

A Real2Sim2Real Method for Robust Object Grasping with Neural Surface Reconstruction

Luobin Wang, Runlin Guo, Quan Vuong, Yuzhe Qin, Hao Su, Henrik Christensen

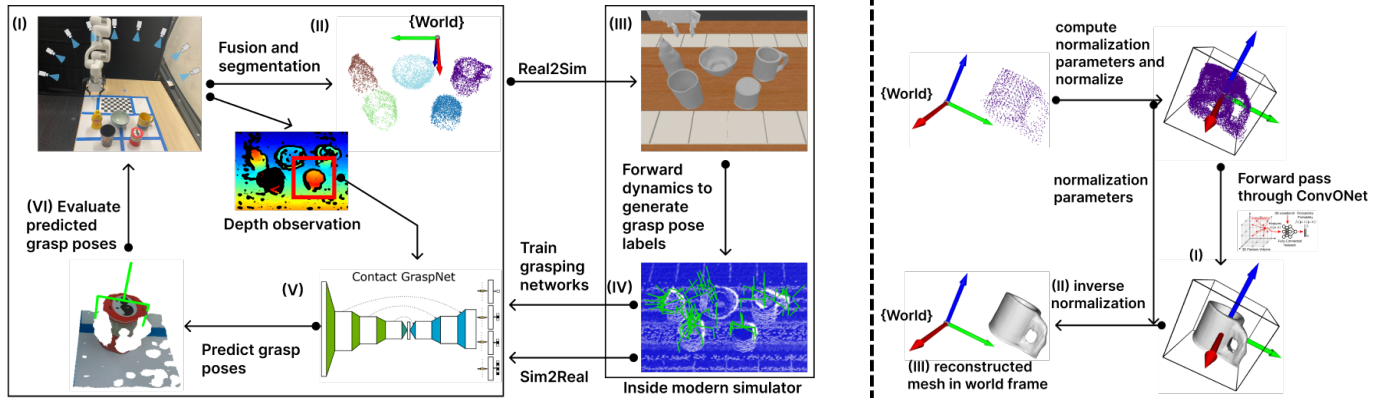


Fig. 1: **Left (our pipeline)**: given a real scene (I), we fuse and segment the camera observations to obtain object level point clouds (II), which we use to construct a digital replica of the real scene (III). The replica is used to generate grasp labels (IV) to obtain trained grasping networks (V). The grasp poses predicted by the trained networks are evaluated in the real scene (VI). **Right (the Real2Sim step)** illustrates how we can automatically place the reconstructed meshes in the digital replica without having to explicitly perform pose estimation. Given an object-level point cloud, we use a trained ConvONet to reconstruct the mesh (I), and then apply the inverse normalization operation (II) to obtain the mesh represented in the world frame (III).

Abstract—We explore an emerging technique, geometric Real2Sim2Real, in the context of object manipulation. We hypothesize that recent 3D modeling methods provides a path towards building digital replicas of real-world scenes that afford physical simulation and support robust manipulation algorithm learning. Since 6 DOF grasping is one the most important primitives for all manipulation tasks, we study whether geometric Real2Sim2Real can help us train a robust grasping network with high sample efficiency. We propose to reconstruct high-quality meshes from real-world point clouds using state-of-the-art neural surface reconstruction method (the Real2Sim step). Because most simulators take meshes for fast simulation, the reconstructed meshes enable grasp pose labels generation without human efforts. The generated labels can train grasp network that performs robustly in real evaluation scenes (the Sim2Real step). In synthetic and real experiments, we show that the Real2Sim2Real pipeline performs better than baseline grasp networks trained with a $10^4 \times$ larger dataset by mimicking geometric shapes of target objects in simulation. We also show that our method has better sample efficiency than training the grasping network with a retrieval-based scene reconstruction method. The benefit of the Real2Sim2Real pipeline comes from 1) decoupling scene modeling and grasp sampling into sub-problems, and 2) both sub-problems can be solved with sufficiently high quality using recent 3D learning algorithms and mesh-based physical simulation techniques. Video presentation available at [this link](#).

I. INTRODUCTION

Learning robotic manipulation skills in simulation and executing the skills in the real world, often termed *Sim2Real*, has fueled many recent advances in robot manipulation [1]–[5]. The paradigm is effective because many recent learning approaches require high volume of interaction samples (e.g., reinforcement learning [6], [7] and simulation-based grasp pose auto-labeling [8]), and obtaining these samples in simulators is much cheaper than in the real world. The *Sim2Real* approach usually requires practitioners to build a digital replica of the real physical scene, where the robots perform the task of interest. To construct the digital replica in simulation, the practitioners have to manually curate the object meshes, calibrate their dynamics parameters, and place them at realistic poses. Even though there are existing approaches that lower the realism requirement of the digital replica, such as domain randomization and domain adaptation [4], [9]–[17], these approaches still require the manual creation of 3D assets before they can be applied. Manual model creation, physical property calibration and scene construction require domain expertise and can be prohibitively costly to scale to multiple large-scale scenes with many objects [18]–[21].

Recognizing scene creation and calibration as a Sim2Real bottleneck, recent researches have attempted to automate this process and dub the problem *Real2Sim2Real* [22]–[24]. In fact, these recent researches tackle *dynamics Real2Sim2Real*, since they focus on estimating the simulation dynamic pa-




Fig. 2: Given objects placed on a tabletop surface, the robot should pick up the objects and move them above the tabletop surface without knocking the neighboring objects over.

rameters. We instead focus on the *geometric Real2Sim2Real* challenge – automatically constructing object geometry from camera observations, placing them at physically-plausible poses that allow for forward simulation, and demonstrating that the reconstructed scenes can train performant neural networks. In manipulation research, constructing object meshes and placing them into simulated scenes can be a highly manual process [18], [25]. Very recently, we have witnessed breakthrough in learning-based 3D modeling. State-of-the-art neural surface reconstruction methods can convert input from 3D sensors to meshes with geometric details and have demonstrated strong cross-scene generalizability [26]–[29].

