RealSense, and the right shows the result of our pipeline. The color scheme in the RealSense image is provided by the RealSense SDK, while the color scheme in the right column is from Matplotlib. We observe that the RealSense depth camera is unable to recover depth from a large portion of the scene, shown in black. On the other hand, the proposed pipeline, while having a few holes, can recover depth for most of the scene.

and supervising with view-dependent emitted radiance. However, while RGB color \mathbf{c} is view-dependent, the volume density σ is not—meaning NeRF has to learn a non-zero σ to represent any color at that spatial location. The usual result is that the transparent object shows up as a “ghostly” or “blurry” version of the original object in rendered RGB images.

When training, a NeRF model learns a density σ of each spatial location. This density corresponds to the transparency of the point, and serves to help learn how much a spatial location contributes to the color of a ray cast through it. Although NeRF converts each σ_i to an occupancy probability $\alpha_i = 1 - \exp(-\sigma_i \delta_i)$, where δ_i is the distance between integration times along the ray, thus implicitly giving α_i an upper bound of 1, it does not place a bound on the raw σ value. Dex-NeRF uses the raw value of σ to determine if a point in space is occupied.

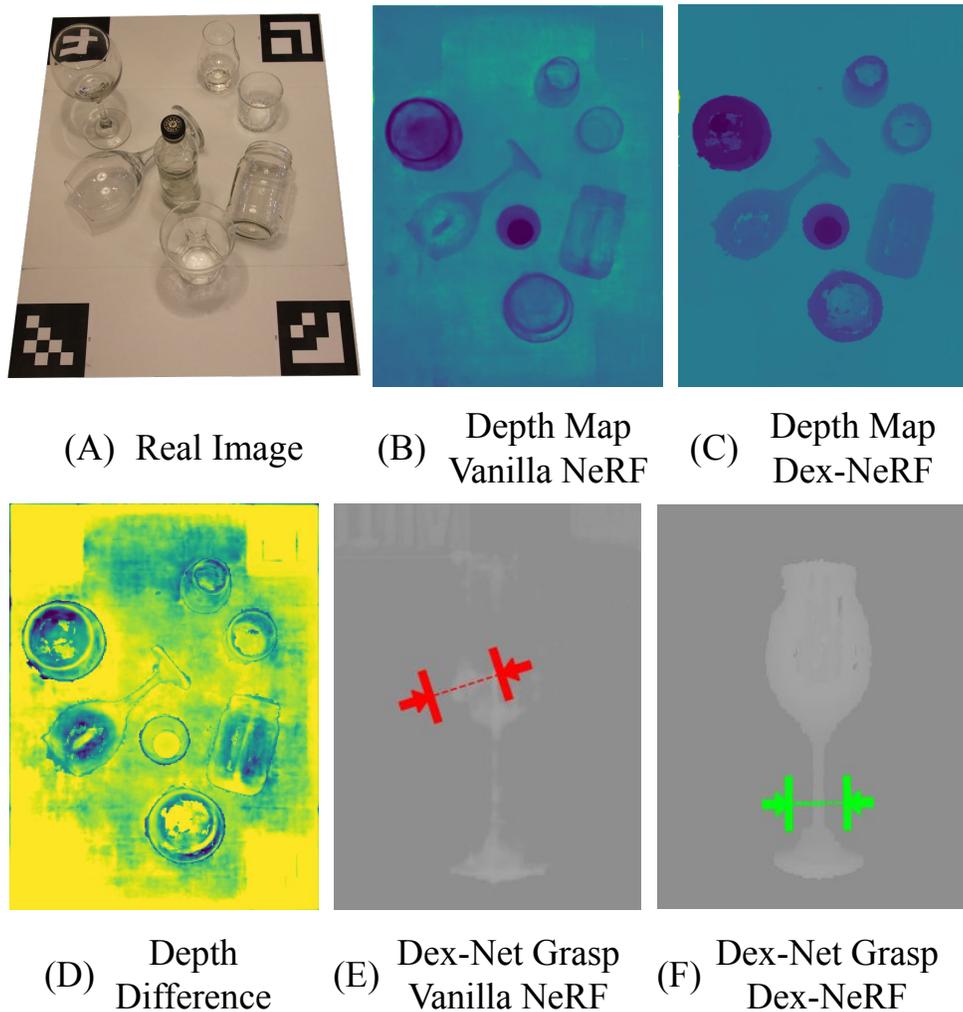


Figure 4.2: **Using NeRF to render depth for grasping transparent objects.** Dex-NeRF uses a transparency-aware depth rendering to render depth maps that can be used for grasp planning. In contrast, Vanilla-NeRF’s depth maps are filled with holes and result in poor grasp predictions.

Rendering Depth for Grasp Analysis

To compute a grasp from a trained NeRF model, we propose to render a depth image and have Dex-Net use it to plan the grasp. To generate a depth image, we consider two candidate reconstructions of depth. First, we use the same depth rendering that NeRF uses. This *Vanilla NeRF* reconstruction first converts σ_i to an occupancy probability α_i . It then applies the transformation $w_i = \alpha_i \prod_{j=1}^{i-1} (1 - \alpha_j)$. To render depth at pixel coordinate $[u, v]$, it computes the sum of sample distances from the camera weighted by the termination

probability $D[u, v] = \sum_{i=1}^N w_i \delta_i$. When applied on transparent objects, however, this results in noisy depth maps, as shown in Fig. 4.2.

