Combining Shape Completion and Grasp Prediction for Fast and Versatile Grasping with a Multi-Fingered Hand

Matthias Humt^{1,2}_{*}, Dominik Winkelbauer^{1,2}_{*}, Ulrich Hillenbrand¹, and Berthold Bäuml^{1,3}

Abstract-Grasping objects with limited or no prior knowledge about them is a highly relevant skill in assistive robotics. Still, in this general setting, it has remained an open problem, especially when it comes to only partial observability and versatile grasping with multi-fingered hands. We present a novel, fast, and high fidelity deep learning pipeline consisting of a shape completion module that is based on a single depth image, and followed by a grasp predictor that is based on the predicted object shape. The shape completion network is based on VQDIF and predicts spatial occupancy values at arbitrary query points. As grasp predictor, we use our twostage architecture that first generates hand poses using an autoregressive model and then regresses finger joint configurations per pose. Critical factors turn out to be sufficient data realism and augmentation, as well as special attention to difficult cases during training. Experiments on a physical robot platform demonstrate successful grasping of a wide range of household objects based on a depth image from a single viewpoint. The whole pipeline is fast, taking only about 1s for completing the object's shape (0.7 s) and generating 1000 grasps (0.3 s). (https://dlr-alr.github.io/2023-humanoids-completion/)

I. INTRODUCTION

A major step towards real-world autonomy of robot systems in general, and of versatile humanoids in particular, is their ability to grasp objects at a first-time encounter, that is, a successful handling of novel objects. This scenario is likely to occur, e.g., in assistive robotics operating in human living and working environments, or indeed in any environment that is not strongly constrained in terms of objects and tasks (hospitals, elderly care homes, shop floors, offices).

Accordingly, the goal of unknown object grasping has been pursued for more than a decade. In early studies, one main approach was to transfer grasps from a database of known objects to novel ones with similar shapes [1, 2]. However, only the grasp pose and the preshape of the robot hand or gripper were retrieved in this way, without having control over the grasp contacts finally made. Grasp poses were also inferred from probabilistic models [3], and the preshape from Eigengrasp planners [4]. Going one step further, in [5, 6], we have warped individual finger contacts onto novel objects, but also being restricted to a high level of shape similarity. All in all, the early approaches were suffering from a lack of generality and control. More recently, there has been progress on grasping of unknown objects from a wider range [7], but being tailored to parallel jaw grippers and hence simple specific grasp types.

- ¹ The author is with the DLR Institute of Robotics and Mechatronics.
- 2 The author is with the Technical University of Munich (TUM).
- ³ The author is with the Deggendorf Institute of Technology (DIT).
- * Equal contribution.
- Contact: matthias.humt | berthold.baeuml@dlr.de



Fig. 1. Grasping the YCB bleach bottle using our grasping pipeline: The object is first perceived using Agile Justin's [8] Kinect camera to obtain a single depth image. Afterward, the shape completion network predicts the full object shape, based on which the grasping network generates a stable grasp. The grasp is then executed on the real robot using whole-body motion planning [9, 10] for positioning the hand relative to the object and a kinematically calibrated robotic system [11, 12]

Using a dexterous multi-fingered robot hand allows a wide repertoire of complex grasps, leading to applications ranging from pick-and-place tasks [13] to object manipulation [14] and to the utilization of tools via functional grasps [15]. However, the exploitation of its full potential under the condition of very limited prior object knowledge requires a detailed reconstruction of the complete object shape and, based on this, a shape-specific prediction of the finger configuration. Using these two critical components, we have integrated the first system that can predict and execute in real-time stable grasps with a fully actuated four-finger hand on a wide range of initially unknown object shapes.

Accurate and complete object shape reconstruction from partial and noisy depth data as can be acquired from a single viewpoint is challenging. We here take shape completion to the real world by introducing a procedure for synthetic training data generation and the training itself that can bridge the gap to the real sensing conditions. Likewise, we use the same synthetic data generation procedure together with the trained shape completion network to acquire realistic object reconstructions for training the grasp network. We further describe an extension to the grasp planner which increases the robustness of the generated training grasps against uncertainties in the relative pose between object and hand.

We show what is important in training shape completion and grasp prediction as a grasping pipeline for the real world. Critical factors are the careful generation of training data as well as a training procedure that emphasizes difficult cases. Our main contributions are:



Fig. 2. Our pipeline consists of two steps during inference. From a partial point cloud of an object obtained through rendering (training) or a depth sensor (inference), we use a shape completion network [16] to predict the full object geometry implicitly as occupancy probabilities at query points on a grid from which we can optionally extract the completed mesh using Marching Cubes [17]. Once this network is trained, the predicted occupancy probabilities are used as input for the grasp predictor [18] both during training on ground truth (GT) grasps as well as during deployment on our robot Agile Justin.

