

# Shape Completion with Prediction of Uncertain Regions

Matthias Humt<sup>1,2</sup>, Dominik Winkelbauer<sup>1,2</sup>, and Ulrich Hillenbrand<sup>1</sup>

**Abstract**—Shape completion, i.e., predicting the complete geometry of an object from a partial observation, is highly relevant for several downstream tasks, most notably robotic manipulation. When basing planning or prediction of real grasps on object shape reconstruction, an indication of severe geometric uncertainty is indispensable. In particular, there can be an irreducible uncertainty in extended regions about the presence of entire object parts when given ambiguous object views. To treat this important case, we propose two novel methods for predicting such uncertain regions as straightforward extensions of any method for predicting local spatial occupancy, one through postprocessing occupancy scores, the other through direct prediction of an uncertainty indicator. We compare these methods together with two known approaches to probabilistic shape completion. Moreover, we generate a dataset, derived from ShapeNet [1], of realistically rendered depth images of object views with ground-truth annotations for the uncertain regions. We train on this dataset and test each method in shape completion and prediction of uncertain regions for known and novel object instances and on synthetic and real data. While direct uncertainty prediction is by far the most accurate in the segmentation of uncertain regions, both novel methods outperform the two baselines in shape completion and uncertain region prediction, and avoiding the predicted uncertain regions increases the quality of grasps for all tested methods. Web: <https://github.com/DLR-RM/shape-completion>

## I. INTRODUCTION

Object grasping and manipulation is the primary task of many robotic systems. As a challenging example, consider an assistance robot operating in human living environments with various objects (household objects, tools) that cannot all be known in advance. A geometric reconstruction would be a good starting point for planning all kinds of interaction. However, acquiring a full view of an object is often not practical or even impossible due to unavoidable occlusions by other objects.

The present work concerns a single view of an unknown object instance of a known category, which, in practice, may be known from an object detector. We investigate shape completion and its adequacy for predicting stable grasps on the object. This has been the subject of study in several works [2, 3, 4]. Unlike the previous works, however, we focus on treating significant uncertainties that arise for objects with partial symmetry. Specifically, there can be identical views for objects with partial symmetry for various object poses as observed in [5]. Thus a unique object pose is not implied for such views, and the position of occluded object parts may therefore vary within a wide spatial range.

Examples of such objects from the household domain are mugs, jugs, and pitchers with partial rotational symmetry and a handle as the symmetry breaker. In fact, an extended

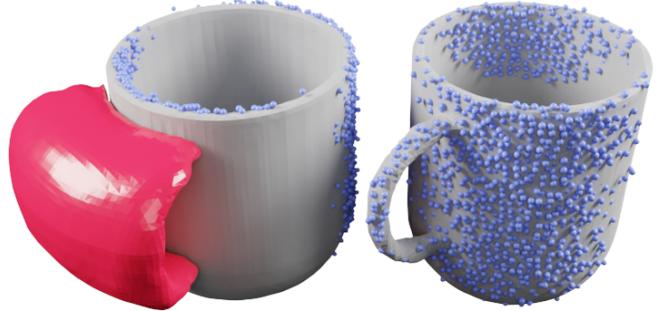


Fig. 1. Shape completion of a mug. Left: If the handle is occluded from the camera view, we predict the uncertain region (red) that contains the handle, resulting from pose ambiguity. The mug is reconstructed in the region not affected by pose ambiguity (gray). Right: If any part of the handle is visible, also the handle is reconstructed, and no uncertain region is predicted.

region may contain an occluded handle and should hence be avoided by the robot gripper; see Fig. 1 for an example. Note that this kind of uncertainty exists even for a fully known object shape or a perfect shape predictor that would provide an accurate reconstruction from another viewpoint. This uncertainty arises from the object shape and the given viewpoint, not from any model uncertainty of the predictor, that is, epistemic uncertainty. In this sense, it is an *objective uncertainty*.

In this work, we consider the representative example of mug-like objects. We propose two different methods for predicting the uncertain spatial regions that result from pose ambiguity. These methods are used here as an extension of [6] for shape completion but can be combined with any completion method that predicts spatial occupancy scores (probabilistic or otherwise).

We conduct an extensive set of experiments on known and novel object instances and both on synthetic and real data. We compare our approaches to two existing methods for probabilistic shape completion that, after little adaptation, lend themselves to predicting viewpoint-induced uncertain regions [7, 5]. Our proposed methods outperform the two baselines on a range of metrics for occupied and uncertain region prediction and improve grasp quality by avoiding the predicted uncertain regions.

We make the following contributions.

- Introduction of two novel methods for predicting uncertain regions in shape completion as extensions of any predictor of spatial occupancy scores.
- Validation and comparison of these new methods and two baselines on the task of uncertainty prediction in an extensive set of experiments, showing superior performance of the new methods.

<sup>1</sup> The author is with the DLR Institute of Robotics and Mechatronics.

<sup>2</sup> The author is with the Technical University of Munich (TUM).