Instead, we consider a second, transparency-aware method that searches for the first sample along the ray for which $\sigma_i > m$, where m is a fixed threshold. The depth is then set to the distance of that sample δ_i . We explore different values for m , and observe that low values result in a noisy depth map while high values create holes in the depth map. In our experiments we set $m = 15$ (see Fig. 5.5).

Improving Reconstruction with Light Placement

For NeRF to learn the geometry of a transparent object, it must be able to “see” it from multiple camera views. If the transparent object is not visible from any views, then it will have no effect on the loss function used in training, and thus not be learned. We thus look for a way to improve visibility of transparent objects to NeRF.

One property that transparent objects share (e.g., glass, clear plastic) is that they are glossy and thus produce specular reflections when the camera view direction is opposite to the surface normal of the incident direction of light. To the NeRF model, a specular reflection viewed from multiple angles will appear as a bright point on a solid surface—e.g., $\mathbf{c} = [1, 1, 1]^T$ and $\sigma > 0$, while from other angles it will appear as $\sigma \leq 0$. As σ is view-independent, NeRF learns a σ between fully opaque and fully transparent for such points.

By placing additional lights in the scene, we create more angles from which cameras will see specular reflections from transparent objects—this results in NeRF learning a model that fills holes in the scene. While the number and placement of lights for optimal training is dependent on both the expected object distribution and camera placement, in experiments (Sec. 5) we show that increasing from 1 light to a 5x5 array of lights improves the quality of the learned geometry.

Chapter 5

Experiments

We experiment in both simulation and on a physical ABB YuMi robot. We generate multiple datasets, where each dataset consists of images and associated camera transforms of one static scene containing one or more transparent objects.

We train NeRF models using a modified JaxNeRF [10] implementation on 4 Nvidia V100 GPUs. We use an existing pre-trained Dex-Net model for grasp planning without modification or fine-tuning. We can do this since NeRF models can be rendered to depth maps from arbitrary camera intrinsics and extrinsics, thus we match our NeRF rendering to the Dex-Net model instead of training a new one.

Datasets

As existing NeRF datasets do not include transparent objects, and existing transparent-object-grasping datasets do not include multiple camera angles, we generate new datasets using 3 different methods: synthetic, Cannon EOS 60D camera with a Tamron Di II lens with a locked focal length, and an Intel RealSense.

For synthetic datasets, we use Blender 2.92’s physically-based Cycles renderer with path tracing set to 10240 samples per pixel, and max light path bounces set to 1024. We chose these settings by increasing them until renderings were indistinguishable from the previous setting—finding that lower settings lead to dark regions and smaller specular reflections. For glass materials, we set the index of refraction to 1.45 to match physical glass. We include 8 synthetic datasets of transparent objects: 2 scenes with clutter: light array and single light; 4 singulated objects from Dex-Net: Pipe Connector, Pawn, Turbine Housing, Mount; and 2 household objects: Wineglass upright and Wineglass on side. As these computationally demanding to render due to the high quality settings, we distribute these as part of the contribution.

For the Cannon EOS and RealSense real-world datasets, we place ArUco markers in the scene to aid in camera pose recovery and take photos around the objects using a fixed ISO, f-stop, and focal length. We use bundle adjustment from COLMAP [34, 35] to refine the camera poses and intrinsics to high accuracy. We include 8 physical datasets of transparent