- A procedure for synthetic data generation and training of shape completion for improved sim2real transfer.
- A quantitative evaluation on synthetic and real data of the effects of the improved data and training procedure.
- An extension of the analytical grasp planner to generate grasps that are more resilient against uncertainties in the relative pose between hand and object.
- An adaptation of the grasp network architecture to handle ambiguities in the dataset without the requirement of generating multiple labels for each training sample.
- A validation of the proposed training procedures and of the system pipeline for reliable grasp generation and execution through a series of grasping experiments on our humanoid Agile Justin with a large range of initially unknown objects.

II. RELATED WORK

Our work is related to deep learning based completion or reconstruction of partial 3D data. There is a large number of approaches to handle this modality from 3D convolutional neural networks [19, 20], over point cloud-based methods [21, 22, 23] to 3D triangle mesh methods [24, 25]. In recent years, implicit function representations [26, 27, 28, 29, 30, 16] have gained traction due to their strong performance on diverse problems. While these works excel on benchmarks, they are seldom extended and evaluated under realistic conditions like varying viewpoints, object diversity and rotation, self-occlusion, and sensor artifacts. We make use of a network architecture introduced in [16] but train it on a large, realistic dataset to enable generalization to real-world scenarios.

There is a strong connection to the robotic grasping literature, specifically works that rely on additional geometric understanding to perform grasps [31, 32, 33, 34]. Varley et al. [31] are the most similar to our work but only train on a limited number of objects (480 vs. 57,000) with a

low-resolution voxel grid representation $(40^3 \text{ vs. continuous}$ implicit representation) and thus lack fidelity in completing unknown-and especially-real objects. Also, the grasps are optimized online using an analytical grasp planner, while we instantly predict stable grasps using our learning-based approach. Yan et al. [32] only shows qualitative shape completion results and uses a parallel jaw gripper for grasping. Also, Chen et al. [34] does not consider multi-fingered hands. der Merwe et al. [33] use an implicit function representation of the completed object and a multi-fingered hand but only show a single low-fidelity qualitative shape completion result on real data. Furthermore, the grasping prediction requires an online optimization, leading to long planning times.

In this work, we also propose a new method to handle ambiguities in the grasp training data. This issue has been usually approached by generating multiple labels for training samples and then using the one closest to the network's prediction for the computation of the loss [18, 35, 36]. While this reduces the ambiguity problem, it makes the data generation process more computationally expensive. In contrast to that, our proposed approach only requires one label per training sample and therefore does not suffer from that problem.

III. SHAPE COMPLETION

We use implicit functions to represent the 3D surface of an object as the continuous decision boundary of a *Deep Neural Network* (DNN) classifier predicting occupancy probability at each spatial location. Compared to discrete representations, implicit functions have the advantages of infinite resolution at a low memory footprint and are agnostic to the genus of the object geometry. During inference, the implicit function can be queried densely on a voxel grid and meshed using *Marching Cubes* [17] to extract an explicit mesh representation of the object. These networks are trained by giving partial observations like depth images or point clouds as input and supervising with occupancy labels at sampled spatial locations.

Specifically, we use the VQDIF [16] (Vector Quantized Deep Implicit Functions) network architecture which improves on Convolutional Occupancy Networks [26, 27] using a more compact sequence-based shape representation.

A. Training data generation

We leverage the ShapeNetCore (v1) subset of the ShapeNet [37] dataset containing more than 57,000 3D models across different categories to generate training data.

Many of the provided meshes are, however, not watertight. We, therefore, first pre-process the entire dataset using the mesh-fusion [38] pipeline with optimizations for speed and parallel execution to allow for the processing of large-scale datasets. We apply quadratic edge collapse decimation [39] and small disconnected component removal to reduce the triangle count by 95% and remove unwanted mesh-fusion artifacts. We sample 100,000 points from the surface of the watertight meshes. For supervision of the occupied space, we use multiple sampling strategies. We uniformly sample 100,000 points in the unit bounding cube of each object with a small padding of 0.1 on each side and determine their occupancy value using a triangle intersection test as in Mescheder et al. [26]. We further sample 10,000 points from the surface of the object and add Gaussian noise with various standard deviations for a total of an additional 100,000 points to increase the resolution of the supervision signal close to the surface similar to Xu et al. [29]. Additionally, we sample 100,000 points from spheres centered on the object with various radii to prevent missing occupancy information inside the unit cube under object rotation. We refer to these points as query points in the remainder of this article.