Contact: [matthias.humt@dlr.de](mailto:matthias.humt@dlr.de)

- In these experiments, demonstration of the utility of the predicted uncertain regions as a constraint for high-quality grasps prediction.
- A dataset of realistically rendered object views, derived from ShapeNet [1], including ground-truth annotation of the corresponding uncertain regions.

## II. RELATED WORK

### A. Shape Completion

In the context of shape completion, i.e., reconstructing the entire 3D object shape from a partial 3D observation, learning architectures that can deal with unordered 3D point clouds [8, 9, 10] are of great interest, as this is the modality directly obtained from depth sensors. Point clouds can easily be voxelized in different ways and then processed with voxel-based methods [11, 12, 7, 5]. The advantage is that proven techniques from CNNs can be used in the ordered grid, but at the cost of high memory when high resolution is required.

Apart from the input modality, the shape representation generated by a shape completion process can make a difference for downstream tasks. In particular, for the present context of robotic grasp generation, triangle surface meshes are the preferred choice. Some methods directly produce meshes as output [13, 14]. But the task is challenging because of the complex mesh structure with varying topology, and hence the generated meshes often need postprocessing to be usable.

More recently, the representation of shapes through an implicit function has emerged as a new paradigm in the shape completion community [15, 16, 6, 17]. Instead of predicting the shape of the object explicitly, the problem is cast as a point-wise classification problem where the object surface is implicitly defined as the decision boundary between its interior and exterior [15, 6, 17]. It can also be formulated as a regression problem where the task is to predict the signed distance from the surface [16] for each query point.

A significant advantage of the implicit representation is its ability to be queried at continuous locations, e.g. in a grid of arbitrary resolution to extract a detailed surface mesh. Further, it can be flexibly combined with encoders to match the input data type, such as a simple PointNet [18] to generate complete meshes from partial point clouds. This is the desired combination for our use case.

Our work is further concerned with shape completion as a basis for geometrically informed robotic grasping, which has recently attracted wider attention [2, 7, 5, 3, 4]. Contrary to previous work however, we focus on the specific problem of extended regions of objective uncertainty caused by viewpoint-dependent ambiguity in the input data.

While [7] are also concerned with uncertainty in shape completion in robotic grasping, their focus lies on epistemic or model uncertainty. More related to our work, [5] focus on ambiguity arising from partial observations and the generation of diverse yet plausible shape completions. But they omit to extract the resulting regions of uncertainty and to directly use them as a constraint for grasp generation, which we propose to do in this work. We quantitatively compare our proposed methods to these two prior works, adapting

them to the here-considered problem of predicting extended uncertain regions.

### B. Uncertainty Quantification in Pose Estimation

The studied uncertainty for shape completion arises from uncertainty in the objects pose for partially symmetric objects. The quantification of pose uncertainty has also been addressed in some works on object pose estimation.

The uncertainty in an estimate can be quantified as the level of confidence in the prediction, sometimes called predictive uncertainty [19, 20]. A more expressive representation of pose uncertainty is predicting a probability density of poses. There has been work on modeling rotational uncertainty with parametric densities and their mixtures [21, 22]. The density parameters and mixture weights are predicted. However, to determine the spatial effect of pose uncertainty, a sufficient sample of poses would need to be drawn from the predicted densities, and for each pose, the object model needs to be transformed into the scene for mapping of the resulting occupancy. This extra effort may easily prevent real-time performance.

An interesting alternative to a parametric density model is through a deep neural network (DNN) that predicts the probability for a given image and every queried object orientation [23]. Another alternative is density representation through a particle filter for pose tracking [24]. However, the effort of sampling many poses and transforming object models into the scene for each of them remains.

The above approaches mainly target known objects; no uncertainty quantification under shape variations is studied. In the present work, we aim at the generalization across shapes.

## III. METHOD

This section describes the generation of realistic synthetic training data and training conditions required to generalize to real-world data. We then briefly describe the data representation and network architecture to predict complete shapes from partial inputs. Finally, a method for finding ambiguous views leading to uncertain regions in the shape completion is presented, and two methods for the prediction of these uncertain regions are introduced.

### A. Realistic Synthetic Data

Point clouds projected from depth images are the most natural representation of real-world 3D data but come with unique challenges beyond permutation equivariance and inhomogeneous density. Self-occlusion only allows for partial views of captured objects, while sensor characteristics and noise remove additional information.

While these effects are captured naturally for data collected with sensors in the real world, this data acquisition process is extremely time-consuming and does not scale to a diverse and potentially changing set of objects and scenarios.

We, therefore, resort to synthetic data generation through rendering from various viewpoints and subsequent online augmentation with noise. As sensors experience stronger

artifacts on surface regions perpendicular to the viewing direction, we use the cosine similarity with the per-point normal direction as the probability of removing points and adding additional noise.

### B. Realistic Training Conditions

The camera pose relative to fixed world coordinates, i.e., the *object* coordinate frame, is frequently unknown on a (mobile) robotic system. For this reason, shape completion approaches geared towards real-world application are usually trained in *camera* coordinates, which can yield better generalization to unknown shapes, albeit at the cost of reduced fidelity on known examples [25].

