

SIMstack: A Generative Shape and Instance Model for Unordered Object Stacks

Zoe Landgraf¹, Raluca Scona¹, Tristan Laidlow¹, Stephen James¹
Stefan Leutenegger² and Andrew J. Davison¹ *†‡

Abstract

By estimating 3D shape and instances from a single view, we can capture information about an environment quickly, without the need for comprehensive scanning and multi-view fusion. Solving this task for composite scenes (such as object stacks) is challenging: occluded areas are not only ambiguous in shape but also in instance segmentation; multiple decompositions could be valid. We observe that physics constrains decomposition as well as shape in occluded regions and hypothesise that a latent space learned from scenes built under physics simulation can serve as a prior to better predict shape and instances in occluded regions. To this end we propose SIMstack, a depth-conditioned Variational Auto-Encoder (VAE), trained on a dataset of objects stacked under physics simulation. We formulate instance segmentation as a centre voting task which allows for class-agnostic detection and doesn't require setting the maximum number of objects in the scene. At test time, our model can generate 3D shape and instance segmentation from a single depth view, probabilistically sampling proposals for the occluded region from the learned latent space. Our method has practical applications in providing robots some of the ability humans have to make rapid intuitive inferences of partially observed scenes. We demonstrate an application for precise (non-disruptive) object grasping of unknown objects from a single depth view.

1. Introduction

While humans are intuitively able to interpret partially observed scenes using geometric reasoning and prior experience, estimating 3D shape from RGB or depth images is challenging in computer vision due to ambiguity — many 3D shapes can explain a 2D observation. The classical approach to generate 3D shapes from depth images involves

*Research presented in this paper has been supported by Dyson Technology Ltd.

†¹Zoe Landgraf, Raluca Scona, Tristan Laidlow, Stephen James and Andrew J. Davison are with the Dyson Robotics Laboratory, Department of Computing, Imperial College London, UK. zoe.landgraf15@imperial.ac.uk

‡²Stefan Leutenegger is with the Smart Robotics Lab, Department of Computing, Imperial College London, UK.




Figure 1: **Top to bottom:** SIMstack outputs object shapes and instances from a single depth view on two YCB sequences and two real data examples, one with a fully occluded object supporting a leaning box (3rd row), for which our model predicts a plausible proposition. **Right:** Grasping demo setup (green: target object).

taking images from all sides of an object and fusing the re-projected points into a common 3D representation such as a truncated signed distance function (TSDF) [29]. However, apart from the exhaustive nature of the task, it is often impossible to reach all required viewpoints to generate a watertight surface reconstruction [12].

This drawback has led researchers to explore learning based approaches to reconstruct 3D shapes, such as learning to complete partial reconstructions [12, 11, 40] and predicting scenes [39] or objects [52, 51] from a single depth image. In parallel, researchers explore 3D shape prediction for scenes [30, 15, 38] and objects [9, 49, 53] from single or multi-view RGB data. While most of these approaches work with pointclouds, voxel occupancy grids or TSDF representations, some have explored alternative representations such as parametric surface elements [16], 2D sketches [45], graph neural networks [13], and the increasingly popular implicit neural surface representations [31, 26, 5].

To our knowledge, no existing work has explored 3D shape prediction with instance segmentation for multiple objects from a single depth image. Single view shape prediction allows to quickly estimate 3D occupancy, while instance segmentation is crucial for interactive tasks such as

object manipulation. We aim to solve this task for a tabletop scene with a variable number of stacked household objects and propose a VAE whose latent space is learned from stable (under physics simulation) scenes, with the aim to better reason about shape and instance decomposition in occluded regions. Conditioned on depth, the VAE learns to predict 3D shape and instances given a depth view, completing occluded regions using its latents, which can be thought of as an ‘intuitive physics’ prior. Our refinement method further optimises the reconstruction for collision-free decomposition. We aim to reconstruct scenes composed of unknown objects and train our model on randomly assembled piles of 3D parametric shapes (superquadrics). At test time, our conditional VAE (C-VAE) generates 3D shape and instance segmentation, as well as realistic reconstruction proposals for occluded regions in one forward pass. Our method can also integrate multiple views for improved reconstruction: the VAE can be conditioned on multiple views and our latent space can be further optimised against novel views using differentiable rendering. We show an application of our method for non-disruptive grasping using a robot arm.