To simulate realistic partial views, we render depth, surface normal, and simulated Kinect depth images from views distributed on the upper hemisphere facing each object using an optimized version of the *Blensor* [40, 41] Kinect sensor model. Figure 3 shows a comparison of a scene rendered using a standard renderer, the Kinect simulator and the real Kinect sensor data. The simulation more closely resembles the real depth image, showing similar noise patterns and regions of missing data compared to the standard depth rendering. At the same time, there are still considerable differences and room for improvement to more accurately simulate real sensor artifacts.

For each object, we render 100 views with the camera positioned at varying distances for a total of around 5.7 million views which we split 9 to 1 per object category into training and validation sets. A test set is not necessary as we perform testing on novel objects and real data, as explained later. We also generate additional data where we randomly scale each object along its upright axis by up to 20% for every rendered view to further increase the object diversity, especially for categories with few examples.

The complete data generation pipeline is summarized in algorithm 1.



Fig. 3. Comparison of standard depth rendering (left), Kinect depth simulation (middle), and the real Kinect depth image (right). The object mesh is obtained via a laser scanner, and its pose through a registration step.

Algorithm 1 Data Generation

- 1: Obtain ShapeNetCore subset containing N meshes
- 2: for i = 1 to N do
- 3: Make mesh watertight
- 4: Normalize mesh
- 5: Sample 100,000 points from mesh surface
- 6: Sample 100,000 random points in unit cube
- 7: **for** r = 1 **to** 5 **do**
- 8: Sample 100,000 random sphere points (radius r)
- 9: end for 10: for s in $S \in$
- 0: for s in $S \in \{0.001, 0.25\}$ do
- 11: Sample 10,000 points from the mesh surface
- 12: Sample and add noise $\mathcal{N}(0,s)$

13: **end for**

- 14: Compute occupancy for all points
- 15: **for** v = 1 **to** 100 **do**
- 16: Apply random scale to mesh (optional)
- 17: Sample view on upper hemisphere
- 18: Render depth, normal and Kinect depth images
- 19: end for
- 20: end for
- 21: Split data into 90% training and 10% validation sets

B. Training

We train our shape completion networks in a fully supervised manner. For each training example, we randomly sample one of the ShapeNet objects and then one of its corresponding rendered depth or simulated Kinect depth images. We refer to the DNN trained on the simulated Kinect depth images as KINECT and the one trained on data with randomly scaled objects during data generation as KINECT SCALE. The depth image is projected to a 3D point cloud based on the known camera intrinsic parameters used during rendering. For the BASIC training setup, we add Gaussian noise to all points and additional noise to edge points which we partially remove as in Humt et al. [42]. We center the input data at its origin and scale its longest side to unit length.

We load the uniform, noisy surface and sphere query points for the given object, crop points outside the unit bounding cube, subsample 2048 queries, and transform them to match the scale and camera frame of the normalized input data. Since the robot's base frame of reference is known during deployment, we further transform all points from camera coordinates to the robot's world coordinate system using the robot's kinematic chain. This reduces the complexity of the learning problem by removing camera tilt. At each query point the occupancy probability is predicted and the ground truth binary labels are then used to compute the binary cross-entropy loss, which we use as training signal.

Once the training is complete, we improve the model performance further through an automated finetuning step. Akin to *Hard Negative Mining*, we employ an importance sampling strategy by oversampling object views proportional to the magnitude of their obtained loss. To mitigate the high class imbalance of ShapeNet (there are approximately 35 times more cars than mugs), we further resort to weighted batch sampling using the inverse class frequency as weights. The effects of this additional step are exemplified in Figure 7.

All methods are implemented in the C++ and Python programming languages. We use PyTorch [43] and Lightning [44] as our primary Deep Learning framework. Networks are Trained with the AdamW [45] optimizer with a learning rate of 1e-4 and weight decay of 0.01. All networks are trained for a maximum of 100 epochs with learning rate reduction when performance on the validation set plateaus for 10 epochs and early stopping after 30 epochs without improvement.

IV. GRASP PREDICTION

To efficiently predict grasps for a given unknown object, we are building on top of the learning-based approach proposed by Winkelbauer et al. [18]. Here, grasping is approached as a supervised learning problem, learning the mapping from a given observed object o to a set of grasps $\{x_1, ..., x_n\}$. Each grasp x_i is composed of a 6D hand pose h and a joint configuration $q \in \mathbb{R}^{12}$. The ground truth grasp annotations for synthetic training objects are generated using an analytical grasp planner. The grasp planner formulates grasping as an optimization problem by using a grasp quality metric as objective. Our quality metric is based on the epsilon quality metric [46], which is defined as the magnitude of the minimal external wrench that would break the grasp. As learning architecture, a two-stage neural network is used. The first stage is a generative network for hand poses that learns the mapping $\mathcal{G}: o \mapsto p(h|o)$. The second stage is a regressive network which learns the mapping $\mathcal{J}: (o,h) \mapsto$ $(q, \Delta h, s)$, predicting a joint configuration q, a small hand pose correction Δh , and a grasp quality s for each generated hand pose h. For further details, we refer the reader to Winkelbauer et al. [18].