Therefore, it is time to ask the hypothesis whether the quality of the generated 3D meshes from state-of-the-art neural reconstruction algorithms can be used for physical simulation and learning of manipulation algorithms. In fact, autonomous driving researches which require coarser environment geometry have already benefited tremendously from efforts to automatically reconstruct digital clones [30]–[34]. However, manipulation research requires finer details for simulation, including precise geometric meshes and accurate physical properties. Using 6-DoF object clutter grasping task with point cloud input as the case study, we demonstrate experimental evidences to support our hypothesis that modern neural reconstruction algorithms can provide sufficient details of objects to support simulation and learning. We choose grasping as the case study application because grasping is an essential primitives for all manipulation tasks and enabling grasping with our compact framework provides strong promises for more complex manipulation tasks.

Our pipeline uses a performant and recent neural surface reconstruction network that learns to convert point clouds to detailed meshes. To create a digital replica of the scene, we propose a simple yet effective method to place the object meshes into the scene automatically without having to explicitly perform pose estimation. We demonstrate our results in the tabletop grasping-in-clutter task and across two recent state-of-the-art grasp pose prediction networks [35], [36]. The grasping networks have different input modality and network architecture, providing additional evidence for the generality of our results.

We compare our method with a recent scene reconstruction baseline [37] which uses object pose estimation followed by retrieval-based method for building the digital scene replica. The results show that we can achieve higher sample

Algorithm 1 Step-by-step description of our reconstruction framework to create a digital replica of the test scenes from camera observations

Input: depth and segmentation maps, camera poses

Output: A digital replica of the test scene

- 1: Convert the depth maps to point clouds and fuse the point clouds
 - 2: Extract object-level point clouds using the semantic segmentation maps
 - 3: **for** Each object-level point cloud **do**
 - 4: Perform point cloud outlier removal
 - 5: Use [Algorithm 2](#) to reconstruct an object mesh
 - 6: Place the object mesh into the simulation scene
 - 7: **end for**
-

efficiency if the shapes of digital replicas are closer to the real world scenario.

To put the performance of our Real2Sim2Real pipeline into perspective, we also compare the pre-trained variant and the train-with-reconstruction variant of the best-performing grasping network we used. Although the pre-trained model used a $10^4 \times$ larger training dataset, we observe similar grasping performance in simulation. We can also achieve the same robust grasping performance when adapting our compact framework to a new gripper in the real world without having to relabel the entire large dataset. In summary, our method decouples 3D modeling and grasp pose sampling, and both sub-problems can be solved with quality and generalizability using state-of-the-art methods.

II. RELATED WORK

3D Reconstruction methods are generally explicit [38]–[40] or implicit. Implicit representation encodes the shape as continuous function [26], [41] and can, in principle, handle more complex geometries. For Real2Sim, we utilize ConvONet [26], which provides a flexible implicit representation to reconstruct 3D scenes with fine-grained details.

Sim2Real gap is a common problem when transferring a simulator policy to the real-world. While existing techniques, such as domain randomization and system identification, can reduce the Sim2Real gap [9], [42], [43], they focus on generalization with respect to the low-dimensional dynamics parameters. We instead study geometric Real2Sim2Real, aligning the high-dimensional geometry to improve Sim2Real transfer.

Geometric Real2Sim2Real is a generative problem: given observations of a real scene, an algorithm should *generate* the corresponding digital replica. Generative algorithms typically fall into two classes. Firstly, the algorithm retrieves from an external database during the generation process, an approach taken by [37]. The second class of approaches relies solely on the trained generative model at test time. A representative algorithm is [44]. While [37], [44] can reconstruct objects and scenes, they do not demonstrate that the reconstruction can be used to train networks to perform robotic manipulation tasks. The ability to use the reconstructed scenes to

train networks is an important step towards solving *geometric Real2Sim2Real*. Moreover, [44] requires *manual* placement of the reconstructed meshes into simulation. They also affix the reconstructed meshes to their supporting surface to ensure the meshes do not fall over. Such practice is not suitable for applications that require separation between the object meshes and their supporting surface, such as pick-and-place. On the other hand, [37] performs reconstruction by retrieval from a database, and hence the reconstructed objects might not well approximate the real object shape. Also differently from [45], we demonstrate that our reconstructions can enable training robust neural networks.

III. PROBLEM STATEMENT

Our study investigates whether state-of-the-art implicit reconstruction methods can generate digital replicas that enable automated grasp labelling and the training of robust manipulation models. Demonstrating the ability of the Real2Sim2Real approach to train robust grasping networks is crucial since 6 DoF grasping is a fundamental building block for many manipulation skills.

We focus on grasping common household objects on a typical kitchen tabletop with a parallel gripper robot arm. The objective is to grasp and lift the target object out of the structured clutter without knocking over other neighboring objects. Examples of simulated and real scenes are shown in [Figure 2](#). Our synthetic object dataset includes 546 object instances from four categories (bottle, bowl, can, and mug), taken from PartNet-Mobility and ShapeNetCorev2 [46], [47]. Each of our 30 simulated scenes contains 5 to 10 randomly chosen objects and is constructed by rejection sampling to ensure that all objects are standing upright with no collision. In the real-world scenes, we also used five objects from the same four categories. We refer to the simulated scenes and real-world scenes collectively as test scenes in our subsequent discussion.