In summary, our contributions are:

- A depth-conditioned VAE for scenes of stacked objects which can generate 3D shape, instance segmentation and probabilistic reconstruction proposals for occluded regions.
- A center voting scheme based on 3D Hough Voting allowing for class-agnostic 3D instance segmentation for scenes with an arbitrary number of objects.
- A shape refinement procedure to generate a compact scene representation of parametric shapes for downstream applications.

2. Related Work

Latent Representations for 3D Objects The shape of the occluded region of a 3D object is ambiguous and requires prior knowledge to estimate. Some methods use geometrical assumptions such as symmetry [56], but most tackle the problem using generative models and learn 3D shape priors using a latent representation [46, 7]. Zhang *et al.* [54] showed that such priors can generalize to unseen classes and Sucar *et al.* [41] used class conditioned latent models to generate 3D object shapes by optimising the latent code against depth views.

Instance Segmentation of 3D Objects 2D instance segmentation has progressed significantly with state-of-the-art methods such as Mask R-CNN [18] and DETR [4], but segmenting objects in 3D remains challenging. Hou *et al.* [20] extend the idea of region proposals to 3D, combining backprojected 2D features with 3D partial scans. Another method for segmenting 3D pointclouds into objects is center

voting [35, 34, 17]. We implement instance segmentation similarly, applying Hough Voting to 3D voxelgrids. The instance segmentation approach of Xie *et al.* [48] is closely related to ours — their context is a cluttered table top and they use Hough Voting. However, their final output is a 2D instance segmentation whereas we aim at 3D output.

Object Decomposition Previous work on decomposing compound objects into parts has focused on rigid objects such as furniture with significant structure and symmetry. Decomposing such objects into parts can help to leverage symmetry for shape completion [42] or discover structure in unseen data [43]. Most approaches decompose into geometric primitives [14, 43] using a library of shapes. Others use hierarchical representations and graph neural networks [27]. Paschalidou *et al.* [32] showed that using superquadrics over shape primitives improves reconstruction quality. We model stacked objects as a compound object and leverage superquadrics to approximate a large variety of composing shapes. Unlike manufactured objects such as chairs, object stacks have no symmetry or class-specific structure. However, we show that their decomposition can be learned using a lower-dimensional representation.

3. Method

Task definition Given a single depth image of objects stacked on a tabletop, our task is to estimate the complete 3D shape of the group and segment the objects into instances. We learn a latent space that describes realistic pile configurations using a VAE, to improve segmentation and reconstruction in occluded regions. To allow the latents to focus on learning the occluded regions, we condition our VAE on features from one or multiple depth images. Our goal is to train a generative model in the form of $p(x, \gamma, z) = p(z)p(x|z, \gamma)$ where $p(z)$ is a zero-mean multi-variate Gaussian prior describing realistic pile configurations. $p(x|z, \gamma)$ is the likelihood of pile configuration x given latents z and depth encoding γ which we model with our depth conditioned, generative decoder. We represent 3D geometry as a truncated signed distance function (TSDF) at a resolution of 64^3 , and instances using a center-voting vector field 3.2.1 at the same resolution. **Pipeline** At test time our method takes as input one or more depth images and generates the complete TSDF of the object stack and a center-voting vector field, which, processed by our 3D Hough Voting algorithm, provides instance segmentation. We refine our output by fitting a superquadric to each predicted mesh (overview in Figure 2).

3.1. A C-VAE for 3D Shape Instance Prediction

3.1.1 Network Architecture and training

Our network architecture consists of a 3D VAE which learns a latent space for realistic pile configurations and a condi-

tioning network whose encoding layers feed their learned feature maps to the VAE, conditioning it on partial observations. Both encoder and decoder of the VAE split into task specific branches in the first and last layers, to allow for separate encoding and decoding for instance and TSDF prediction. To learn descriptive depth features, the conditioning network is trained by auto-encoding a partial TSDF and is trained jointly with the VAE. We choose Residual blocks [19] with Squeeze and Excitation (SE) [21] over conventional convolutional layers in the joint encoder and decoder of the VAE to avoid gradient underflow and favour context integration. Since we compress both TSDF and instance segmentation into one latent code of size 96, it learns a consistent joint representation of the geometry and instance segmentation of a scene. Our architecture can be seen in Figure 2. **Training** To be effective, reconstruction has to be possible from any viewpoint. We therefore train our model on random viewpoints uniformly sampled during training in the range $(-1\text{m}, 1\text{m})$ for lateral positions and $(10\text{cm}, 1\text{m})$ for camera height. We generate a partial TSDF for every novel viewpoint as described in Section 3.1.2).