A. Object pose uncertainty

When executing a grasp on the real system, multiple small inaccuracies in the grasping pipeline lead to an uncertainty in the actual relative positioning of object and hand. These inaccuracies can occur in the calibration of the camera and kinematic structure, the shape completion, and also



Fig. 4. The DLR-Hand II [47] which we use in our experiments. Each red dot represents a potential contact point. In the left image, all contact points used in the grasp planner are shown, while the right image shows only the central contact points that are used to verify that a finger is still in full contact with the object after slight pose variations.

the grasp prediction itself. Some grasps, however, require a very precise placement of the fingers to be successful. So, due to small inaccuracies, a finger might miss or slide over an edge, which can lead to a distortion of the grasp and even to dropping the object. To prevent such grasps, we adapt our grasp planner's objective in a way that we diminish the quality of grasps that are prone to missing the object due to small changes in its relative pose. To detect such grasps, we check the grasp under multiple small object pose deviations. In detail, we move the object pose using a translation Δt_i . To then resolve possible collisions with the hand base, the hand is moved backward along the approach direction until any collision with the hand base is resolved. Afterward, the fingers are closed until they touch the object or until they reach the limits of the joints. For each joint, we check whether the difference between the resulting joint configuration and the original one is smaller than 23° . We further check whether each finger, that was in contact with the object before, still contacts the object in at least one contact point close to the center of the finger (see Fig. 4). This detects situations where a finger is still in contact with the object but only touches its surface slightly with the side of the finger. If one of the checks fails, the grasp quality is decreased to 30% of its original computed value. To keep the grasp planner's objective computationally efficient and deterministic, we test six translations Δt_i , each moving the object along one axis for $2 \,\mathrm{cm}$ in both directions. In summary, grasps that are unstable when there is object pose uncertainty are devalued and therefore, such grasps are prevented by the grasp planner, if possible.

B. Ambiguities in predicting joint configurations

As mentioned in the introduction of Section IV, the second stage of the grasp prediction network, is trained to predict for a given object o and hand pose h the joint configuration q with the highest resulting grasp quality. This mapping is ambiguous, however, meaning there are usually multiple valid joint configurations that lead to the same or very similar grasp qualities. When using standard supervised training, this leads to learning the mean of all valid joint configurations,



Fig. 5. Adapted joint predictor network architecture with multiple heads. Each head predicts one configuration together with a logit l_k , which is used to determine which head to use at inference.

which, however, might not be valid. The way this is handled by Winkelbauer et al. [18], is by explicitly having multiple ground truth annotations for each training sample and then using the one closest to the network prediction for computing the loss. The disadvantage of this approach is that it is computationally expensive to find this set of joint configurations. Also, this might lead to including joint configurations with less quality and consequentially to the network learning not the top grasp.

We here present an alternative solution to the ambiguity problem, which is based on adapting the network architecture as shown in Figure 5. Instead of predicting just one configuration, we let the network predict N configurations using N heads. Each head k predicts a joint configuration q_k , a small hand pose correction Δh_k , and a grasp quality s_k . To compute the loss, we only use the prediction of the head j that is closest to the ground truth \hat{q} : $j = \arg \min_k ||\hat{q} - q_k||^2$. The joint loss term is hence defined as

$$L_{\text{joint}} = \left\| \widehat{q} - q_j \right\|^2 \,. \tag{1}$$

The loss terms for the predicted grasp quality s_j and hand pose correction Δh_j are formulated in an analog manner. To decide which head's prediction to use at inference time, we additionally model the relative frequency of each head's prediction in the training dataset. In detail, each head k additionally predicts a logit l_k , which are together transformed into a probability distribution using softmax: $c_k = \frac{e^{l_k}}{\sum_{i=1}^{K} e^{l_j}}$ Using the cross entropy loss function, this distribution is optimized to model the true underlying distribution of the predicted modes in the training dataset:

$$L_{\text{class}} = -\log c_j \tag{2}$$

At inference time, we select the prediction of the head k with the largest l_k , which corresponds to taking the prediction that is the most prominent in the training dataset. Using this architectural adaptation, it is possible for the network to handle ambiguities in the dataset, while only one ground truth label for each training sample is necessary. In this work, we are using N = 5 heads, which we found to be enough to cover the main ambiguities contained in our training data.