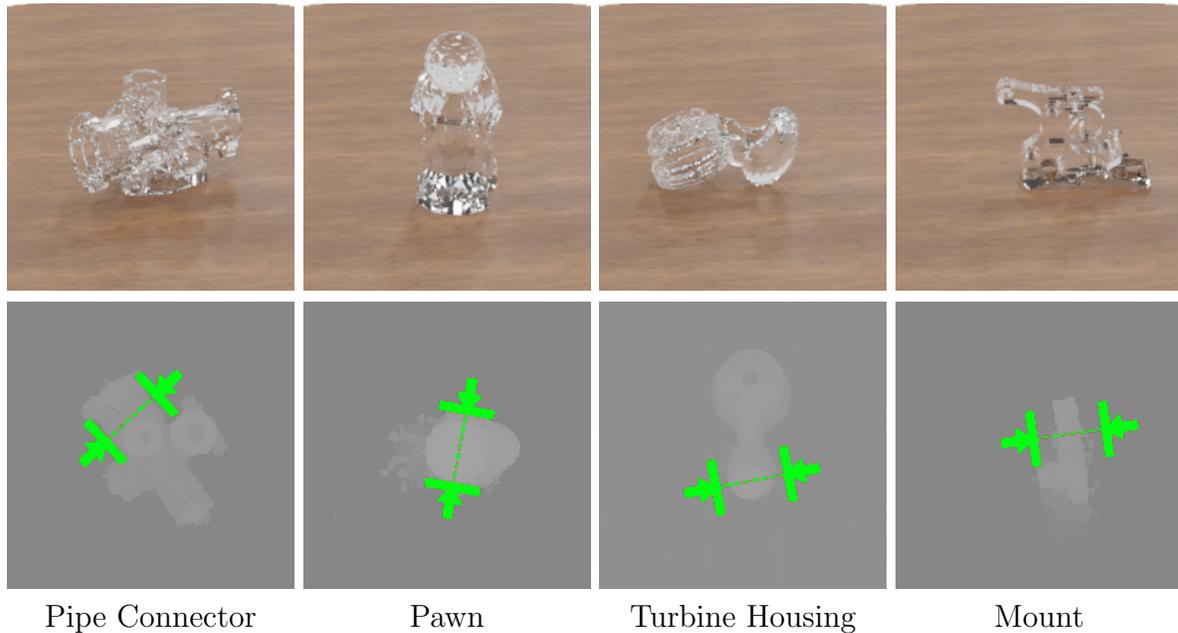


Figure 5.1: **Synthetic singulated objects** used in simulation experiments. **Top row:** image of the object in the training data. **Bottom row:** computed depth map and candidate grasp.

objects with a variety of camera poses: table with clutter, Dishwasher, Tape Dispenser, Wineglass on side, Flask, Safety Glasses, Bottle upright, Lion Figurine in clutter. The main difficulty in generating these datasets is calibration and computing high-precision camera poses.

The datasets (at <https://sites.google.com/view/dex-nerf>) differ from prior work in their focus on scenes with transparent objects in a graspable setting, with over 70 camera poses each.

Synthetic Grasping Experiments

We test the ability of Dex-NeRF to generate grasps on the synthetic singulated transparent Dex-Net object datasets.

For each dataset, we evaluate the grasp in simulation using a wrench resistance metric measuring the ability of the grasp to resist gravity [22]. Fig. 5.1 shows images of the synthetic objects, Dex-NeRF-generated depth map, and an example sampled grasp for each. To measure the effect of training time on grasp-success rate, we simulate and record grasps over the course of training. In Fig. ??, we observe that grasp-success rate improves with training time, but plateaus between 80 % and 98 % success rate at around 50k to 60k iterations. This suggests that there may be a practical fixed iteration limit to obtain high grasp success rates.

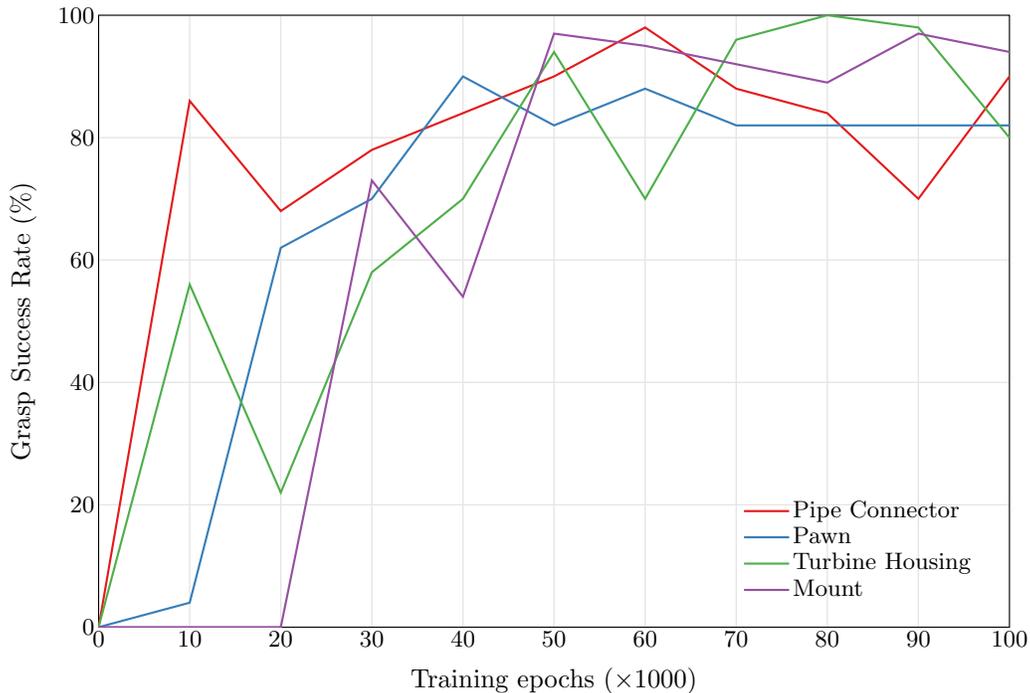


Figure 5.2: **Grasp-success rate vs training epochs.** As opposed to view-synthesis, which requires over 200k epochs, we observe high grasp success rates after 50k to 60k epochs.

We test Dex-NeRF on a scene of a tabletop cluttered with transparent objects. In this experiment, the goal is to grasp a transparent object placed in a stable pose in close proximity to other transparent objects. The challenge is twofold: the depth rendering quality should be sufficient for both grasp planning and collision avoidance.

Fig. 1.1 shows the robot and scene in the upper left, and the overhead image, depth, and computed grasp inline in the pipeline, and the final computed grasp with simulated execution is in the upper middle image. The final grasp contact point was accurate to a 2mm tolerance, suggesting that Dex-NeRF with sufficient images taken from precisely-known camera locations may be practical in highly cluttered environments.