3.1.2 Conditioning on a Depth Image

Conditioning a 3D decoder on 2D depth image features would require the network to implicitly learn a reprojection task. To let it focus on the main task, we convert the input depth image(s) into a partial TSDF by reprojecting depth into the scene. We use inverse ray tracing for this and sample along each ray to fill the voxels in the camera view frustum with the closest distance to a surface. Parallelized, our algorithm generates a partial TSDF in under 0.8 seconds and can be used in online training.

3.1.3 Loss Definition

When predicting a TSDF, areas close to the surface with lower values will incur a smaller penalty than more distant positions. Other work [50] has proposed masking the loss by applying a different weight near the surface. We observe that this approach adds a discontinuity to the loss as well as an additional hyperparameter δ for the mask area. Instead we propose a loss which adds the inverse of the TSDF itself as a weighting factor:

$$L_{\text{TSDF}} = \sum_{x=0}^X \sum_{y=0}^Y \sum_{z=0}^Z \frac{1}{|v_{x,y,z}| + \epsilon} \|v_{x,y,z} - \hat{v}_{x,y,z}\|_2, \quad (1)$$

where $v_{x,y,z}$ and $\hat{v}_{x,y,z}$ are the ground truth and predicted TSDF values at voxelgrid index (x, y, z) , ϵ prevents division by 0 (set to $1e^{-9}$) and $\|\cdot\|_2$ is the L_2 -norm. We use the same loss for the partial TSDF prediction and refer to it as $L_{p,\text{TSDF}}$. For the instance vector voting task, we use the L_2 -norm between predicted and ground truth center vectors

\hat{c} and \bar{c} :

$$L_{\text{center_votes}} = \sum_{x=0}^X \sum_{y=0}^Y \sum_{z=0}^Z \|\bar{c}_{x,y,z} - \hat{c}_{x,y,z}\|_2. \quad (2)$$

We found that this led to better convergence for the instance segmentation task than using cosine similarity. To approximate our prior, we empirically found that training using the Maximum-Mean Discrepancy D_{MMD} [55] led to better convergence than using the standard KL-divergence [24]. Our final loss is composed of 4 components:

$$L_{\text{total}} = \alpha L_{\text{TSDF}} + \beta L_{p,\text{TSDF}} + \gamma L_{\text{center_votes}} + \delta D_{\text{MMD}}. \quad (3)$$

We set $\alpha = \beta = \gamma = 1$ and δ to $1e^5$.

3.1.4 A dataset of SuperQuadratic Piles

Instead of training on specific shapes from datasets such as ShapeNet [6] or ModelNet [47], we use generic shapes to favour generalisation. Superquadratics offer a general shape description, extending quadratics to multiple exponents, and are defined as the solution to the implicit equation:

$$f(x, y, z) = \left| \frac{x}{a_1} \right|^r + \left| \frac{y}{a_2} \right|^s + \left| \frac{z}{a_3} \right|^t = 1. \quad (4)$$

For exponents $r, s, t \geq 1$ the shapes are convex and varying the scale parameters a_1, a_2, a_3 generates very flat, small, or bulky shapes, ideal for approximating everyday household objects. We leverage this continuous parametrisation to sample a large variety of shapes (3500 individual shapes with exponents ranging between 2 to 100 and scales between 5 cm and 30 cm) and generate 10,000 realistic object piles using PyBullet [10]. For each scene we randomly place between 3 and 4 superquadratics (see Figure 3).