V. EXPERIMENTS

We conduct a systematic evaluation of our shape completion pipeline through multiple experiments on synthetic and real test data. First, we quantitatively assess the generalization capabilities to novel objects using volumetric and surface metrics on the completed shapes from simulated depth images. Next, we validate performance on real partial scans from a Kinect sensor and evaluate the ability to transfer from synthetic training data. Finally, we provide a qualitative grasp analysis by attempting robotic grasping of the completed shapes from real partial views. These experiments analyze the utility of our approach for enabling robotic manipulation of unknown objects based on incomplete observations from commodity RGB-D sensors. The robust performance demonstrates accurate shape inference on synthetic and real partial inputs, along with successful grasping by leveraging the completed geometry.

A. Shape completion

The performance of shape completion is evaluated on both synthetic and real test data. First, we assess the ability to generalize to novel objects not seen during training. We use 29 objects which were selected and scanned with a laser scanner for an Automatica trade fair demonstration as well as 38 objects of the YCB [48] dataset, for which we render 100 simulated Kinect depth images per object, each covering different viewpoints. We quantitatively evaluate the completed shapes using volumetric metrics like intersection over union (IoU), F1-score, precision, and recall computed by sampling one million spatial points uniformly inside the unit cube (Table I). Additionally, we extract surface meshes for each completion and measure point-based versions of F1, precision, recall, as well as Chamfer (L1) distance [26, 49] between points sampled from the surface of the ground truth and generated meshes (II).

TABLE I Volumetric results on Automatica/YCB dataset

Model	IoU↑	F1↑	Precision↑	Recall ↑
KINECT FINETUNE	66.7	75.7	73.5	83.3
KINECT SCALE	60.4	71	75.5	73.6
KINECT	58	68.9	75.7	70.6
BASIC	49.5	61.1	74	59.5

Table I shows the volumetric evaluation results on the Automatica/YCB dataset for different variants of our shape completion model. We observe that fine-tuning the network (KINECT FINETUNE) as explained in section III-B leads to the best performance with an IoU of 66.7% and F1-score of 75.7%. The network trained on Kinect data but without fine-tuning (KINECT) performs worse, justifying the additional step. Interestingly, the network trained on simulated Kinect data with additional random scaling augmentation (KINECT SCALE) leads to considerable improvement. Finally, the network trained only on synthetic data without simulation (BASIC) performs significantly worse, with an IoU of just 49.5%, highlighting the domain gap between synthetic and

real data. Overall, the results validate our design choices of using Kinect simulation, scaling augmentation, and finetuning, which enable accurate shape completion from real partial scans.

TABLE II Mesh surface results on Automatica/YCB dataset. Chamfer-L1 (CD)×10.

Model	CD↓	F1↑	Precision↑	Recall↑
KINECT FINETUNE	0.258	41.5	43.4	40.2
KINECT SCALE	0.276	43.4	44.6	42.8
KINECT	0.29	42.7	44.1	42
BASIC	0.455	34.1	36.8	32.6

Table II shows the mesh-based metrics on the Automatica/YCB dataset. We again observe the benefit of Kinect pretraining and finetuning, with KINECT FINETUNE achieving the lowest Chamfer-L1 distance of 0.258. The network trained only on synthetic data (BASIC) has significantly higher CD, revealing the domain gap. Overall, the mesh metrics confirm that leveraging simulated Kinect data along with finetuning is crucial for accurate shape completion from real partial observations, enabling tasks like robotic grasping.

Second, we validate performance on real sensor data from a Kinect camera (Table III and Figure 8). We use a limited subset of 25 objects and a single view. To align the ground truth meshes with the real data, we use a global-to-local, coarse-to-fine pose estimation pipeline¹ with manual visual inspection of the result. Note that this is only necessary for quantitative evaluation of the results and *not* during the inference phase when deploying the pipeline on the robot.

We separate the object points from the complete point cloud obtained from the depth image projection through plane segmentation. We first fit a plane into the scene corresponding to the table surface and then only keep points that are above. Due to noise, this procedure removes points close to the surface, like the bottom of bowls. To prevent this, we compute the convex hull of the object points, slightly increase its size, and translate it towards the table. We then use it as a mask to cut points from the table surface and reintroduce them to the object points. The reconstructed shapes are evaluated using the same set of volumetric and surface metrics to quantify the sim-to-real transferability.