When working with a robotic system, we identify a third alternative to camera and object coordinates: the *robot world* coordinates. Through the robot’s kinematic chain, which is usually known, the camera pose can be projected to this world frame, removing the effects of camera rotation, thus reducing the complexity of the learning problem to object pose variations.

### C. Implicit Function Learning

We briefly summarize the shape completion problem as implicit function learning following [15] which we make use of in our first approach and extend for the second one. Given an observation in the form of a partial point cloud  $x \in \mathcal{X}$  and an arbitrary 3D point  $p \in \mathbb{R}^3$ , we want to find a parametric function  $f_\theta : \mathbb{R}^3 \times \mathcal{X} \rightarrow [0, 1]$  where  $f_\theta$  is a DNN, and its output represents the occupancy probability at  $p$ . The task of learning the parameters of the model, therefore, reduces to the standard binary classification problem which is trained using the binary cross-entropy cost function,

Evaluating  $f_\theta$  at points on a 3D grid, we extract a mesh from the resulting occupancy grid as the approximate isosurface at threshold  $\tau$  with  $f_\theta(p, x) = \tau$  using the *Marching Cubes* [26] algorithm. At  $\tau = 0.5$ , the extracted isosurface represents the decision boundary between the two classes (*free* and *occupied*), defining the surface of the object implicitly.

### D. Uncertain Shape Completion

Quantifying uncertainty in deep learning is of great interest [27], especially if the model is deployed in the real world where wrong predictions made with high certainty, i.e., overconfident predictions, can have grave consequences.

The predictive uncertainty of a model is the consequence of epistemic and aleatoric uncertainty. Dealing with objects in 3D space, additional *pose uncertainty* can arise. Given an ambiguous object view, this form of uncertainty remains even under conditions of perfect knowledge regarding model parameters and object geometry: it is irreducible and objective. This effect can be observed for objects with partial symmetry (continuous or discrete) broken only from specific viewpoints.

Mugs are a prominent real-world example. When the handle is hidden from view by the mug’s body, many possible orientations of the mug yield the same observation. As a

result, there is an extended region that may contain the mugs handle, which we call *uncertain region*. Crucially, however, this uncertainty is precisely defined by all rigid transformations  $T \in \text{SE}(3)$  that yield the same observation.

Similar to [5], we estimate this set of transformations, and thereby the region in  $\mathbb{R}^3$  affected by pose uncertainty by randomly transforming the object and comparing the resulting 2D projection to the initial one. An identical or, in practice, similar view that places the handle at a different position enlarges the uncertain region.

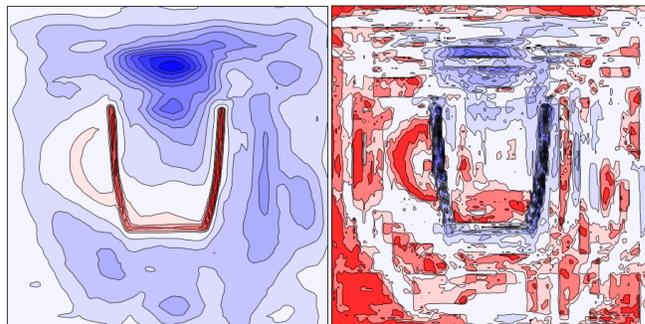
Given this ground truth information about uncertain regions for each observation, we extend the binary classification task of implicit function learning with classes *free* and *occupied* by a third class *uncertain*. We then train a new model  $f_\theta : \mathbb{R}^3 \times \mathcal{X} \rightarrow [0, 1, 2]$ .

While this direct approach is intuitive and promising, it requires the costly identification of ground truth uncertain regions for model training. A second approach that does not need extra annotation relies on the per-class score to measure the uncertainty in a predicted local occupancy.

To extract regions of low certainty from the occupancy grid, we can set  $\tau < 0.5$ . Unfortunately, each isosurface at  $\tau_1$  is entirely surrounded by another with  $\tau_2 < \tau_1$ . Thus, we need a way to differentiate between narrow uncertain regions necessarily occurring near the object’s surface and more extended ones reaching farther out.

One such criterion is  $|\nabla_p \hat{y}|$ , the magnitude of the gradient of the occupancy probability at a query point  $p$  with  $\hat{y} := p(y = 1|x, p, \theta)$  being the predicted occupancy probability at point  $p$  conditioned on the observation and the model parameters. Being able to obtain such gradients is a unique and valuable property of implicit function models [15].

We expect this gradient to be large near the decision boundary defining the object’s surface, whereas it should be small for points farther away. Using an upper threshold on the norm of the gradient for each point in the occupancy grid, we extract approximate uncertain regions from a trained model, as shown in Fig. 2, without relying on ground truth *uncertain* labels.



(a) Predicted occupancy probability (b) Gradient of (a)

Fig. 2. Slice of a side view from the predicted occupancy probability grid (a). Using a small lower threshold, a region possibly containing a handle appears behind the mug (light red), but the mug itself is also contained (dark red). Using the gradient of the predicted occupancy probability (b) with its average as upper threshold, this unwanted region can be discarded by only considering the intersection of regions shown in red as uncertain.