We extensively studied the grasping problem with a hand-to-eye setting, where we position RGB-D cameras above the table and looked at the table center to capture posed depth images from the test scenes. Given posed depth images and segmentation masks of the test scenes as inputs, the reconstruction algorithm outputs posed object meshes which are used to construct a replica in a physical simulator. We then use the reconstructed scene to train grasping networks, such as Contact-GraspNet [35] or Volumetric Grasping Network [36]. The trained grasping network is evaluated in the original test scenes and we use the grasp success rate to measure the performance of our Real2Sim2Real pipeline.

We demonstrate our results using test scenes in both simulation and real-world. For simulation environments, to account for realistic depth camera noises, we adapted SimKinect, which is a popular noise injection algorithm to benchmark SLAM systems [48]–[50] to add noise to the simulated depth images. Our simulator SAPIEN [46] uses PhysX [51], a state-of-the-art physics engine, to simulate dynamic physics. To compare the grasping success of our method, we include a retrieval-based reconstruction method [37] as

Algorithm 2 Obtaining reconstructed mesh from object-level point cloud

Input: $p_WOlist \in R^{n \times 3}$, 3D positions of n object points expressed with respect to the world frame W

Input: Trained Convolutional ONet, represented by f

Output: A mesh constructed from the input point cloud p_WOlist , expressed in the world frame $mesh_W$

1: Obtain the max and min bounds of the input point cloud

$$bound_max = np.max(p_WOlist, axis = 0) \in R^3$$

$$bound_min = np.min(p_WOlist, axis = 0) \in R^3$$

2: Compute normalization parameters

$$center = (bound_max + bound_min)/2$$

$$scale = np.max(bound_max - bound_min)$$

3: Normalize point cloud and obtain mesh

$$p_WOlist_norm = (p_WOlist - center)/scale$$

$$mesh_O = MISE(f(p_WOlist_norm))$$

4: Transform mesh to world frame

$$T_WO = \begin{bmatrix} scale * I & center \\ 0 & 1 \end{bmatrix}$$

$$mesh_W = mesh_O.apply_transform(T_WO)$$

a baseline in our simulated test scenes. The retrieval-based method belongs to a different category of reconstruction method that does not use implicit representations to recover object shapes. The retrieval-based method therefore serves as a useful baseline to understand if *any* reconstruction method would lead to good reconstruction and grasping performance, or whether the particular choice of the reconstruction method leads to significant differences in grasping success.


IV. USING DIGITAL REPLICA TO TRAIN NETWORKS

We next provide description of the reconstruction framework in [subsection IV-A](#), the grasping algorithm in [subsection IV-B](#) and implementation details in [subsection IV-C](#).

A. Reconstruction framework without pose estimation

The first step of our framework consists of converting the depth maps and segmentation masks into object-level point clouds. That is, after this first step, each object-level point cloud only contains points that lie on the surface of the same object instance. For each object-level point cloud, we use a trained ConvONet [26] to produce an implicit representation of the object surface. The object mesh can then be extracted by querying the learned implicit representation using the Multiresolution Isosurface Extraction procedure introduced in [52].

We next describe how we obtain the object poses to place the reconstructed meshes into digital replica of the test scene. To do so, we first describe an important pre-processing step used in ConvONet. Given an object-level point cloud, represented with respect to the world frame,



(a) Without applying outlier removal, the outlier point drastically affect the centering and scaling operations, leading to incorrect reconstruction (the big blob under table). (b) Reconstruction of the same scene with outlier removal applied. The previously erroneously reconstructed mesh shown in the left was more accurate reconstruction.

Fig. 3: The figure illustrates the importance of applying point cloud outlier removal to the observed point cloud to improve the robustness of the reconstruction framework.

centering and scaling operations are applied to the object-level point cloud before inputted into ConvONet [26]. This step transforms the object-level point cloud into a canonical frame whose origin coincides with the origin of the world frame. Therefore, the inverse of the centering and scaling operations can be interpreted as the object pose. We can place the reconstructed mesh into simulation without ever having to estimate the object poses explicitly. This is a significant advantage considering how challenging object pose estimation can be. Even though the normalization operations discussed in this paragraph are commonly used in 3D vision algorithms as a pre-processing step, we posit that the usage of the inverse normalization operations to place objects into the digital replica is a less explored and useful technique that is of independent interest.

We describe the procedure in Algorithm 2 and visually in Figure 1 (Right). The step-by-step description of our reconstruction framework is in Algorithm 1. The object orientation is implicitly represented by the point cloud and hence the reconstructed object is already oriented correctly.

In addition to these steps, we apply two additional techniques to improve the reconstruction quality. First, we use statistical point cloud outlier removal [53] to remove outlier points in the object-level point cloud before the centering and scaling operations. Doing so is helpful because the centering and scaling parameters are sensitive to point cloud noises. Figure 3 provides an illustrative example of the benefit of noise removal. Second, given the object mesh constructed by the Multiresolution Isosurface Extraction (MISE) procedure, we perform approximate convex decomposition to obtain a simplified and smoother object mesh representation [54], since the mesh obtained by MISE often has complex geometry, leading to slow collision detection during simulation.

B. Neural networks used for comparing grasping success

We evaluate the reconstruction quality using the success rates of grasping networks trained using the reconstructed scenes. This protocol differs from the standard evaluation metrics used in the computer vision community, which only measure how well the reconstructed shape approximates the ground truth shape, e.g. Chamfer distance. Using the

grasping success as the evaluation metrics is important because we are interested in studying whether the reconstructed scenes can be used to train performant grasping networks. We next describe why Contact-GraspNet (CGN) [35] and Volumetric Grasping Network (VGN) [36] are reasonable grasping network choices for our use case.