TABLE III Results on real scans from Automatica/YCB objects. Chamfer-L1 (CD) $\times 10.$

Model	CD↓	IoU↑	F1↑	Precision↑	Recall ↑
KINECT FINETUNE	0.134	74.7	83	79.2	89.2
KINECT SCALE	0.147	72	81.1	79.8	84.2
KINECT	0.128	74.2	82.5	81.3	85.3
BASIC	0.143	71.2	80.4	82.3	80.7

Table III shows the evaluation on real test scans from Automatica/YCB objects. We observe that the models on

¹https://github.com/hummat/easy-o3d

simulated Kinect data achieve lower Chamfer distances and higher volumetric scores compared to the basic training, consolidating the results from the synthetic data evaluation and the qualitative results, showing failure cases without Kinect simulation (Figure 6) and without finetuning (Figure 7) and high-fidelity completions for a diverse set of objects and viewpoints (Figure 8). The KINECT model performs best in terms of Chamfer distance but worse than the finetuned model on the volumetric metrics. The scale-augmented Kinect model performs worst in this setting. Due to the small sample size and the resulting high variance–all methods show a standard deviation of around 0.05 for Chamfer distance and 2.4% (absolute) for IoU–further investigation is needed to differentiate between the different Kinect methods.



(a) Basic depth rendering (b) With Kinect simulation

Fig. 6. Effects of Kinect depth simulation. Training on unrealistic depth data leads to failure in challenging regions far away from input data and training distribution (left). Accurately simulating sensor characteristics resolves this (right).



Fig. 7. Effects of finetuning. The standard training methodology can lead to failure cases on challenging objects and viewpoints. Automated finetuning on difficult examples through importance sampling alleviates or greatly reduces this problem.

B. Grasping

Using the adapted grasp planner described in Section IV-A, we generate a dataset of 190 000 grasps across 12 000 ShapeNet objects [37]. For the training of the second stage, we generate an additional 100 000 grasps using hand poses generated by the trained first stage. This makes sure the second stage can correctly process the hand poses produced by the generative network. For each ShapeNet object in the training dataset, we generate a simulated Kinect depth image, based on which the full shape is again reconstructed using the shape completion network KINECT FINETUNE, presented in the last section. The two-stage grasp network, including the improvements described in Section IV-B, is now trained



Fig. 8. Qualitative results on the Automatica/YCB dataset. Our approach yields detailed object geometry from partial, noisy inputs. The complete meshes correspond to our KINECT FINETUNE model and its quantitative results in table III.

using the generated shape completions as input and the optimized stable grasps as labels.

We evaluate the grasping stage in simulation on the same dataset as used in Section V-A. Each object is grasped 10 times, each time rotated randomly around the up-axis. Of all tested 660 grasps, 95.2% were successful. The most prominent failure case is the object being too small to be grasped while keeping a distance to the table. Same as with the shape completion network, all evaluated objects shown here or in the video accompanying the paper are unknown to the grasping network and not part of the training dataset.

The full pipeline of predicting 1000 grasps for a given object takes only about 1 s with shape completion 0.7 s and grasp prediction 0.3 s^{-2} .

As visible in the video, even for unusual objects (e.g., the stuffed dog), the completion only gradually breaks down, leading to locally correct completions that still allow for a successful grasp.

In the following subsections, the two adaptations proposed to the grasp network are examined more closely:

1) Object pose uncertainty: When not taking the object pose uncertainty into account, fingers are sometimes placed close to or even on object edges, as they usually increase the grasp quality. However, when performing the grasp on the real system, such placed fingers might slide over the edge and involuntary rotate the object or even break the grasp Using the changes proposed in Section IV-A, the grasps in our training dataset and, therefore, also the grasps predicted by the network, are chosen in a way that fingers are placed further from edges, if possible. This can be seen in



Fig. 9. Comparing grasp predictions with and without considering object pose uncertainty during training data generation: When taking pose uncertainty into account, it leads to fingers being placed further away from edges.

Fig. 9. The network trained without considering object pose uncertainty is trying to fit as many fingers as possible on the small surface of the cup, leading to one finger being placed on the edge of the cup The network trained with object pose uncertainty, however, instead only uses three fingers, which allows them to be placed at sufficient distance to the object edges. So even with slight variations in the object pose, the grasp would stay stable.

2) Joint configuration ambiguity: By applying the architectural changes described in Section IV-B to the grasping network, we are able to train our network using only one label for each training sample. When training the model architecture used in Winkelbauer et al. [18] on the same data, ambiguities cannot be distinguished and the network learns the mean of all valid joint configurations. This becomes especially clear when there exist distinct modes in the training labels. Fig. 10 shows an example where the object can be grasped with three fingers or four fingers. While the multi-output network can predict separate valid modes, the single-output network predicts a mix between the modes leading to one finger intersecting the object. This example also shows the importance of classifying which predicted modes are active. For this kind of object, there seem to exist two main types of grasps in the training dataset. One using four fingers and one using only three. The network predicts both modes, each having a large frequency classification, while the other three outputs are not used here and therefore are classified as close to 0%.