While the approach with the gradient criterion does not need extra data annotation, it does require a proper adjustment of the gradient threshold, in the general setting covering all possible uncertain region sizes with their gradients.

Another possible criterion is the variance of multiple stochastic forward passes using *Monte Carlo Dropout* [28] as employed by [7].

Similarly, we are also able to extract an uncertain region from [5] through the variance of multiple samples from the Variational Autoencoder (VAE).

For all methods, a low threshold  $\tau$  extracts a larger uncertain region but often introduces artifacts in the form of small free-floating pieces of geometry. As we want to extract large continuous areas, we cluster the connected components and remove those too small.

#### IV. EXPERIMENTS

We conduct three experiments to evaluate (i) the generalization ability to novel views of known objects, (ii) the ability to generalize to novel instances of the same class, and (iii) the generalization from synthetic to real data (sim2real). The evaluation is done regarding shape completion, prediction of uncertain regions, and its effect on the robotic task of object grasping. We compare the two newly proposed methods, i.e., the *binary* with a gradient criterion and the *trinary*, against each other and the two adapted baselines, i.e., *dropout* and VAE, described in the previous section.

##### A. Datasets

We render 200k train, 20k validation, and 40k test depth and normal maps offline using *BlenderProc2* [29] and then augment them online during training as shown in Algorithm 1 and explained in the previous section.

---

##### Algorithm 1 Generate datum

---

- 1: Sample camera position  $c$  randomly in upper hemisphere
  - 2: Scale object  $\mathcal{U}(0.05, 0.15)$  and  $\pm 20\%$  along z-axis
  - 3: Render depth and normal image
  - 4: Project depth, normals to point cloud  $\mathcal{P} = \{(p_i, n_i)\}_{i=0}^N$
  - 5: **for**  $p, n \in \mathcal{P}$  **do**
  - 6:    $p = p + \mathcal{N}(0, 0.005)$  // Add noise
  - 7:    $s = n \cdot c / \|n\| \|c\|$  // Compute cosine similarity
  - 8:    $\mathcal{P} \setminus p$  if  $|s| < \mathcal{U}(0, 1)$  // Remove point
  - 9:    $p = p + \mathcal{N}(0, 0.01)$  if  $|s| < \mathcal{U}(0, 1)$  // Add noise
  - 10: **end for**
  - 11: Transform point cloud from camera to robot world frame
  - 12: **return** augmented point cloud
- 

We consider two datasets for our experiments. For the novel view and novel instance experiments, we use the *mugs* category of the *ShapeNet* [1] dataset. Of the 214 available mugs, we select the 201 with a handle. Watertight meshes and query points are generated similarly to [15]. For the

novel instance experiment, we split the 201 objects into training (70%), validation (10%), and test (20%) sets.

The sim2real experiment is conducted on a subset of the datasets from the *BOP Challenge* [30, 31], which provide real-world RGB-D images with ground truth per-object pose, segmentation mask and mesh.

Five out of 12 datasets (*Linemod (LM)*, *HomebrewedDB (HB)*, *YCB-Video (YCB-V)*, *Toyota Light (TYOL)*) contain a total of nine different mugs. We replace the low quality *LM* mug mesh with a higher quality mesh of the featured *IKEA FÄRGRIK* mug and transform the *HB* mug from being half full to empty. HB provides RGB-D data from both the Primesense ( $HB_{pri}$ ) and Kinect ( $HB_{kin}$ ) sensors. We use both YCB-V sequences (48 and 55) in which the mug is present.

As the provided ground truth poses and/or camera intrinsic parameters are inaccurate, we perform a few steps of the *Iterative Closest Point* [32] algorithm between projected real and rendered depth maps to improve the ground truth pose.

We use the provided per-object segmentation masks to remove non-object points. Those masks are inaccurate, however, leading to unwanted pixels at the boundaries, which we remove by scaling the convex hull of the object mesh by a small factor and only keeping points on the inside. We only evaluate views with an object visibility of at least 85%.

##### B. Metrics

We evaluate the performance in region segmentation, both occupied and uncertain, using the standard volumetric metrics *IoU*, *Precision*, *Recall*, *F1-Score*. We further build three grasping-related metrics from the confusion matrix constituents.

$$GCR = (FN_{occ} + FN_{unc}) / (TP_{occ} + FN_{occ} + TP_{unc} + FN_{unc})$$

$$GMR = FP_{occ} / (FP_{occ} + TP_{occ})$$

$$GER = |\mathcal{FP}_{occ} \cup \mathcal{FP}_{unc}| / (FP_{occ} + TN_{occ})$$

Here FP, FN, TP, and TN are the usual true and false counts for classifying occupied (occ) and uncertain (unc) vs. free regions, respectively. Confusion between occupied and uncertain regions does not affect these metrics.  $\mathcal{FP}$  denotes the set of false positives to prevent accounting twice for the same location. The *Grasp Collision Risk* (GCR) is a straightforward measure of the probability of unexpected collision when closing the hand around a planned grasp region; *Grasp Miss Risk* (GMR) is for the opposite case when closing the hand, i.e., an expected contact is not made; *Grasp Exclusion Risk* (GER) quantifies the risk of rejecting feasible grasp candidates because they are blocked by false prediction of occupied or uncertain regions. All volume measures are quantified by counting uniform point samples inside the considered volumes.