Physical Grasping Experiments

To test the Dex-NeRF in a physical setup, we place transparent singulated objects in front of an ABB YuMi robot, and have the robot perform the computed grasps. We compare to 2 baselines: (1) *PhoXi*, in which a PhoXi camera provides the depth map; and (2) *Vanilla NeRF*, in which we use the original depth rendering from NeRF. The PhoXi camera is normally able to generate high-precision depth maps for non-transparent objects. All methods use the same pre-trained Dex-Net model, and both Vanilla NeRF and Dex-NeRF

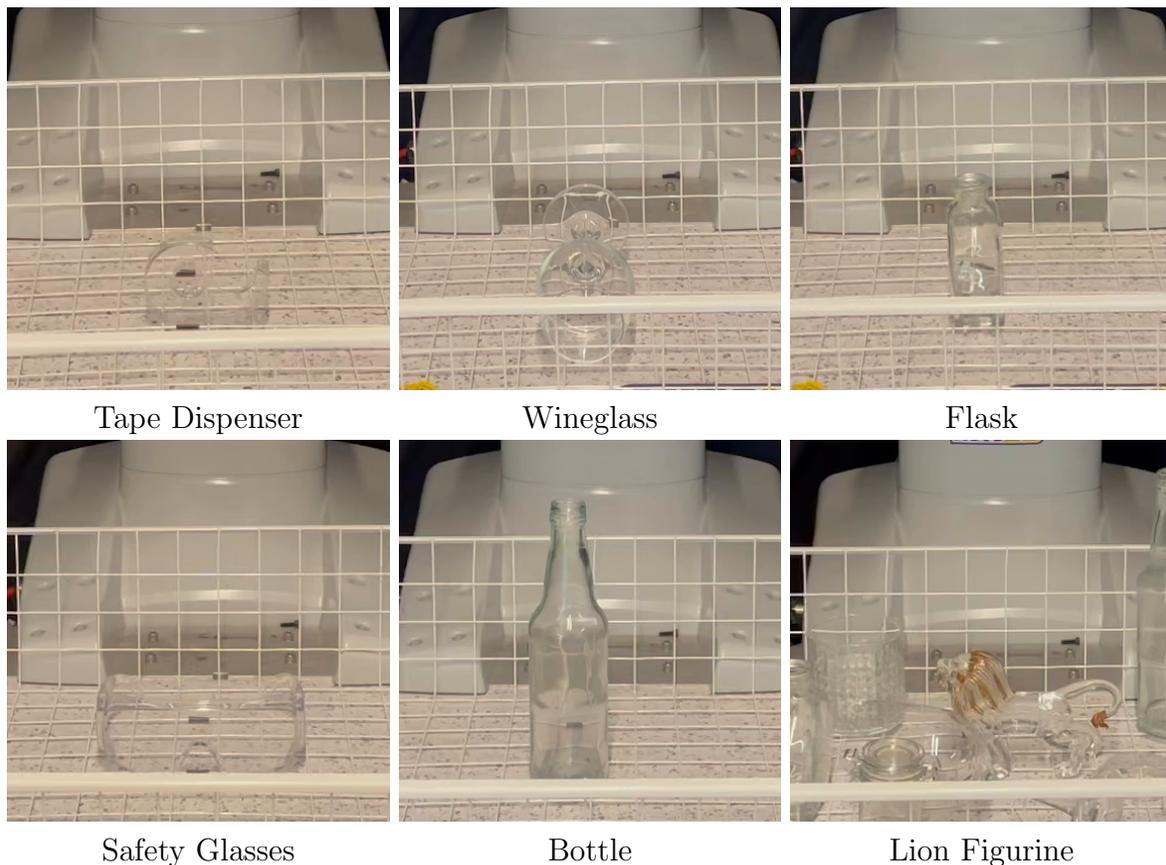


Figure 5.3: **Physical grasps objects.** The objects are placed on a rack, and in the background is the base of the YuMi robot.

use the same NeRF model—the only difference is the depth rendering. We test with 6 objects (Fig. 5.3), and compute and execute 10 different grasps for each and record the success rate. A grasp is successful if the robot lifts the object. In Table 5.1, we see that Dex-NeRF gets 90% and 100% success rates for all objects, while the baselines get few successful grasps. The PhoXi camera is unable to recover any meaningful geometry which causes Dex-Net predictions to fail. The Vanilla NeRF depth maps often have unpredictable protrusions that result in Dex-Net generating unreliable grasps.