CGN takes point cloud as input and predicts for every point in the point cloud a corresponding grasp pose and the probability of grasp success. VGN accepts a TSDF representation and directly predicts for each 3D voxel the grasp pose and quality. As such, the input modalities of CGN and VGN only require depth information, and not color. Such requirement is suitable for our use case because our reconstruction framework only reconstructs the shape and not the color of the objects. We therefore only have access to the depth maps of the scenes and not color images. CGN and VGN are also recent state-of-the-art grasping networks with excellent reported performance. Their authors also release pre-trained models, e.g. a pre-trained CGN model trained using 17.7 million simulated grasps [55]. The pre-trained models serve as strong baseline for us to calibrate the performance of the models trained using our reconstructions.

C. Implementation Details

To perform grasping using CGN, given a desired pose of the gripper, we use the MPlib library [56] to plan a trajectory in joint space that moves the robot gripper from the current to the desired pose while avoiding collision. The VGN model takes as input a cubic voxel grid of dimension 40 with voxel size $7.5mm$. The voxel grid spans a 3D cube of size $30cm$, which is not sufficient to include all the objects in our scene. To evaluate the pre-trained CGN model without changing the input voxel grid dimension, for each object in a test scene, we obtain a different cubic voxel grid of dimension 40. The center of the voxel grid coincides with the center of the object axis aligned bounding box. We use the pre-trained CGN models to obtain predicted grasp successes and poses for each voxel grid, after which their union constitute the predicted grasp successes and poses for the entire scene.

To train ConvONet, we generate the training data from 80 training scenes using the same methodology we use to generate the test scenes. The objects used to generate the training and test scenes form disjoint sets. We refer readers to [26], [44], [52] for details on the ConvONet training.

V. EXPERIMENTAL RESULTS

Our experimental results answer the following questions:

- 1) What are the performances of representative grasping networks trained with our pipeline against baseline models that are trained with retrieval-based reconstruction methods or pre-trained models that are trained on significantly larger datasets?
- 2) What changes do we need to make to representative grasping networks training recipe to better accommodate the proposed pipeline?
- 3) Do the networks trained using the reconstructions generalize to unseen scene configurations?

Grasp Proposal Network	Grasp Success
CGN trained using reconstructed scenes (Ours)	0.93 (0.25)
CGN trained using Acronym dataset	0.0 (0.0)

TABLE I: The same grasping network is more robust when trained using our reconstructed scenes compared to the model trained using millions of grasps. The model trained using Acronym over-fits to geometry of the Panda robotic gripper used during training and does not generalize to the xArm gripper, which we use in our experiment. Qualitatively, the model generates grasps that are always in collision with the objects, and therefore has poor grasp success rate.

Dataset for training grasping models	CGN	VGN
ShapeNet groundtruth models and scenes	0.77 (0.16)	0.1 (0.04)
Reconstructed scenes only (ours)	0.76 (0.15)	0.34 (0.14)

TABLE II: The table demonstrates the quality of the reconstructed scene. The CGN and VGN model trained using our reconstructed scenes performs favorably relatively to the model pre-trained using ShapeNet. Since the CGN models trained using our scenes are trained from scratch, their good performance demonstrates that the reconstructed scenes allow for training performant grasping networks.

A. Real-world 6-DOF grasping in clutter experiments

The experiments using simulated scenes allow us to produce reproducible and detailed analyses in the upcoming sections. To provide evidence that our framework can generalize to real-world data, even though the reconstruction and grasping neural networks are only trained in simulation, we additionally run real-world grasping experiments using an xArm 6, illustrated in Figure 2 (Right). Table I demonstrates that CGN pre-trained with Panda gripper cannot generalize to the real scene where the grasping is executed with xArm gripper. However, by training CGN with grasp pose labels generated using the xArm gripper in the reconstructions, the trained model can reach reasonable grasping performance.

B. Simulated experiments comparing against pre-trained grasping models and retrieval-based reconstruction method

Result for CGN Table II illustrates the quality of our reconstructed scenes, in the sense that the CGN model trained using our reconstructed scenes performs well relatively to the publicly available pre-trained model. This is a particularly exciting result because the publicly released model was trained using 17.7 million simulated grasps. For the results presented in Table II, we train the CGN models from scratch and do not use external grasping dataset or finetune from existing network weights to avoid confounding factors when interpreting the success of the trained CGN models. We also train a different CGN model for each reconstructed scene. That is, for each reconstructed scene, we overfit a different CGN models to the reconstructed scene. As such, the good performance of the CGN models trained using the reconstructed scenes in Table II can be mainly attributed to the high quality reconstruction.

Reconstruction	CGN Training Epoch	Grasp Success
Ours	5000	0.76 (0.15)
Retrieval-based	2000	0.61 (0.1)
	3000	0.64 (0.1)
	5000	0.66 (0.1)
	6000	0.65 (0.1)
	8000	0.62 (0.1)

TABLE III: We outperform a retrieval-based reconstruction baseline, in terms of the grasping success of the grasping networks trained using the reconstruction.

Comparison to a retrieval-based method We next compare the performance of CGN when trained using our reconstruction framework versus using a retrieval-based scene reconstruction method [37]. When training using the scenes reconstructed by our method, we simply overfit the CGN model, i.e. perform gradient descent until all training losses are minimized. We find that such model selection scheme is sufficient. However, when training CGN using the scenes reconstructed by the retrieval-based method [37], we find that simply overfitting the CGN models to the reconstructions do not lead to the highest performance - possibly because the retrieved models fail to approximate the geometry of the objects in the test scenes well. We therefore evaluate the CGN models trained with different number of epochs for fair comparison. To obtain the RGBD trajectories that the retrieval-based method expects as input, we use an optimization-based exploration method to find camera trajectory that minimizes the area of occlusion [57]. Table III illustrates that CGN has higher performance when trained using our reconstruction framework, demonstrating that not all reconstruction method leads to good performance when training neural networks.