VI. CONCLUSION

In this paper, we have proposed a novel, complete, and fast pipeline (only 1s from looking to grasping) for robotic grasping of unknown objects with a multi-fingered hand based only on partial shape observations from a commodity depth sensor. Our method combines deep implicit shape completion (0.7 s) with data-driven multi-finger grasp prediction (0.3 s for 1000 grasps). This enables efficient grasping across objects with significant shape variation, a key capability for robotic manipulation in unconstrained real-world environments. Through systematic experiments on synthetic and real data, we have validated the benefits of our training procedures for improved sim-to-real transfer. Robotic grasping trials further demonstrate the utility of our approach by successfully manipulating a diverse range of novel household objects based on single-view depth images.

²Intel Xeon 6144 with 16 cores and a single GPU NVIDIA V100.



Fig. 10. Comparing single-output network (left) with multi-output network (right) on predicting the grasp for a butter box: The prediction of the singleoutput network is a mixture of all modes leading to one finger intersecting the object. The prediction of the multi-output network can explicitly consider several modes. In this case, two modes are used: The second one using all four fingers and the fourth one using only three.

Future work will include further improvements to sensor simulation, network architectures and training procedures.

REFERENCES

- N. Curtis and J. Xiao, "Efficient and effective grasping of novel objects through learning and adapting a knowledge base," in *IEEE International Conference on Intelligent Robots and Systems*, 2008.
- [2] C. Goldfeder, M. Ciocarlie, H. Dang, and P. Allen, "The Columbia grasp database," in *IEEE International Conference on Robotics and Automation*, 2009.
- [3] D. Song, C. Ek, K. Huebner, and D. Kragic, "Ebodiment-specific representation of robot grasping using graphical models and latent- space discretization," in *IEEE International Conference on Intelligent Robots and Systems*, 2011.
- [4] M. Ciocarlie and P. Allen, "Hand posture subspaces for dexterous robotic grasping," Int. J. Robotics Research, vol. 28, no. 7, pp. 851–867, 2009.
- [5] U. Hillenbrand and M. A. Roa, "Transferring functional grasps through contact warping and local replanning," in *IEEE International Conference on Intelligent Robots and Systems*, 2012.
- [6] T. Stouraitis, U. Hillenbrand, and M. A. Roa, "Functional power grasps transferred through warping and replanning," in *IEEE International Conference on Robotics and Automation*, 2015.
- [7] M. Sundermeyer, A. Mousavian, R. Triebel, and D. Fox, "Contact-GraspNet: Efficient 6-DoF grasp generation in cluttered scenes," in *IEEE International Conference on Robotics and Automation*, 2021.
- [8] B. Bäuml et al., "Agile justin: An upgraded member of DLR's family of lightweight and torque controlled humanoids," in Proc. IEEE International Conference on Robotics and Automation, 2014.
- [9] J. Tenhumberg and B. Bäuml, "Speeding up optimization-based motion planning through deep learning," in *Proc. Int. Conf. Intelligent Robots and Systems*, 2022.
- [10] —, "Efficiently learning the inverse kinematics with obstacles," in Proc. IEEE-RAS International Conference on Humanoid Robots, 2023.
- [11] —, "Self-contained calibration of an elastic humanoid upper body using only a head-mounted rgb camera," in *IEEE-RAS International Conference on Humanoid Robots*, 2022.
- [12] J. Tenhumberg, L. Sievers, and B. Bäuml, "Self-contained and automatic calibration of a multi-fingered hand
- using only pairwise contact measurements," in Proc. IEEE-RAS International Conference on Humanoid Robots, 2023.
- [13] B. Wu *et al.*, "Generative Attention Learning: a "GenerAL" framework for highperformance multi-fingered grasping in clutter," *Autonomous Robots*, vol. 44, no. 6, pp. 971–990, 2020.
- [14] L. Sievers, J. Pitz, and B. Bäuml, "Learning purely tactile in-hand manipulation with a torque-controlled hand," in *IEEE International Conference on Robotics* and Automation, 2022.
- [15] R. Wu *et al.*, "Functional grasp transfer across a category of objects from only one labeled instance," *IEEE Robotics and Automation Letters*, vol. 8, no. 5, pp. 2748–2755, 2023.
- [16] X. Yan et al., "Shapeformer: Transformer-based shape completion via sparse representation," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2022.
- [17] W. E. Lorensen and H. E. Cline, "Marching cubes: A high resolution 3d surface construction algorithm," in ACM SIGGRAPH Computer Graphics, 1987.
- [18] D. Winkelbauer et al., "A two-stage learning architecture that generates highquality grasps for a multi-fingered hand," in *IEEE International Conference on Intelligent Robots and Systems*, 2022.
- [19] Z. Wu et al., "3d shapenets: A deep representation for volumetric shapes," in IEEE Conference on Computer Vision and Pattern Recognition, 2014.
- [20] A. Dai, C. Qi, and M. Nießner, "Shape completion using 3d-encoder-predictor cnns and shape synthesis," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
- [21] W. Yuan et al., "Pcn: Point completion network," in International Conference on 3D Vision, 2018.
- [22] P. Xiang et al., "Snowflakenet: Point cloud completion by snowflake point deconvolution with skip-transformer," in *IEEE International Conference on Computer Vision*, 2021.
- [23] X. Yu et al., "Pointr: Diverse point cloud completion with geometry-aware transformers," in *IEEE International Conference on Computer Vision*, 2021.