Intuitively, precision is the fraction of correctly predicted occupied space, and recall the fraction of actual occupied space found by the model.

In addition to the volumetric quantities, we also include the *Chamfer-L1* distance (CD) [15] to measure shape difference and displacement between the surface meshes of the

ground-truth and predicted occupied and uncertain regions, respectively. The length scale is set by normalizing object size to fit in a unit cube.

To measure the effect of uncertain region prediction on object grasping, we use the network from [33] to predict 1024 grasps for a humanoid four-finger hand (DLR-Hand II [34]) per object from its completion. or the uncertainty-aware variant we use the predicted uncertain region as a filter, removing grasps that intersect with it. We evaluate the grasps’ quality using the *Improved Epsilon Quality* (IEQ) [33] metric, which measures the minimal external force applied to the object (ground truth mesh) that would break the grasp where the highest IEQ in the sample is the achieved grasp quality.

### C. Network, Training & Hyperparameters

We apply our proposed methods to the *Convolutional Occupancy Network* (ConvONet) [6], which strikes a good balance between performance and size.

We implement all methods in *Python* and *PyTorch* [35], building on the code base of [15, 6] for ConvONet and adapt [7] for our first baseline. For our second baseline, we convert the *TensorFlow* [36] implementation of [5] into *PyTorch*. We use the *Adam* [37] optimizer with default parameters to train the models until convergence or until their validation error starts to increase.

When extracting occupied and uncertain regions from a binary occupancy grid, two hyperparameters must be chosen: the threshold  $\tau$  on the occupancy score  $\hat{y}$  deciding whether a point lies on the interior of the shape and another threshold  $\tau_u < \tau$  determining whether a point location is regarded as uncertain. The canonical evaluation in binary classification setting requires  $\tau = 0.5$ , but we can also view it as a hyperparameter to be optimized on the validation set for maximum IoU [15], shown in Fig. 3. This can be seen as a form of post hoc calibration of the predicted probabilities similar to temperature scaling [38].

In the trinary case, the class prediction is usually the *argmax* of the predicted class scores. As in the binary case, however, we regard it as a hyperparameter and again optimize it on the validation set for maximum IoU to mitigate potential overconfidence of the DNN [38].

One additional hyperparameter must be chosen for the *binary*, *dropout* and *VAE* methods. For a location with  $\hat{y} < \tau_u$  to be considered part of the uncertain region, we threshold the magnitude of the gradient to be below average for the *binary* method and the variance of the samples from the *dropout* and *VAE* methods to be above average.

### D. Novel View Generalization

Table I shows the quantitative evaluation of all methods on the test set. The implicit function models (*binary* and *trinary*) perform similarly in predicting the occupied space and significantly outperform the voxel-based baselines. The *dropout* model performs worse than the *VAE*, which uses a higher-resolution voxel grid. Predicting the uncertain region

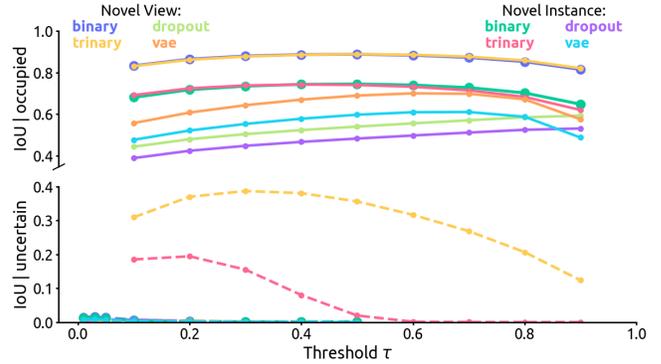


Fig. 3. Effect of varying the threshold parameter  $\tau$  on IoU for occupied and uncertain region predictions.

TABLE I. **Novel Views:** Quantitative shape completion results. All segmentation metrics are shown separately for the predicted occupied and uncertain regions. All grasp-related metrics are shown for the cases where the predicted uncertainty is ignored and considered. For metrics with  $\uparrow$  higher is better, for those with  $\downarrow$  lower is better. All volumetric measures are in %, CD is  $\times 100$ . Best values are in **bold**.

	trinary	binary	dropout	vae	trinary	binary	dropout	vae
	occupied region				uncertain region			
IoU $\uparrow$	78.83	<b>79.00</b>	54.92	63.07	<b>31.05</b>	9.32	4.92	4.86
F1 $\uparrow$	86.08	<b>86.29</b>	69.27	75.38	<b>49.53</b>	31.76	19.38	18.72
Prec. $\uparrow$	84.97	<b>85.59</b>	64.66	72.32	<b>44.63</b>	13.05	7.19	7.85
Rec. $\uparrow$	<b>90.92</b>	90.60	79.50	82.72	<b>49.51</b>	45.46	26.16	23.58
CD $\downarrow$	2.85	<b>2.77</b>	5.24	4.50	<b>15.77</b>	22.86	21.81	23.56
	uncertainty ignored				uncertainty considered			
IEQ $\uparrow$	<b>1.79</b>	1.64	1.66	1.66	<b>2.25</b>	2.05	2.20	1.73
GER $\downarrow$	2.48	<b>2.31</b>	4.89	3.78	<b>2.83</b>	4.73	7.93	6.27
GCR $\downarrow$	<b>18.07</b>	18.37	28.25	24.78	14.49	<b>13.60</b>	25.63	22.37
GMR $\downarrow$	18.06	<b>17.07</b>	33.16	26.77	18.06	<b>17.07</b>	33.16	26.77