Comparison to RealSense Depth

We qualitatively compare the rendered depth map of the proposed pipeline against a readily-available depth sensor on scenes with transparent objects in real-world settings (Fig. 4.1). We select the Intel RealSense as it is common to robotics applications, readily available, and high-performance. The RealSense, like most stereo depth cameras, struggles with transparent

Object	PhoXi	Vanilla NeRF	Dex-NeRF
Tape Dispenser	0/10	0/10	10/10
Wineglass	0/10	0/10	9/10
Flask	0/10	1/10	9/10
Safety Glasses	0/10	0/10	10/10
Bottle	0/10	10/10	10/10
Lion Figurine	0/10	3/10	10/10

Table 5.1: **Physical grasp success rate.** For each object, we compute a depth map using a PhoXi camera, unmodified Vanilla NeRF, and Dex-NeRF for grasping transparent objects. From the depth map, Dex-Net computes a 10 different grasps, and an ABB YuMi attempts the grasp. Successful grasps lift the object.

objects as they are unable to compute a stereo disparity between pixels from different cameras when the pixels are specular reflections or the color of the object behind the transparent object. The RealSense optionally projects a structured light pattern on the scene to aid in computing depth from textureless surfaces; however, in experiments, we observed no qualitative difference with and without the light pattern emitter enabled. We use a Canon EOS for NeRF, and use a RealSense for a depth image. In this experiment, we observe that the RealSense cannot compute the depth of most transparent objects and often produces regions of unknown depth (shown in black) where transparent objects are. On the other hand, the proposed pipeline produces high-quality depth maps with only a few noisy areas.

One vs Many Lights

We experiment with different light setups to test the effect of specular reflections on the ability of NeRF to recover the geometry of transparent objects. We create two scenes (Fig. 5.4), one with a single bright light source directly above the work surface, and another with an array of 5x5 (25) lights above the work surface. We set the total wattage of the lights in each scene to be the same. Since most lights in the multiple light scene are further away from the work surface than the single light source, the scene appears darker, though more evenly illuminated. The effect of the specular reflections is prominent on the lightbulb in the lower part of the image. In the single light source, there is a single specular reflection, while in the multiple light scene, the reflection of the array of lights is visible.

With the same camera setup for both scenes, we train NeRF models with the same number of iterations. We show the depth rendering in Fig. 5.4 and circle a glass and a wineglass on their side. In the single-light source image, the closer surfaces of the glasses are missing, while in the multiple-light source depth image, the glasses are nearly fully recovered. This suggests that additional lights in the scene can help NeRF recover the geometry of transparent objects better.

Workcell Setup

We experiment with a potential setup for a robot workcell in which a grid of overhead cameras captures views of the cluttered scene so that a robot manipulator arm can then perform tasks with transparent objects in the workcell. We propose that a grid of overhead cameras would be practical to setup and would not obstruct manipulator tasks nor operator interventions. The objective is to determine how many overhead cameras would be needed to recover a depth map of sufficient accuracy to perform manipulation tasks.

We place a 2 m by 2 m grid of cameras 1 m above the work surface, and have them all point at the center of the work surface. Each camera has the same intrinsics, and are evenly spaced along the grid. We experiment with grids having 4, 9, 16, 25, 36, and 49 cameras. The environment has the same 5x5 grid of lights as before. For each camera grid, we train JaxNeRF for 50k iterations and compare performance.

After training, we observe increasing peak signal to noise ratios (PSNR) and structural similarity (SSIM) scores with increasing number of cameras. The 2x2 grid of cameras produces a high train-to-test ratio for PSNR, likely indicating overfitting to training data, and results in a depth map without apparent geometry. This ratio decreases with additional cameras. The minimum number of cameras for this proposed setup appears to be around 9 (3x3) as its depth map is usable for grasp planning, while the 5x5 grid shows better PSNR and SSIM and ratio between train and test PSNR, and the 7x7 grid is the best. See Fig. 5.6 for a visual comparison. Additionally, we trained 9x9, 11x11, and 13x13 grids, observing no statistically significant improvement beyond the 7x7 grid.