- [24] T. Groueix et al., "AtlasNet: A Papier-Mâché Approach to Learning 3D Surface Generation," in IEEE Conference on Computer Vision and Pattern Recognition, 2018.
- [25] Y. Liao, S. Donné, and A. Geiger, "Deep marching cubes: Learning explicit surface representations," 2018.
- [26] L. M. Mescheder et al., "Occupancy networks: Learning 3d reconstruction in function space," in *IEEE Conference on Computer Vision and Pattern Recogni*tion, 2018.
- [27] S. Peng et al., "Convolutional occupancy networks," in European Conference on Computer Vision, 2020.
- [28] J. Chibane, T. Alldieck, and G. Pons-Moll, "Implicit functions in feature space for 3d shape reconstruction and completion," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2020.
- [29] Q. Xu et al., "Disn: Deep implicit surface network for high-quality single-view 3d reconstruction," in Advances in Neural Information Processing Systems 32, 2019.
- [30] J. J. Park et al., "Deepsdf: Learning continuous signed distance functions for shape representation," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2019.
- [31] J. Varley et al., "Shape completion enabled robotic grasping," in IEEE International Conference on Intelligent Robots and Systems, 2016.
- [32] X. Yan et al., "Learning 6-dof grasping interaction via deep geometry-aware 3d representations," in IEEE International Conference on Robotics and Automation, 2017.
- [33] M. V. der Merwe et al., "Learning continuous 3d reconstructions for geometrically aware grasping," in *IEEE International Conference on Robotics and Automation*, 2019.
- [34] W. Chen et al., "Improving object grasp performance via transformer-based sparse shape completion," Journal of Intelligent & Robotic Systems, vol. 104, pp. 1–14, 2022.
- [35] M. Liu et al., "Generating Grasp Poses for a High-DOF Gripper Using Neural Networks," in *IEEE International Conference on Intelligent Robots and Systems*, 2019.
- [36] —, "Deep differentiable grasp planner for high-dof grippers," in *Robotics: Science and Systems*, 2020.
- [37] A. X. Chang et al., "Shapenet: An information-rich 3d model repository," ArXiv, vol. abs/1512.03012, 2015.
- [38] D. Stutz and A. Geiger, "Learning 3d shape completion under weak supervision," International Journal of Computer Vision, vol. 128, pp. 1162–1181, 2020.
- [39] P. Cignoni et al., "MeshLab: an Open-Source Mesh Processing Tool," in Eurographics Italian Chapter Conference, 2008.
- [40] M. Gschwandtner, R. Kwitt, A. Uhl, and W. Pree, "Blensor: Blender sensor simulation toolbox," in *International Symposium on Visual Computing*, 2011.
- [41] J. Bohg, J. Romero, A. Herzog, and S. Schaal, "Robot arm pose estimation through pixel-wise part classification," in *IEEE International Conference on Robotics and Automation*, 2014.
- [42] M. Humt, D. Winkelbauer, and U. Hillenbrand, "Shape completion with prediction of uncertain regions," in *IEEE International Conference on Intelligent Robots and Systems*, 2023.
- [43] A. Paszke et al., "Pytorch: An imperative style, high-performance deep learning library," in *Neural Information Processing Systems*, 2019.
- [44] W. Falcon and The PyTorch Lightning team, "PyTorch Lightning," Mar. 2019. [Online]. Available: https://github.com/Lightning-AI/lightning
- [45] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," 2014.
- [46] C. Ferrari and J. Canny, "Planning optimal grasps," in *IEEE International Conference on Robotics and Automation*, 1992.
- [47] J. Butterfaß, M. Grebenstein, H. Liu, and G. Hirzinger, "DLR-Hand II: Next generation of a dextrous robot hand," in *IEEE International Conference on Robotics and Automation*, 2001.
- [48] B. Calli et al., "Benchmarking in manipulation research: Using the yale-cmuberkeley object and model set," *IEEE Robotics Automation Magazine*, vol. 22, no. 3, pp. 36–52, 2015.
- [49] M. Tatarchenko et al., "What do single-view 3d reconstruction networks learn?" in IEEE Conference on Computer Vision and Pattern Recognition, 2019.