TABLE II. **Novel Instances:** Metrics as for *Novel Views*; cf. Table I.

	trinary	binary	dropout	vae	trinary	binary	dropout	vae
	occupied region				uncertain region			
IoU $\uparrow$	69.24	<b>69.94</b>	52.52	58.02	<b>19.36</b>	6.56	4.93	5.46
F1 $\uparrow$	80.13	<b>80.74</b>	67.63	71.55	<b>35.51</b>	23.11	18.89	20.07
Prec. $\uparrow$	81.63	<b>83.68</b>	65.25	72.43	<b>32.04</b>	7.72	7.08	7.70
Rec. $\uparrow$	<b>82.91</b>	81.86	75.19	75.20	38.20	<b>52.07</b>	27.62	28.19
CD $\downarrow$	3.69	<b>3.64</b>	5.56	5.11	<b>17.05</b>	25.32	21.95	25.64
	uncertainty ignored				uncertainty considered			
IEQ $\uparrow$	1.86	<b>1.97</b>	1.62	1.71	2.30	<b>2.35</b>	2.07	1.94
GER $\downarrow$	2.57	<b>2.05</b>	4.48	3.58	<b>2.77</b>	7.12	8.05	6.92
GCR $\downarrow$	<b>26.70</b>	27.55	33.34	31.25	25.79	<b>21.98</b>	30.64	28.33
GMR $\downarrow$	19.85	<b>16.60</b>	32.14	26.87	19.85	<b>16.60</b>	32.14	26.87

is a more challenging task where the *trinary* model outperforms all other methods by a large margin, followed by the *binary* model.

The IEQ increases substantially across all methods when the predicted uncertain region is considered with the *trinary* model taking the lead. In this case, the GER also increases, but the *trinary* model is affected the least due to its ability to accurately predict the uncertain region. However, the *binary* model achieves the lowest GCR and GMR. Again, both baselines are outperformed across the grasp metrics, except

TABLE III. **Sim2Real**: The segmentation metrics are shown for the predicted occupied region. IEQ is shown when predicted uncertainty is ignored and considered. All volumetric measures are in %, CD is  $\times 100$ . Best mean values across the different test datasets are in **bold**.

	trinary							binary							dropout	vae
	HB <sub>pri</sub>	HB <sub>kin</sub>	LM	TYOL	YCBV <sub>48</sub>	YCBV <sub>55</sub>	mean	HB <sub>pri</sub>	HB <sub>kin</sub>	LM	TYOL	YCBV <sub>48</sub>	YCBV <sub>55</sub>	mean	mean	mean
<i>occupied region</i>																
IoU $\uparrow$	49.26	32.22	33.83	39.64	23.08	19.54	<b>32.93</b>	49.51	30.45	32.70	39.92	22.80	20.25	32.60	29.37	29.04
F1 $\uparrow$	65.82	46.40	48.43	55.87	37.16	31.35	<b>47.50</b>	66.03	44.65	47.23	56.09	36.84	32.26	47.18	44.07	43.16
Prec. $\uparrow$	68.34	49.47	46.65	59.23	33.04	38.26	49.16	73.13	47.21	47.27	60.69	37.79	41.28	<b>51.23</b>	46.02	44.05
Rec. $\uparrow$	64.31	48.02	55.46	54.73	48.51	28.02	<b>49.84</b>	60.53	48.25	52.00	53.69	38.40	28.43	46.88	47.61	46.83
CD $\downarrow$	4.14	13.47	16.85	5.05	21.29	13.35	<b>12.36</b>	4.01	13.93	18.89	5.28	21.38	16.32	13.30	14.37	14.33
GMR $\downarrow$	32.13	53.68	53.81	43.00	69.97	57.44	51.67	26.62	56.48	53.04	41.04	64.00	54.17	<b>49.23</b>	54.73	55.90
<i>uncertainty ignored</i>																
IEQ $\uparrow$	0.85	1.15	0.72	1.86	0.24	0.79	0.93	1.09	1.40	1.11	2.07	0.69	0.65	<b>1.17</b>	1.07	1.10
<i>uncertainty considered</i>																
IEQ $\uparrow$	1.00	1.23	0.78	1.86	0.22	0.79	0.98	1.26	1.58	1.13	1.83	0.65	0.61	1.17	1.17	<b>1.22</b>

the *dropout* model beating the *binary* in IEQ.