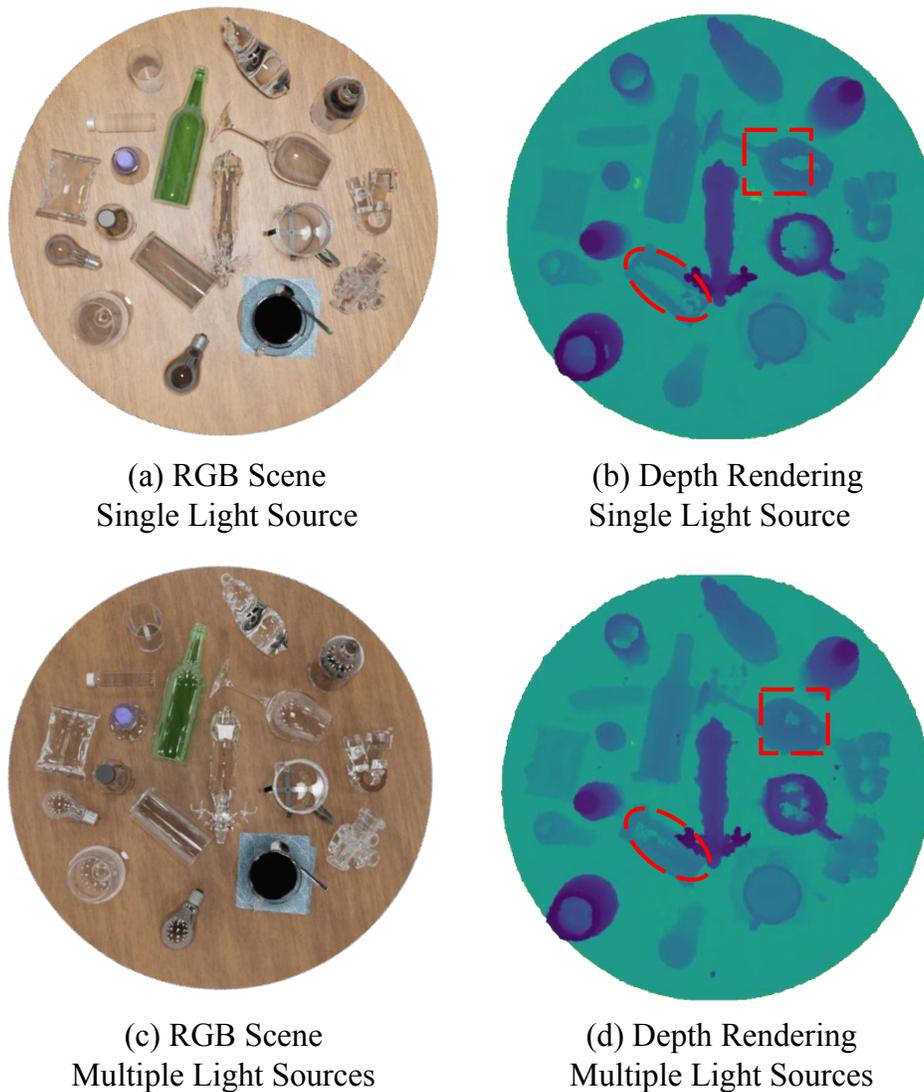


Figure 5.4: More lights mean more specular reflections, and result in better NeRF depth estimation of transparent surfaces. In (a) and (b), we show a scene lit by a single overhead high-intensity light. In (c) and (d) we show the same scene lit by an overhead 5x5 array of lights. The combined light wattage is equal in both scenes. Images (a) and (c) are views of the scene, and (b) and (d) are the corresponding depth images obtained from the pipeline. Two glasses on their sides are missing top surfaces (outlined in dashed red) in (b), while the effect is reduced in (d) due to the additional light sources.

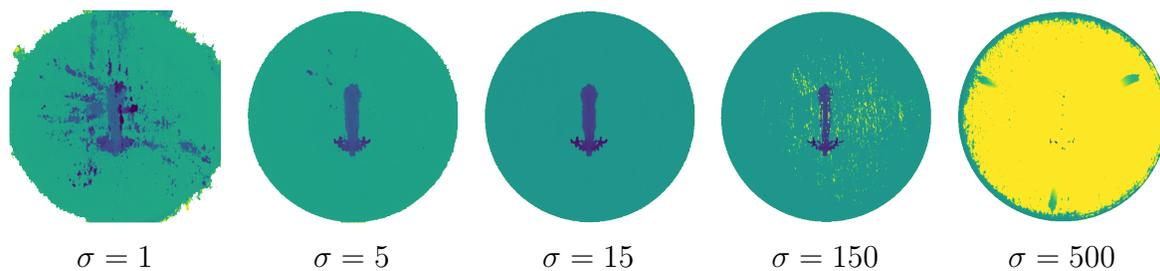


Figure 5.5: **depth rendering using NeRF with different thresholds** Here we show the effect of the threshold value on the depth rendering of an isolated deer figurine. Values too low result in excess noise, while values too high cause parts of the scene to disappear.

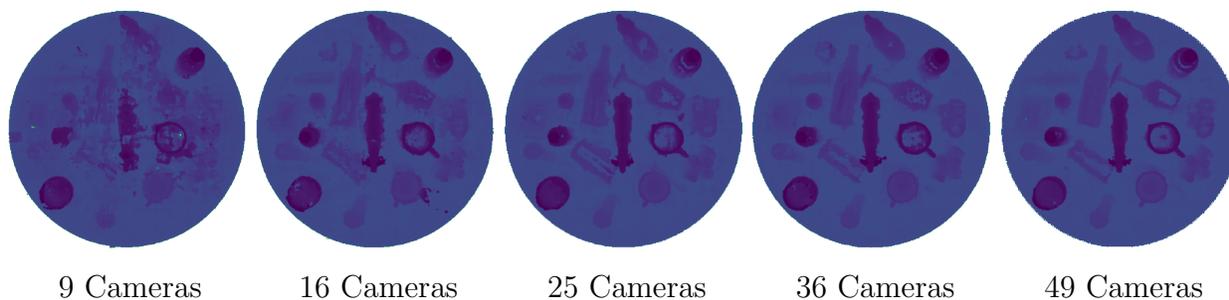


Figure 5.6: **Depth rendering using a grid of overhead cameras.** Using increasing amounts of overhead cameras improves the quality of the depth map and its utility in grasping, however, beyond a certain number of cameras there is a diminishing return.

Chapter 6

Limitations

The main drawback of Dex-NeRF is the training time required to obtain a NeRF model. NeRF Training time for a single scene is on the order of hours, and every change in the scene requires training a new NeRF model. If training stops early, the resulting model generates noisy depth maps with *floaters*, spurious regions of density floating in space.

Furthermore, in order to compute a high quality depth reconstruction with NeRF, we collect images of the scene from a large variety of angles on a hemisphere around the workspace. Therefore the location of the objects on the workspace is limited to a region in the origin of the hemisphere, while the size of the workspace is limited by the robot reachability.