Result for VGN The VGN grasping networks trained from scratch using our reconstructions performs competitively with the pre-trained model released by VGN’s authors. The results are illustrated in Table II. Since we train the VGN networks from scratch and only use the reconstructed scenes for label generation, we can conclude that the good performance of the VGN models are due to the quality of the reconstructions. We note that the grasping performance of VGN and CGN are not directly comparable because we use motion planning to avoid collision when evaluating the predicted grasps for CGN, but not VGN. In the original publication, VGN does not use motion planning to avoid collision and therefore we follow the same protocol. For the pre-trained VGN model, the predicted grasps often collide with the target or neighboring objects during grasp execution. To provide more insights into the performance of VGN, we include the grasp success per object category in Table IV.

C. Training Modifications to CGN and VGN

Grasping networks are often trained using a large dataset of grasp labels. For example, CGN was trained using ACRONYM [8], a dataset of 17.7M simulated parallel-jaw grasps of 8872 objects. In contrast, we train on test scene reconstructions and then evaluate the model in the actual test

Object category	bottle	can	bowl	mug
Grasp Success	0.22 (0.2)	0.31 (0.2)	0.42 (0.3)	0.41 (0.3)

TABLE IV: Grasp success of the VGN models trained using our reconstructions on different object categories. The differences in success rate between the object categories and the high values of the standard deviation indicate that some objects are much harder for the models than others.

Grasp labeling when training using reconstructions	Grasp Success
Filtering using mean success rate	0.76 (0.15)
Using one execution trial	0.716 (0.09)

TABLE V: The table demonstrates the importance of determining grasp success label by executing the grasp multiple times, and only label as successful grasp pose whose mean success rate is above a pre-defined threshold.

scenes. To probe the reconstruction quality, we only train on a much smaller dataset sampled from the reconstructions and thus are required to make some modifications to the training process.

Modifications for CGN: CGN predicts gripper grasp poses. A motion planning algorithms converts the pose into a joint-level trajectory for the robot controller to track. As such, if the planner is stochastic, evaluating a grasp multiple times from the same initial conditions can lead to different results, depending on whether the planned joint-level trajectory leads to collision with objects in the scenes. When we generate the grasp label using the reconstructed scenes, we thus only labels as success grasps that are likely to pick up the target objects under the stochasticity of the planner. More concretely, instead of using one execution trial to evaluate the grasp [8], during grasp label generation, we evaluate a grasp pose multiple times from the same initial scene condition, and only label the pose as successful when the mean success rate is above a pre-defined threshold. [Table V](#) illustrates quantitatively the importance of doing so. The threshold is 0.7 in all our experiments.

Modifications for VGN: VGN takes as input a 3D voxel grid and predicts for each voxel a grasp pose and the corresponding success probability. Instead of only computing the loss and gradients through the voxels in which ground truth grasp label is available, we compute the loss and




Fig. 4: **Left:** VGN assigns high success probability to voxels far away from object surface. **Right:** the network no longer assigns high success probability to these voxels when the loss is computed through all input voxels. The green grasps represent grasps with high predicted success probability.

Input voxels without grasp annotations	Grasp Success
Without assigning negative labels	34.0
With assigning negative labels	60.0

TABLE VI: VGN performs better when we compute the loss through all input voxels and assign negative grasp labels to voxels without annotations. The models are trained with grasp annotation generated from the ground truth test scene.




Fig. 5: **Left:** a test scene **Right:** after rearranging the objects in the left test scene, we evaluate CGN trained in the reconstruction of the **left** in the right scene.

gradients through all voxels in the input voxel grid. For voxels where there is no grasp annotation, we label the voxel with a negative grasp label. Without our modifications, the VGN model can assign high grasp success probability in 3D region where there is no nearby voxels with grasp annotation label, as illustrated in [Figure 4](#) and [Table VI](#).

D. Generalization performance to unseen scenes

We previously train grasping networks using the reconstructions of the test scenes the grasping networks are evaluated in. In this section, we devise additional experiments to understand the generalization behavior of the grasping networks trained from the reconstructions as the dataset size and quality vary. Given the limited time and the amount of manual efforts to rearrange the selected scenes, we select 5 test scenes where the CGN models (trained using the reconstructions of the test scenes the models are evaluated in) have good performance. We then vary the dataset size and quality while training CGN and evaluating the trained models in these 5 test scenes.

Rearrange setting: for each of the 5 test scenes, we rearrange the objects in the test scene into different poses, as illustrated in [Figure 5](#). We then evaluate the CGN models trained using the reconstruction of the test scene (before rearrangement) in the re-arranged scene. In other words, the CGN models were trained on the reconstructions of the objects in the test scenes, but the arrangement of the objects has changed when evaluating the CGN models. This setting represents generalization to unseen scenes with seen objects.

Training using reconstructions of scenes different from test scenes: in previous sections, we train one CGN models per reconstructions of the test scenes. In this experimental setting, we instead train one CGN model on reconstructions of *multiple* test scenes, but evaluate the trained models in test scenes whose reconstructions are not included in the set of reconstructions used for training the CGN models.

The grasp success of the trained CGN models in these two different settings are illustrated in [Figure 6](#). The results suggest that having more reconstructed scenes improve




Fig. 6: The performance of the Contact-GraspNet increases as a function of both the number of training reconstructed scenes and how close the reconstructions resemble the test scene. The number 5, 10, 20, 29 indicates the number of training reconstructed scenes, holding out the test scene reconstruction. *Rearrange* refers to the rearrange setting described in subsection V-D. *Test* refers to the setting where the CGN model is trained on the reconstructed test scene only.

generalization, and the trained models can generalize to rearranged scenes with a minor drop in performance.