### E. Novel Instance Generalization

Table II shows the evaluation of the *novel instance* models on the held-out objects from the test set. The results from the novel view experiment are consolidated, mirroring the similar performance between *binary* and *trinary* on the occupied region and superiority of *trinary* in predicting the uncertain region. Nonetheless, the *binary* model takes the lead in IEQ by a small margin over the *trinary*. Again, both baselines are not competitive. As expected, the general performance is lower than in the novel view case as the learning task is more challenging.

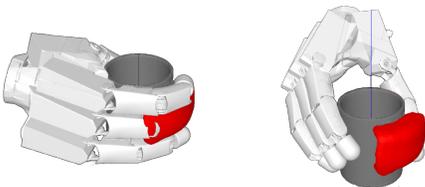


Fig. 4. Grasping the mug with an occluded handle without filtering using the predicted uncertain region leads to collision with the handle (left). Discarding grasps that collide with the uncertain region avoids collision and thus improves the grasp quality (right).

### F. Sim2Real Generalization

Table III shows the performance of the *novel instance* models on real data and thus their *sim2real* capabilities. The average performance of all methods across all metrics decreases noticeably. Many effects observed for real RGB-D sensors are hard to simulate, such as the effect of surface color, opacity, and reflectance on the depth data, leading to a large *sim2real* gap. Visual inspection of the depth data reveals that large parts of the objects are frequently missing although visible in the RGB image, removing in particular thin structures like the mugs’ handle. The effects on shape completion and prediction of uncertain regions are twofold: (i) non-occluded handles are often not reconstructed; (ii) false uncertain regions are often predicted on the back side of mugs. This mainly explains the large drop in performance of all methods.

Some reconstruction errors are especially accentuated by the high sensitivity of volumetric measures in the case of low-volume objects like mugs. However, we can see good CD values for the datasets with a better depth quality, i.e., HB<sub>pri</sub> and TYOL. Moreover, qualitative results shown in Fig. 5 reveal usable completions.



Fig. 5. Qualitative *sim2real* results. From top to bottom: Input point cloud, predicted mesh, ground-truth mesh. From left to right: test datasets HB<sub>pri</sub>, HB<sub>kin</sub>, LM, TYOL, YCBV<sub>48</sub>, YCBV<sub>55</sub>.

The number of images with an occluded handle having a ground-truth uncertain region is smaller than with a non-occluded handle missing from the depth image, hence with an erroneous uncertain region predicted on the back side. Consequently, a quantitative evaluation of the segmentation of uncertain regions is not sensible and therefore omitted.

Despite the—frequently large—discrepancy between the predicted and the actual uncertain regions, e.g. in most cases where handle data is missing in the depth image, we can still see an average positive effect in the IEQ values when considering predicted uncertainty in grasp planning.

### G. Discussion

From Fig. 3 and Tables I and II it is clear that the *trinary* method yields by far the highest accuracy in segmenting the uncertain regions. Nonetheless, the grasp qualities achieved when considering the predicted regions do not generally stand out, which might seem odd at first sight. However, the reason for it can be established as follows.

The robot hand for which we predict the grasp (DLR-Hand II [34]) is large compared to the mug objects. Hence even

if the prediction of the uncertain region has low precision, a moderate recall is sufficient to likely cause an intersection with the large fingers in that region, and thus the grasp is rejected as desired. However, the low precision is still reflected in a higher GER value. Indeed, for the *dropout* and *VAE* methods, there were cases where all predicted grasps were blocked by the predicted uncertain region, including too many false positives, resulting in grasping failure.

In applications with specific grasp requirements, e.g. regarding the approach direction or the object parts to be touched, the advantage of a more accurate prediction of uncertain regions is likely to make a significant difference. In addition to grasping, an accurate prediction of viewpoint-induced uncertainty regions is also relevant in many mobile robot applications.

## V. CONCLUSION

In this work, we introduced and compared methods for shape completion with uncertain regions relevant for safety critical real-world applications. We focused on viewpoint-induced irreducible uncertainty about the positions of object parts, considering the case of mugs with handles and the task of robotic grasping. We showed the advantage of predicting and avoiding uncertain regions in extensive experiments on synthetic and real data and the superiority of two novel methods over two existing, adapted baselines. Future work will consider more object classes and grasping scenarios and further improve sim2real capabilities.