In addition, even given a perfect depth reconstruction, the grasping success rate is as good as Dex-Net, since we used pre-trained weights. More over, these weights were pre-trained in simulation on synthetic depth images, that appear different from depth images rendered from NeRF. This distribution shift may affect the grasping performance, however investigating these differences thoroughly is left for future work.

Chapter 7

Conclusion

In this work, we showed that NeRF can recover the geometry of transparent objects with sufficient accuracy to compute grasps for robot manipulation tasks. NeRF learns the density of all points in space, which corresponds to how much the view-dependent color of each point contributes to rays passing through it. With the key observation that specular reflections on transparent objects cause NeRF to learn a non-zero density, we have Dex-NeRF recover the geometry of transparent objects through a combination of additional lights to create specular reflections and thresholding to find transparent points that are visible from some view directions. With the geometry recovered, we pass it to a grasp planner, and show that the recovered geometry is sufficient to compute a grasp, and accurate enough to achieve 90 % and 100 % grasp success rates in physical experiments on an ABB YuMi robot. We created synthetic and real datasets for experiments in transparent geometry recovery, but we believe these datasets may be of interest to researchers interested in extending NeRF capabilities in other ways and thus contribute them as well. Finally, to test if NeRF could be used in a robot workcell, we experimented with grids of cameras facing a worksurface and their ability to recover geometry in potential setup, and showed the increased capabilities and point of diminishing return for additional cameras.

Bibliography

- [1] Michal Adamkiewicz et al. “Vision-Only Robot Navigation in a Neural Radiance World”. In: *arXiv preprint arXiv:2110.00168* (2021).
- [2] Yahav Avigal et al. “AVPLUG: Approach Vector PLanning for Unicontact Grasping amid Clutter”. In: *2021 IEEE 17th International Conference on Automation Science and Engineering (CASE)*. IEEE. 2021, pp. 1140–1147.
- [3] Yahav Avigal et al. “GOMP-ST: Grasp Optimized Motion Planning for Suction Transport”. In: *arXiv preprint arXiv:2203.08359* (2022).
- [4] Yahav Avigal et al. “Learning Seed Placements and Automation Policies for Polyculture Farming with Companion Plants”. In: *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2021, pp. 902–908.
- [5] Yahav Avigal et al. “Simulating polyculture farming to tune automation policies for plant diversity and precision irrigation”. In: *2020 IEEE 16th International Conference on Automation Science and Engineering (CASE)*. IEEE. 2020, pp. 238–245.
- [6] Yahav Avigal et al. “Speedfolding: Learning efficient bimanual folding of garments”. In: *arXiv preprint arXiv:2208.10552* (2022).
- [7] Antonio Bicchi and Vijay Kumar. “Robotic grasping and contact: A review”. In: *Proceedings 2000 ICRA. Millennium Conference. IEEE International Conference on Robotics and Automation. Symposia Proceedings (Cat. No. 00CH37065)*. Vol. 1. IEEE. 2000, pp. 348–353.
- [8] Michel Breyer et al. “Volumetric grasping network: Real-time 6 DOF grasp detection in clutter”. In: *arXiv preprint arXiv:2101.01132* (2021).
- [9] Zhiqin Chen and Hao Zhang. “Learning implicit fields for generative shape modeling”. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2019, pp. 5939–5948.
- [10] Boyang Deng, Jonathan T. Barron, and Pratul P. Srinivasan. *JaxNeRF: an efficient JAX implementation of NeRF*. 2020. URL: <http://github.com/google-research/google-research/tree/master/jaxnerf>.
- [11] Kangle Deng et al. “Depth-supervised NeRF: Fewer Views and Faster Training for Free”. In: *arXiv e-prints* (2021), arXiv–2107.

- [12] Jeffrey Ichnowski et al. “Deep learning can accelerate grasp-optimized motion planning”. In: *Science Robotics* 5.48 (2020), eabd7710.
- [13] Jeffrey Ichnowski et al. “GOMP-FIT: Grasp-optimized motion planning for fast inertial transport”. In: *2022 International Conference on Robotics and Automation (ICRA)*. IEEE. 2022, pp. 5255–5261.
- [14] Eric Jang et al. “Grasp2vec: Learning object representations from self-supervised grasping”. In: *arXiv preprint arXiv:1811.06964* (2018).
- [15] Daniel Kappler, Jeannette Bohg, and Stefan Schaal. “Leveraging big data for grasp planning”. In: *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2015, pp. 4304–4311.
- [16] Justin Kerr et al. “Evo-NeRF: Evolving NeRF for Sequential Robot Grasping of Transparent Objects”. In: *6th Annual Conference on Robot Learning*.
- [17] May Phyo Khaing and Mukunoki Masayuki. “Transparent object detection using convolutional neural network”. In: *International Conference on Big Data Analysis and Deep Learning Applications*. Springer. 2018, pp. 86–93.
- [18] Kilian Kleeberger et al. “A survey on learning-based robotic grasping”. In: *Current Robotics Reports* (2020), pp. 1–11.
- [19] Po-Jen Lai and Chiou-Shann Fuh. “Transparent object detection using regions with convolutional neural network”. In: *IPPR Conference on Computer Vision, Graphics, and Image Processing*. Vol. 2. 2015.
- [20] Ian Lenz, Honglak Lee, and Ashutosh Saxena. “Deep learning for detecting robotic grasps”. In: *The International Journal of Robotics Research* 34.4-5 (2015), pp. 705–724.
- [21] Jeffrey Mahler et al. “Dex-Net 2.0: Deep learning to plan robust grasps with synthetic point clouds and analytic grasp metrics”. In: *arXiv preprint arXiv:1703.09312* (2017).
- [22] Jeffrey Mahler et al. “Dex-Net 3.0: Computing robust vacuum suction grasp targets in point clouds using a new analytic model and deep learning”. In: *2018 IEEE International Conference on robotics and automation (ICRA)*. IEEE. 2018, pp. 5620–5627.
- [23] Lars Mescheder et al. “Occupancy networks: Learning 3d reconstruction in function space”. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2019, pp. 4460–4470.
- [24] Ben Mildenhall et al. “NeRF: Representing scenes as neural radiance fields for view synthesis”. In: *European Conference on Computer Vision*. Springer. 2020, pp. 405–421.
- [25] Arsalan Mousavian, Clemens Eppner, and Dieter Fox. “6-DOF GraspNet: Variational grasp generation for object manipulation”. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 2019, pp. 2901–2910.