3.2. Postprocessing

3.2.1 Class-agnostic Instance Segmentation

To make our instance segmentation class-agnostic and independent of the number of objects present, we implement instance segmentation using 3D Hough voting. The instance branch of our network predicts a vector field in which each voxel ‘votes’ for the object it belongs to by predicting a 3D unit vector \hat{c} from its own centroid to the object’s centroid. (see Figure 4). At a 10-fold higher resolution than our voxel grid (640×640), we count the number of rays that traverse each voxel by marching each ray through the grid. We use the object bounds obtained from the TSDF prediction to limit raycasting to the inside of objects. A higher resolution allows for more detailed prediction of centers which don’t coincide with voxel centers at the original resolution. Those voxels with a number of traversals larger than μ are




Figure 2: Overview of our method (excluding shape fitting). Networks are trained jointly. Pre-processing steps (grey) are not optimised.



Figure 3: **Left:** 2 object stacks from our Superquadric (SQ) training set. **Right:** 3 stacks from our YCB-object test set. All stacks are generated by random placement under physics simulation.

selected as object center votes and passed to MeanShift to obtain the final center locations. To allocate voxels to their corresponding centers, we compute:

$$\min_{c_1, \dots, c_N} |\arccos(\hat{c} \cdot \hat{c}_n)| + \gamma * \|\hat{c}_n\| \quad (5)$$

whereby \hat{c}_n is the normalised distance from the voxels' center to cluster center n . Setting hyperparameters $\mu = 10$ and $\gamma = 0.1$ gave the best performance. While we observe limited sensitivity to γ , very large (≥ 25) and small (≤ 4) values of μ increase under- and over-segmentation respectively. Parallellized on the GPU our method takes on average 19ms for a scene resolution of 64^3 .




Figure 4: Our Hough Voting method illustrated for a single object center. **Left:** The network output — unit vectors pointing towards the center of the object. **Center:** Raycasting and extracting those voxels with a high number of traversal at a 10-fold voxelgrid resolution. **Right:** Voxels with a high ray-traversal rate are likely close to the object center and selected as 'votes'. To obtain final centers, selected votes are further processed using Mean shift clustering.

3.2.2 Shape Refinement using Superquadric Fitting

Our system outputs a TSDF and voxelised instance segmentation of the compound object; if needed, additional post-processing can be applied (e.g. primitive shape fitting, CAD-model fitting). To generate a compact representation, we propose following refinement procedure, which extracts a set of superquadrics and their poses from the raw output:

First, we run marching cubes on the TSDF and sample from the resulting mesh to get a point cloud. We segment this point cloud according to the voxelised instance predictions. We then independently fit a superquadric to the N_i points belonging to the i th segment by minimising the “mean distortion” [8], $L_{\text{dist},i}$, given by:

$$L_{\text{dist},i} = \frac{1}{N_i} \sum_{k=1}^{N_i} \sqrt{a_{i1} a_{i2} a_{i3}} \cdot d_i(x_{i,k}, y_{i,k}, z_{i,k})^2, \quad (6)$$

where

$$d_i(x_{i,k}, y_{i,k}, z_{i,k}) = \frac{\|\overrightarrow{O_i P_k}\| f_i(x_{i,k}, y_{i,k}, z_{i,k}) - 1}{f_i(x_{i,k}, y_{i,k}, z_{i,k})}, \quad (7)$$

a_{i1}, a_{i2}, a_{i3} are the scale parameters and $f_i(x_{i,k}, y_{i,k}, z_{i,k})$ is the implicit equation from Eq. 4, $d_i(x_{i,k}, y_{i,k}, z_{i,k})$ is the approximate Euclidean distance to the surface [1], $\|\overrightarrow{O_i P_k}\|$ is the distance from the point to the superquadric center.

We then optimise for the parameters of all Q superquadrics together, with additional cost terms to penalise collisions between superquadrics and intersection with the floor plane:

$$L = \sum_i^Q \left[\lambda_{\text{dist}} L_{\text{dist},i} + \lambda_{\text{coll}} \sum_{\substack{j=1 \\ i \neq j}}^Q L_{\text{coll},ij} + \lambda_{\text{floor}} L_{\text{floor},i} \right]. \quad (8)$$