VI. CONCLUSIONS

We explore the geometric Real2Sim2Real pipeline on a tabletop clutter grasping task. Our synthetic and real experiments demonstrate robust grasping results when state-of-the-art reconstruction and grasping networks are combined wisely using a modern physical simulator for bridging. We also show promising performance when adapting existing grasping networks to a new robot system following our pipeline without having to label a large dataset. Our work illustrates the good performance of current neural surface reconstruction methods for physical simulation and the usefulness of retraining recent grasping networks with simulated grasp sampling when deploying onto a new robot system.

REFERENCES

- [1] J. Mahler, J. Liang, S. Niyaz, M. Laskey, R. Doan, X. Liu, J. A. Ojea, and K. Goldberg, “Dex-net 2.0: Deep learning to plan robust grasps with synthetic point clouds and analytic grasp metrics,” *CoRR*, vol. abs/1703.09312, 2017. [Online]. Available: <http://arxiv.org/abs/1703.09312>
- [2] J. Mahler, F. T. Pokorny, B. Hou, M. Roderick, M. Laskey, M. Aubry, K. Kohlhoff, T. Kröger, J. Kuffner, and K. Goldberg, “Dex-net 1.0: A cloud-based network of 3d objects for robust grasp planning using a multi-armed bandit model with correlated rewards,” in *2016 IEEE International Conference on Robotics and Automation (ICRA)*, 2016, pp. 1957–1964.
- [3] J. Mahler, M. Matl, X. Liu, A. Li, D. V. Gealy, and K. Goldberg, “Dex-net 3.0: Computing robust robot suction grasp targets in point clouds using a new analytic model and deep learning,” *CoRR*, vol. abs/1709.06670, 2017. [Online]. Available: <http://arxiv.org/abs/1709.06670>
- [4] D. Ho, K. Rao, Z. Xu, E. Jang, M. Khansari, and Y. Bai, “Retinagan: An object-aware approach to sim-to-real transfer,” *CoRR*, vol. abs/2011.03148, 2020. [Online]. Available: <https://arxiv.org/abs/2011.03148>
- [5] Z. Xu, C. Chi, B. Burchfiel, E. Cousineau, S. Feng, and S. Song, “Dextairity: Deformable manipulation can be a breeze,” 2022. [Online]. Available: <https://arxiv.org/abs/2203.01197>
- [6] W. Zhou and D. Held, “Learning to grasp the ungraspable with emergent extrinsic dexterity,” in *ICRA 2022 Workshop: Reinforcement Learning for Contact-Rich Manipulation*, 2022. [Online]. Available: <https://openreview.net/forum?id=Zrp4wpa9lqh>
- [7] A. Agarwal, A. Kumar, J. Malik, and D. Pathak, “Legged locomotion in challenging terrains using egocentric vision,” in *6th Annual Conference on Robot Learning*, 2022. [Online]. Available: <https://openreview.net/forum?id=Re3NjSwf0WF>
- [8] C. Eppner, A. Mousavian, and D. Fox, “ACRONYM: A large-scale grasp dataset based on simulation,” in *2021 IEEE Int. Conf. on Robotics and Automation, ICRA*, 2020.
- [9] J. Tobin, R. Fong, A. Ray, J. Schneider, W. Zaremba, and P. Abbeel, “Domain randomization for transferring deep neural networks from simulation to the real world,” *CoRR*, vol. abs/1703.06907, 2017. [Online]. Available: <http://arxiv.org/abs/1703.06907>
- [10] F. Sadeghi and S. Levine, “(cad)\$2\$rl: Real single-image flight without a single real image,” *CoRR*, vol. abs/1611.04201, 2016. [Online]. Available: <http://arxiv.org/abs/1611.04201>
- [11] X. B. Peng, M. Andrychowicz, W. Zaremba, and P. Abbeel, “Sim-to-real transfer of robotic control with dynamics randomization,” *CoRR*, vol. abs/1710.06537, 2017. [Online]. Available: <http://arxiv.org/abs/1710.06537>
- [12] Q. Vuong, S. Vikram, H. Su, S. Gao, and H. I. Christensen, “How to pick the domain randomization parameters for sim-to-real transfer of reinforcement learning policies?” *CoRR*, vol. abs/1903.11774, 2019. [Online]. Available: <http://arxiv.org/abs/1903.11774>
- [13] N. Ruiz, S. Schuler, and M. Chandraker, “Learning to simulate,” *CoRR*, vol. abs/1810.02513, 2018. [Online]. Available: <http://arxiv.org/abs/1810.02513>
- [14] W. Yu, C. K. Liu, and G. Turk, “Policy transfer with strategy optimization,” in *International Conference on Learning Representations*, 2019. [Online]. Available: <https://openreview.net/forum?id=H1g6osRcFQ>
- [15] A. Prakash, S. Boochoon, M. Brophy, D. Acuna, E. Cameracci, G. State, O. Shapira, and S. Birchfield, “Structured domain randomization: Bridging the reality gap by context-aware synthetic data,” *CoRR*, vol. abs/1810.10093, 2018. [Online]. Available: <http://arxiv.org/abs/1810.10093>
- [16] S. James, P. Wohlhart, M. Kalakrishnan, D. Kalashnikov, A. Irpan, J. Ibarz, S. Levine, R. Hadsell, and K. Bousmalis, “Sim-to-real via sim-to-sim: Data-efficient robotic grasping via randomized-to-canonical adaptation networks,” *CoRR*, vol. abs/1812.07252, 2018. [Online]. Available: <http://arxiv.org/abs/1812.07252>
- [17] S. Zakharov, W. Kehl, and S. Ilic, “Deceptionnet: Network-driven domain randomization,” *CoRR*, vol. abs/1904.02750, 2019. [Online]. Available: <http://arxiv.org/abs/1904.02750>
- [18] A. Szot, A. Clegg, E. Undersander, E. Wijnmans, Y. Zhao, J. Turner, N. Maestre, M. Mukadam, D. Chaplot, O. Maksymets, A. Gokaslan, V. Vondrus, S. Dharur, F. Meier, W. Galuba, A. Chang, Z. Kira, V. Koltun, J. Malik, M. Savva, and D. Batra, “Habitat 2.0: Training home assistants to rearrange their habitat,” 2021. [Online]. Available: <https://arxiv.org/abs/2106.14405>
- [19] H. Fu, B. Cai, L. Gao, L.-X. Zhang, J. Wang, C. Li, Q. Zeng, C. Sun, R. Jia, B. Zhao *et al.*, “3d-front: 3d furnished rooms with layouts and semantics,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 10933–10942.
- [20] H. Fu, R. Jia, L. Gao, M. Gong, B. Zhao, S. Maybank, and D. Tao, “3d-future: 3d furniture shape with texture,” *International Journal of Computer Vision*, pp. 1–25, 2021.
- [21] A. Kalervo, J. Ylioinas, M. Häikiö, A. Karhu, and J. Kannala, “Cubicasa5k: A dataset and an improved multi-task model for floorplan image analysis,” *CoRR*, vol. abs/1904.01920, 2019. [Online]. Available: <http://arxiv.org/abs/1904.01920>
- [22] V. Lim, H. Huang, L. Y. Chen, J. Wang, J. Ichnowski, D. Seita, M. Laskey, and K. Goldberg, “Planar robot casting with real2sim2real self-supervised learning,” 2021. [Online]. Available: <https://arxiv.org/abs/2111.04814>
- [23] Y. Chebotar, A. Handa, V. Makovychuk, M. Macklin, J. Issac, N. Ratliff, and D. Fox, “Closing the sim-to-real loop: Adapting simulation randomization with real world experience,” 2018. [Online]. Available: <https://arxiv.org/abs/1810.05687>
- [24] Y. Du, O. Watkins, T. Darrell, P. Abbeel, and D. Pathak, “Auto-tuned sim-to-real transfer,” *CoRR*, vol. abs/2104.07662, 2021. [Online]. Available: <https://arxiv.org/abs/2104.07662>