- [26] Richard M Murray, Zexiang Li, and S Shankar Sastry. *A mathematical introduction to robotic manipulation*. CRC press, 2017.
- [27] Jeong Joon Park et al. “DeepSDF: Learning continuous signed distance functions for shape representation”. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2019, pp. 165–174.
- [28] Gang Peng et al. “A self-supervised learning-based 6-DOF grasp planning method for manipulator”. In: *arXiv preprint arXiv:2102.00205* (2021).
- [29] Cody J Phillips, Matthieu Lecce, and Kostas Daniilidis. “Seeing Glassware: from Edge Detection to Pose Estimation and Shape Recovery.” In: *Robotics: Science and Systems*. Vol. 3. 2016.
- [30] Domenico Prattichizzo et al. “Springer handbook of robotics”. In: *Grasping; Springer: Berlin/Heidelberg, Germany* (2008), pp. 671–700.
- [31] Yuzhe Qin et al. “S4G: Amodal single-view single-shot SE(3) grasp detection in cluttered scenes”. In: *Conference on robot learning*. PMLR. 2020, pp. 53–65.
- [32] Joseph Redmon and Anelia Angelova. “Real-time grasp detection using convolutional neural networks”. In: *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2015, pp. 1316–1322.
- [33] Shreeyak Sajjan et al. “Clear Grasp: 3D Shape Estimation of Transparent Objects for Manipulation”. In: *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2020, pp. 3634–3642.
- [34] Johannes Lutz Schönberger and Jan-Michael Frahm. “Structure-from-Motion Revisited”. In: *Conference on Computer Vision and Pattern Recognition (CVPR)*. 2016.
- [35] Johannes Lutz Schönberger et al. “Pixelwise View Selection for Unstructured Multi-View Stereo”. In: *European Conference on Computer Vision (ECCV)*. 2016.
- [36] Vincent Sitzmann, Michael Zollhöfer, and Gordon Wetzstein. “Scene representation networks: Continuous 3d-structure-aware neural scene representations”. In: *arXiv preprint arXiv:1906.01618* (2019).
- [37] Shuran Song et al. “Grasping in the wild: Learning 6 DOF closed-loop grasping from low-cost demonstrations”. In: *IEEE Robotics and Automation Letters* 5.3 (2020), pp. 4978–4985.
- [38] Martin Sundermeyer et al. “Contact-GraspNet: Efficient 6-DoF Grasp Generation in Cluttered Scenes”. In: *arXiv preprint arXiv:2103.14127* (2021).
- [39] Ashish Vaswani et al. “Attention is all you need”. In: *arXiv preprint arXiv:1706.03762* (2017).
- [40] Chen Wang et al. “Densefusion: 6d object pose estimation by iterative dense fusion”. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2019, pp. 3343–3352.

- [41] Enze Xie et al. “Segmenting transparent object in the wild with transformer”. In: *arXiv preprint arXiv:2101.08461* (2021).
- [42] Chi Xu et al. “6DoF Pose Estimation of Transparent Object from a Single RGB-D Image”. In: *Sensors* 20.23 (2020), p. 6790.
- [43] Alex Yu et al. “PlenOctrees for real-time rendering of neural radiance fields”. In: *arXiv preprint arXiv:2103.14024* (2021).
- [44] Zheming Zhou, Xiaotong Chen, and Odest Chadwicke Jenkins. “LIT: Light-field Inference of Transparency for Refractive Object Localization”. In: *IEEE Robotics and Automation Letters* 5.3 (2020), pp. 4548–4555.
- [45] Zheming Zhou et al. “Glassloc: Plenoptic grasp pose detection in transparent clutter”. In: *arXiv preprint arXiv:1909.04269* (2019).
- [46] Luyang Zhu et al. “RGB-D Local Implicit Function for Depth Completion of Transparent Objects”. In: *arXiv preprint arXiv:2104.00622* (2021).