The cost of sample points from superquadric j colliding with superquadric i , $L_{\text{coll},ij}$, is given by:

$$L_{\text{coll},ij} = \frac{1}{N_j} \sum_{k=1}^{N_j} \mathbb{1}_{d_i(x_{j,k}, y_{j,k}, z_{j,k}) < 0} \cdot d_i(x_{j,k}, y_{j,k}, z_{j,k})^2, \quad (9)$$

where $\mathbb{1}$ is the indicator function, and the cost of sample points from superquadric i colliding with the floor plane, $L_{\text{floor},i}$, is given by:

$$L_{\text{floor},i} = \mathbb{1}_{z_{i,k} < 0} \cdot z_{i,k}^2. \quad (10)$$

Experimentally, we found $\lambda_{\text{dist}} = 100$, $\lambda_{\text{coll}} = 10$, and $\lambda_{\text{floor}} = 1$ to work well. For both stages, we minimised the cost function using the dogleg algorithm with rectangular trust regions, bounding the superquadric parameters to enforce convex shapes.

3.3. Multi-View Estimation

Our method supports multi-view reconstruction in two ways: the generative decoder can be conditioned on a partial TSDF generated from fusing multiple depth images (*multi-view conditioning*); or the latent code can be optimised from one or multiple depth views using differentiable depth rendering (*multi-view optimisation*). To implement multi-view conditioning, we backproject multiple depth views into the scene using inverse raytracing as in Section 3.1.2. To demonstrate how our latent code can be optimised against multiple depth views, we implement a differentiable depth renderer for raytracing TSDFs and an optimization method described in the supplementary material.

4. Experiments

We evaluate our method quantitatively on our own test dataset (1419 scenes) and a test dataset of YCB [3] objects (256 scenes) to demonstrate generalisation. For the latter we select YCB objects with IDs: 2, 3, 4, 5, 7, 8, 9, 10, 36, 61 which can be approximated by a single superquadric and generate object stacks using the same procedure we used for our own synthetic scenes (Section 3.1.4). We show dataset examples in Figure 3. We demonstrate qualitative result on our synthetic datasets in Figure 9 as well as on real data examples and sequences of the YCB-Video dataset (Figures 1, 10 and 5) and evaluate our method for multi-view reconstruction. We show how our method generalises to scenes with more objects and demonstrate an application of our method for non-disruptive grasping on unknown objects.




Figure 5: Sampling from the latent space for a real data example. We see multiple pile samples which all look similar from the point of view of the input depth observation (**top**) but significantly different though all plausible from the back (**bottom**).

4.1. Single-view Reconstruction

We are not aware of directly comparable work on 3D instance segmentation for object groups and therefore select the following two baselines to compare our method against. **SSCNet Baseline** The closest existing method is SSCNet [39], which predicts a 3D occupancy grid with semantic labels for multi-object rooms from a single depth image. As this method does not predict instances, we use it as a base-

line only for geometric reconstruction and adapt it as follows: (1) we change the last layer to predict a TSDF; (2) we remove downsampling (required only in large-scale scenes) from layers by adding padding as well as one upsampling layer. We refer to this modified version as *SSCNet***.

Fully Convolutional (FC) Baseline Secondly, we use an ablated version of our model: a fully convolutional network predicting a TSDF and instance vector field from a single depth image. This serves both as a baseline for instance segmentation and an ablation study, showing the advantage of a learned prior over direct prediction. Our ablated model uses the same SE and Residual units as our main model but no feature space compression, as we found it reduced convergence. A detailed description of baseline architectures can be found in the supplementary material.



Figure 6: **Left:** Ground truth with the visible part of the mesh highlighted in red. **Right:** 3 random latent code samples generated by our network. Note how every sample has different reconstructions in the occluded areas, which are plausible but would yield high reconstruction error with respect to the ground truth TSDF.

4.1.1 Reconstruction Accuracy

Comparing the full reconstruction to the ground truth scene unfairly penalises the network in occluded regions for reconstructions which are plausible, but differ from the ground truth (see Figure 6). We therefore evaluate our method on the visible surface area and the full reconstruction separately. To show how performance differs depending on how much of the object stack is visible from a given viewpoint, we report results by surface visibility. Note that given a single view, the maximum surface visibility of a pile is around 50%, since the back will always be occluded. Since our method is viewpoint agnostic, we evaluate every test scene from 3 random viewpoints uniformly sampled from the same ranges used during training (Section 3.1.1). We ensure the distance between camera and scene centers is at least 50cm. For every viewpoint, we generate a 3D scene and compute the average reconstruction accuracy for 3 latent code samples to account for variability. **Visible surface evaluation** To evaluate surface reconstruction accuracy, we extract a mesh from the generated and ground truth TSDFs using marching cubes. We obtain the visible surface by extracting all visible faces using ray-triangle intersection with the generated mesh. We produce 1000 uniform samples from the extracted surface and evaluate the bidirectional Chamfer Distance between ground truth and