- [25] E. Kolve, R. Mottaghi, W. Han, E. VanderBilt, L. Weihs, A. Herrasti, D. Gordon, Y. Zhu, A. Gupta, and A. Farhadi, "Ai2-thor: An interactive 3d environment for visual ai," 2017. [Online]. Available: <https://arxiv.org/abs/1712.05474>
- [26] S. Peng, M. Niemeyer, L. M. Mescheder, M. Pollefeys, and A. Geiger, "Convolutional occupancy networks," *CoRR*, vol. abs/2003.04618, 2020. [Online]. Available: <https://arxiv.org/abs/2003.04618>
- [27] J. J. Park, P. Florence, J. Straub, R. A. Newcombe, and S. Lovegrove, "DeepSDF: Learning continuous signed distance functions for shape representation," *CoRR*, vol. abs/1901.05103, 2019. [Online]. Available: <http://arxiv.org/abs/1901.05103>
- [28] P. Wang, L. Liu, Y. Liu, C. Theobalt, T. Komura, and W. Wang, "Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction," *CoRR*, vol. abs/2106.10689, 2021. [Online]. Available: <https://arxiv.org/abs/2106.10689>
- [29] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, "Nerf: Representing scenes as neural radiance fields for view synthesis," *CoRR*, vol. abs/2003.08934, 2020. [Online]. Available: <https://arxiv.org/abs/2003.08934>
- [30] A. Kar, A. Prakash, M.-Y. Liu, E. Cameracci, J. Yuan, M. Rusiniak, D. Acuna, A. Torralba, and S. Fidler, "Meta-sim: Learning to generate synthetic datasets," 2019. [Online]. Available: <https://arxiv.org/abs/1904.11621>
- [31] J. Devaranjan, A. Kar, and S. Fidler, "Meta-sim2: Unsupervised learning of scene structure for synthetic data generation," 2020. [Online]. Available: <https://arxiv.org/abs/2008.09092>
- [32] A. Gaidon, Q. Wang, Y. Cabon, and E. Vig, "Virtual worlds as proxy for multi-object tracking analysis," 2016. [Online]. Available: <https://arxiv.org/abs/1605.06457>
- [33] Y. Cabon, N. Murray, and M. Humenberger, "Virtual kitti 2," 2020. [Online]. Available: <https://arxiv.org/abs/2001.10773>
- [34] M. Igl, D. Kim, A. Kuefler, P. Mougin, P. Shah, K. Shiarlis, D. Anguelov, M. Palatucci, B. White, and S. Whiteson, "Symphony: Learning realistic and diverse agents for autonomous driving simulation," 2022. [Online]. Available: <https://arxiv.org/abs/2205.03195>
- [35] M. Sundermeyer, A. Mousavian, R. Triebel, and D. Fox, "Contact-graspnet: Efficient 6-dof grasp generation in cluttered scenes," *CoRR*, vol. abs/2103.14127, 2021. [Online]. Available: <https://arxiv.org/abs/2103.14127>
- [36] M. Breyer, J. J. Chung, L. Ott, R. Siegwart, and J. I. Nieto, "Volumetric grasping network: Real-time 6 DOF grasp detection in clutter," *CoRR*, vol. abs/2101.01132, 2021. [Online]. Available: <https://arxiv.org/abs/2101.01132>
- [37] M. Han, Z. Zhang, Z. Jiao, X. Xie, Y. Zhu, S.-C. Zhu, and H. Liu, "Reconstructing interactive 3d scenes by panoptic mapping and cad model alignments," 2021. [Online]. Available: <https://arxiv.org/abs/2103.16095>
- [38] M. Nießner, M. Zollhöfer, S. Izadi, and M. Stamminger, "Real-time 3d reconstruction at scale using voxel hashing," *ACM Trans. Graph.*, vol. 32, no. 6, pp. 169:1–169:11, 2013. [Online]. Available: <https://doi.org/10.1145/2508363.2508374>
- [39] R. Chen, S. Han, J. Xu, and H. Su, "Visibility-aware point-based multi-view stereo network," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 43, no. 10, pp. 3695–3708, 2021. [Online]. Available: <https://doi.org/10.1109/TPAMI.2020.2988729>
- [40] N. Wang, Y. Zhang, Z. Li, Y. Fu, W. Liu, and Y. Jiang, "Pixel2mesh: Generating 3d mesh models from single RGB images," in *Computer Vision - ECCV 2018 - 15th European Conference, Munich, Germany, September 8-14, 2018, Proceedings, Part XI*, ser. Lecture Notes in Computer Science, V. Ferrari, M. Hebert, C. Sminchisescu, and Y. Weiss, Eds., vol. 11215. Springer, 2018, pp. 55–71. [Online]. Available: https://doi.org/10.1007/978-3-030-01252-6_4