## REFERENCES

- [1] A. X. Chang *et al.*, “Shapenet: An information-rich 3d model repository,” *ArXiv*, vol. abs/1512.03012, 2015.
- [2] J. Varley *et al.*, “Shape completion enabled robotic grasping,” *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 2442–2447, 2016.
- [3] M. V. der Merwe *et al.*, “Learning continuous 3d reconstructions for geometrically aware grasping,” *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 11 516–11 522, 2019.
- [4] 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.
- [5] B. Saund and D. Berenson, “Diverse plausible shape completions from ambiguous depth images,” in *Conference on Robot Learning*, 2020.
- [6] S. Peng *et al.*, “Convolutional occupancy networks,” in *European Conference on Computer Vision (ECCV)*. Cham: Springer International Publishing, Aug. 2020.
- [7] J. Lundell, F. Verdoja, and V. Kyrki, “Robust grasp planning over uncertain shape completions,” *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 1526–1532, 2019.
- [8] H. Fan, H. Su, and L. J. Guibas, “A point set generation network for 3d object reconstruction from a single image,” *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2463–2471, 2016.
- [9] L. P. Tchapmi *et al.*, “Topnet: Structural point cloud decoder,” *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 383–392, 2019.
- [10] W. Yuan *et al.*, “Pcn: Point completion network,” *2018 International Conference on 3D Vision (3DV)*, pp. 728–737, 2018.
- [11] Z. Wu *et al.*, “3d shapenets: A deep representation for volumetric shapes,” *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1912–1920, 2014.
- [12] A. Dai, C. Qi, and M. Nießner, “Shape completion using 3d-encoder-predictor cnns and shape synthesis,” *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 6545–6554, 2016.
- [13] T. Groueix *et al.*, “A papier-mache approach to learning 3d surface generation,” *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 216–224, 2018.
- [14] Y. Liao, S. Donné, and A. Geiger, “Deep marching cubes: Learning explicit surface representations,” *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2916–2925, 2018.
- [15] L. M. Mescheder *et al.*, “Occupancy networks: Learning 3d reconstruction in function space,” *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 4455–4465, 2018.
- [16] J. J. Park *et al.*, “DeepSDF: Learning continuous signed distance functions for shape representation,” *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 165–174, 2019.
- [17] J. Chibane, T. Alldieck, and G. Pons-Moll, “Implicit functions in feature space for 3d shape reconstruction and completion,” *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 6968–6979, 2020.
- [18] C. Qi, H. Su, K. Mo, and L. J. Guibas, “Pointnet: Deep learning on point sets for 3d classification and segmentation,” *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 77–85, 2016.
- [19] J. Tremblay *et al.*, “Deep object pose estimation for semantic robotic grasping of household objects,” in *Conference on Robot Learning (CoRL)*, 2018.
- [20] G. Shi *et al.*, “Fast uncertainty quantification for deep object pose estimation,” *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 5200–5207, 2020.
- [21] S. Prokudin, P. Gehler, and S. Nowozin, “Deep directional statistics: Pose estimation with uncertainty quantification,” in *European Conference on Computer Vision*, 2018.
- [22] I. Gilitschenski *et al.*, “Deep orientation uncertainty learning based on a bingham loss,” in *International Conference on Learning Representations*, 2020.
- [23] K. A. Murphy *et al.*, “Implicit-pdf: Non-parametric representation of probability distributions on the rotation manifold,” in *Proceedings of the 38th International Conference on Machine Learning*, 2021, pp. 7882–7893.
- [24] X. Deng *et al.*, “Poserbpf: A rao-blackwellized particle filter for 6d object pose estimation,” *Robotics: Science and Systems XV*, 2019.
- [25] D. Shin, C. C. Fowlkes, and D. Hoiem, “Pixels, voxels, and views: A study of shape representations for single view 3d object shape prediction,” *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 3061–3069, 2018.
- [26] W. E. Lorensen and H. E. Cline, “Marching cubes: A high resolution 3d surface construction algorithm,” *Proceedings of the 14th annual conference on Computer graphics and interactive techniques*, 1987.
- [27] Y. Gal, “Uncertainty in deep learning,” 2016.
- [28] Y. Gal and Z. Ghahramani, “Dropout as a bayesian approximation: Representing model uncertainty in deep learning,” in *Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48*, 2016, p. 1050–1059.
- [29] M. Denninger *et al.*, “Blenderproc2: A procedural pipeline for photorealistic rendering,” *Journal of Open Source Software*, vol. 8, no. 82, p. 4901, 2023. [Online]. Available: <https://doi.org/10.21105/joss.04901>
- [30] T. Hodaň *et al.*, “BOP: Benchmark for 6D object pose estimation,” *European Conference on Computer Vision (ECCV)*, 2018.
- [31] —, “BOP challenge 2020 on 6D object localization,” *European Conference on Computer Vision Workshops (ECCVW)*, 2020.
- [32] P. J. Besl and N. D. McKay, “A method for registration of 3-d shapes,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 14, pp. 239–256, 1992.
- [33] D. Winkelbauer *et al.*, “A two-stage learning architecture that generates high-quality grasps for a multi-fingered hand,” *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 4757–4764, 2022.
- [34] J. Butterfaß, M. Grebenstein, H. Liu, and G. Hirzinger, “DLR-Hand II: Next generation of a dextrous robot hand,” in *Proceedings of the IEEE International Conference on Robotics & Automation*, 2001, p. 109–114.
- [35] A. Paszke *et al.*, “Pytorch: An imperative style, high-performance deep learning library,” in *Neural Information Processing Systems*, 2019.
- [36] M. Abadi *et al.*, “TensorFlow: Large-scale machine learning on heterogeneous systems,” 2015, software available from tensorflow.org. [Online]. Available: <https://www.tensorflow.org/>
- [37] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *CoRR*, vol. abs/1412.6980, 2014.
- [38] C. Guo, G. Pleiss, Y. Sun, and K. Q. Weinberger, “On calibration of modern neural networks,” in *Proceedings of the 34th International Conference on Machine Learning - Volume 70*, 2017, p. 1321–1330.