predicted point sets, P_m and P_k :

$$CD = \frac{1}{N_m} \sum_{x_i \in P_m} \min_{x_j \in P_k} (\mathbf{x}_i - \mathbf{x}_j) + \frac{1}{N_k} \sum_{x_i \in P_k} \min_{x_j \in P_m} (\mathbf{x}_j - \mathbf{x}_i), \quad (11)$$

where \mathbf{x}_i and \mathbf{x}_j are the 3D coordinates of sampled points from P_m and P_k respectively. **Full reconstruction evaluation** We evaluate the full reconstruction in terms of surface reconstruction accuracy and predicted voxel occupancy. For the former we compute the Chamfer Distance between 1000 samples from the full predicted and ground truth mesh surfaces. For the latter we compute the Binary Cross Entropy (BCE):

$$-\frac{1}{N} \sum_{\Omega} y_i \log\left(\frac{1}{S} \sum_{j=0}^{j=S} \hat{y}_j\right) + (1 - y_i) \log\left(1 - \frac{1}{S} \sum_{j=0}^{j=S} \hat{y}_j\right), \quad (12)$$

where N is the voxelgrid resolution 64^3 , S is the number of latent code samples taken per viewpoint, which we set to 3 and \hat{y} and y are the predicted and ground truth occupancy value respectively, computed as follows:

$$Occupancy_{x,y,z} = \begin{cases} 0 & \text{if } TSDF_{x,y,z} > 0 \\ 1 & \text{if } TSDF_{x,y,z} \leq 0 \end{cases} \quad (13)$$

We also report the Intersection over Union, whereby we use a threshold of 0.5 to distinguish between (label) occupied and unoccupied space. We evaluate our baselines using the same procedure and display our results in Figure 7 and Table 1. Note that we set S to 1 in (12) for our baselines and just use the prediction \hat{y} . Our method outperforms both baselines for full reconstruction (Chamfer Distance evaluated for the full mesh surface, and expected occupancy) and visible surface reconstruction, showing its overall advantage for reconstructing and decomposing a 3D scene. Note that the gap between visible and full surface reconstruction is largest for *SSCNet***, but more comparable between the FC baseline and our method. As expected for all methods, with lower visibility full surface reconstruction accuracy drops faster than visible surface accuracy. The drop is however most pronounced for *SSCNet***. These observations suggest that while using a prior improves overall performance, a large improvement in predicting occluded space is achieved by jointly predicting shape and instances.

4.1.2 Scene decomposition (instance segmentation)

Similarly to the reconstruction accuracy, it is impossible to make an exact quantitative evaluation of our method’s instance segmentation since decomposition estimates in occluded regions can be plausible even when very different from the ground truth. In Figure 5, for example, our method hypothesises plausible hidden objects, although the ground




Figure 7: Chamfer Distance for 1000 samples on our test datasets of SuperQuadric shapes (**left**) and object stacks composed of YCB objects (**right**). We plot results for different surface mesh visibility ratios. **VS**: Visible surface. **FS**: Full surface

truth does not contain one. We therefore decided to evaluate instance segmentation as a valid scene decomposition which generates stable piles in physics simulation. **Stability evaluation under physics simulation** We evaluate stability by loading our generated 3D meshes into the PyBullet physics engine and simulating 10000 steps with a gravity setting of $10 \frac{m}{s^2}$ along $-z$, a friction coefficient of 1 and assuming uniform density. We compare our method against our FC baseline (which provides instance segmentation) in terms of object center displacement after simulation and report results in Table 1. Our method outperforms the fully convolutional baseline by over $50cm$ on both our test datasets indicating that the collections generated by our method are generally more stable and hence plausible. Note that in the case of rolling objects as well as intersecting objects (causing objects to be pushed apart due to contact force), average object displacement can become large.