- [41] J. J. Park, P. Florence, J. Straub, R. A. Newcombe, and S. Lovegrove, "DeepSDF: Learning continuous signed distance functions for shape representation," in *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*. Computer Vision Foundation / IEEE, 2019, pp. 165–174. [Online]. Available: http://openaccess.thecvf.com/content_CVPR_2019/html/Park_DeepSDF_Learning_Continuous_Signed_Distance_Functions_for_Shape_Representation.CVPR.2019_paper.html
- [42] OpenAI, M. Andrychowicz, B. Baker, M. Chociej, R. Józefowicz, B. McGrew, J. Pachocki, A. Petron, M. Plappert, G. Powell, A. Ray, J. Schneider, S. Sidor, J. Tobin, P. Welinder, L. Weng, and W. Zaremba, "Learning dexterous in-hand manipulation," *CoRR*, vol. abs/1808.00177, 2018. [Online]. Available: <http://arxiv.org/abs/1808.00177>
- [43] Z. Xie, X. Da, M. van de Panne, B. Babich, and A. Garg, "Dynamics randomization revisited: A case study for quadrupedal locomotion," in *IEEE International Conference on Robotics and Automation, ICRA 2021, Xi'an, China, May 30 - June 5, 2021*. IEEE, 2021, pp. 4955–4961. [Online]. Available: <https://doi.org/10.1109/ICRA48506.2021.9560837>
- [44] Z. Jiang, C.-C. Hsu, and Y. Zhu, "Ditto: Building digital twins of articulated objects from interaction," 2022. [Online]. Available: <https://arxiv.org/abs/2202.08227>
- [45] J. Lv, Q. Yu, L. Shao, W. Liu, W. Xu, and C. Lu, "SAGCI-system: Towards sample-efficient, generalizable, compositional, and incremental robot learning," in *ICRA 2022 Workshop: Reinforcement Learning for Contact-Rich Manipulation*, 2022. [Online]. Available: <https://openreview.net/forum?id=Fdj08qJzXh>
- [46] F. Xiang, Y. Qin, K. Mo, Y. Xia, H. Zhu, F. Liu, M. Liu, H. Jiang, Y. Yuan, H. Wang, L. Yi, A. X. Chang, L. J. Guibas, and H. Su, "SAPIEN: A simulated part-based interactive environment," *CoRR*, vol. abs/2003.08515, 2020. [Online]. Available: <https://arxiv.org/abs/2003.08515>
- [47] A. Chang, T. Funkhouser, L. Guibas, P. Hanrahan, Q. Huang, Z. Li, S. Savarese, M. Savva, S. Song, H. Su, J. Xiao, L. Yi, and F. Yu, "Shapenet: An information-rich 3d model repository," 12 2015.
- [48] A. Handa, T. Whelan, J. McDonald, and A. J. Davison, "A benchmark for rgb-d visual odometry, 3d reconstruction and slam," *ICRA*, 2014.
- [49] J. T. Barron and J. Malik, "Intrinsic scene properties from a single rgb-d image," *CVPR*, 2013.
- [50] J. Bohg, J. Romero, A. Herzog, and S. Schaal, "Robot arm pose estimation through pixel-wise part classification," *ICRA*, 2014.
- [51] Nvidia, "PhysX physics engine," <https://www.geforce.com/hardware/technology/physx>.
- [52] L. M. Mescheder, M. Oechsle, M. Niemeyer, S. Nowozin, and A. Geiger, "Occupancy networks: Learning 3d reconstruction in function space," *CoRR*, vol. abs/1812.03828, 2018. [Online]. Available: <http://arxiv.org/abs/1812.03828>
- [53] Q.-Y. Zhou, J. Park, and V. Koltun, "Open3d: A modern library for 3d data processing," 2018, cite arxiv:1801.09847Comment: <http://www.open3d.org>. [Online]. Available: <http://arxiv.org/abs/1801.09847>
- [54] K. Mamou, "Volumetric approximate convex decomposition," in *Game Engine Gems 3*, E. Lengyel, Ed. A K Peters / CRC Press, 2016, ch. 12, pp. 141–158.
- [55] "Cgn," https://github.com/NVlabs/contact_graspnet.
- [56] L. Minghua and J. Gu, "Mplib," <https://github.com/haosulab/MPLib>, 2021.
- [57] G. Kahn, P. Suján, S. Patil, S. Bopardikar, J. Ryde, K. Goldberg, and P. Abbeel, "Active exploration using trajectory optimization for robotic grasping in the presence of occlusions," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*, 2015, pp. 4783–4790.