Figure 8: Example of multi-view reconstruction comparing multi-view conditioning (MV cond), multi-view optimisation (MV opt) to TSDF-Fusion. We show raw results (no SQ fitting)

4.1.3 Evaluating on Real Data

We qualitatively evaluate the proposed method on real data. We collect RGB-D images of structures (leveraging ORB-SLAM [28] poses) and post-process them by (1) segmenting the floor plane using RANSAC [37] and (2) rectifying the roll and pitch of the camera pose by aligning the floor plane normal with the vertical axis. Our results (Figure 10, 1 and 5) show that our method generates correct 3D reconstruction, realistic occlusion proposals and instance seg-




Figure 9: Comparing our method against our baselines on examples from our SQ test dataset and our YCB test dataset. More qualitative examples are shown in the supplementary material.

mentation for a number of real world examples. Our example in Figure 5 shows that our method realistically fills in the occluded area below the box by extending its shape to the floor. When sampling from the latents, it is able to hypothesise different occluded objects. In particular, our example in Figure 1 (3rd row) shows that our network is capable to hypothesise a supporting object which is completely occluded but required to support a leaning box. We also evaluate our method on YCB-Video sequences (Figure 1).




Figure 10: Qualitative results on real data examples. Our method generates 3D shape and instances from single depth views.

4.2. Multi-View Reconstruction

We evaluate the performance of our method using multi-view conditioning and multi-view optimisation. We use simple TSDF Fusion as a baseline which we implement using the method described in 3.1.2. For a fixed sequence of 6 viewpoints (see Figure 13) we compare the reconstruction quality for our SQ test dataset. For multi-view conditioning, we average reconstruction accuracy over 2 latent code samples. Note that the first view covers on average 30% of the test scenes. For multi-view optimization, we condition

	ours		FC		SSCNet**	
Avg. Chamfer Distance (m)						
	(μ)	(σ)	(μ)	(σ)	(μ)	(σ)
visible surface (SQ dataset)	0.016	0.0002	0.044	0.002	0.036	0.002
full surface (SQ dataset)	0.028	0.0003	0.062	0.107	0.003	0.002
visible surface (YCB dataset)	0.036	0.002	0.076	0.006	0.065	0.008
full surface (YCB dataset)	0.062	0.007	0.098	0.011	0.116	0.009
Expected occupancy						
Binary Cross Entropy						
SQ dataset	0.089	0.003	0.175	0.005	0.247	0.008
YCB dataset	0.112	0.006	0.236	0.029	0.242	0.015
Intersection over Union						
SQ dataset	0.829	0.004	0.635	0.005	0.424	0.004
YCB dataset	0.784	0.003	0.542	0.012	0.460	0.006
Avg. object centre displacement (m)						
SQ dataset	0.593	1.28	1.112	5.441	–	–
YCB dataset	1.216	4.734	1.756	6.494	–	–

Table 1: Comparison in terms of Chamfer Distance, expected occupancy and stability under physics simulation. We report the average results for 3 viewpoints per scene.

on the first view and optimise the latent code against the additional views (see supplementary material for details). Our experiments show that multi-view conditioning generates the best reconstruction and that both our multi-view methods outperform TSDF Fusion, even at very high visibility. We attribute this to the fact that our generative method also reconstructs the bottom of objects, which will always be occluded and therefore cannot be reconstructed by TSDF Fusion. Our example in Figure 8 shows how a scene reconstruction is updated with additional views. Please refer to our supplementary material and video for additional qualitative examples and runtime analysis.

4.2.1 Additional qualitative experiments

More objects and non-convex objects Although only trained on 3-4 objects, SIMstack generalises well to scenes with up to 7 objects (see Figure 12). We attribute this to the fact that our formulation of instance segmentation is independent of the number of objects in the scene and that our model learns from scenes with varying number of objects. Our method also shows some ability to generalise to non-convex objects (see our supplementary material for examples). **Latent code analysis** Our multi-view optimization experiments show that our latent code can be optimised against novel views, generating consistent shape and instance predictions (Figure 8). This suggests that our shape encoding is smooth as well as consistent between shape and instance representation. We provide further qualitative examples demonstrating this in the supplementary material.

4.3. Precise Robot Manipulation

Robotic grasping is a well studied problem [2] and an active area of research [33, 25, 23]. However, many state-of-the-art grasping systems perform indiscriminate grasping, with no regard of how the grasp might effect surrounding objects. However, multi-object reasoning is essential when grasping in cluttered scenes; grasping in an imprecise man-




Figure 11: Two examples from our grasping system. Target objects are highlighted in green. The first example shows the robot sliding the target object sideways, leaving the other objects undisturbed. The second example grasps the object from the side and pulls it out such that the top resting object slides onto a lower supporting object, causing minimum disruption to the stack.




Figure 12: Qualitative results on scenes with 7 objects, demonstrating the ability of our method to generalise to more objects.




Figure 13: Multi-view estimation. **Left:** The 6 viewpoints for which we evaluate. **Right:** Chamfer Distance (1000 points) for increasing views of TSDF Fusion, multi-view conditioning (MV Conditioning) and multi-view optimisation (MV Optimisation).

ner may topple a stack and cause damage to fragile objects. In this section, we show how SIMstack can be used to perform precise 6D grasping of a target object while minimising disruption to surrounding objects.

Our grasping demo consists of a real stack of (unknown) objects, a Franka Panda robot arm, and a suction gripper (Figure 1). An image of the scene is fed to SIMstack which outputs meshes and poses. These, along with a target object, are loaded into CoppeliaSim/PyRep [36, 22] where 5 virtual cameras are used to create a pointcloud of the visible surface of the stack. The target object’s pointcloud is extracted, and grasping locations are sampled based on sur-

face normals. We exhaustively simulate each valid grasp and measure the mean displacement of all objects (excluding the target object) after the grasp has been made. The grasp which produces the lowest mean displacement is chosen to run on the real platform. Note that due to COVID restrictions, we were unable to show the grasps running on the real robot. Two examples of successful grasps in simulations based on SIMstack reconstructions of real object piles are shown in Figure 11, and a demo of the system can be seen in the supplementary video.

5. Discussion and Conclusion

We propose a novel method (SIMstack) to generate 3D shape and class-agnostic instance segmentation for multiple stacked objects from a single depth image. We learn a joint shape and instance encoding, trained on a dataset of parametric shapes, randomly assembled under physics simulation. This encoding acts as an intuitive physics prior for realistic object stacks, improving reconstruction and segmentation in occluded regions. SIMstack can generate a 3D shape and instance decomposition of a collection of (convex) objects from a single depth view. This output allows for a quick estimate of shape and instance decomposition, which could be used for downstream applications (e.g to initialise a multi-view scanning system to capture more detail [44]) and allows for fast multi-object reasoning, useful for interactive tasks such as the precise (non-disruptive) grasping we demonstrate. Our method can leverage multi-view information for improved reconstruction which makes it by design a candidate for incremental reconstruction systems.

We believe our approach can play an important role in rapid scene understanding to help embodied AI systems make physically intuitive interpretations of ambiguous scenes. Promising future research directions include integrating RGB information and extending our approach to non-convex objects.

References

- [1] Eric Bardinet, Laurent D. Cohen, and Nicholas Ayache. A parametric deformable model to fit unstructured 3D data. *RR-2617, INRIA*, 1995. 4
- [2] Jeannette Bohg, Antonio Morales, Tamim Asfour, and Danica Kragic. Data-driven grasp synthesis—a survey. *IEEE Transactions on Robotics (T-RO)*, 2014. 7
- [3] B. Calli, A. Singh, A. Walsman, P Srinivasa S. and, Abbeel, and A. M. Dollar. The ycb object and model set: Towards common benchmarks for manipulation research. In *International Conference on Advanced Robotics (ICAR)*, pages 510–517, 2015. 5
- [4] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 2
- [5] Rohan Chhabra, Jan Eric Lenssen, Eddy Ilg, Tanner Schmidt, Julian Straub, Steven Lovegrove, and Richard Newcombe. Deep local shapes: Learning local SDF priors for detailed 3D reconstruction. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 1
- [6] Angel X. Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, Jianxiong Xiao, Li Yi, and Fisher Yu. ShapeNet: An information-rich 3D model repository. *arXiv preprint arXiv:1512.03012*, 2015. 3
- [7] Zhiqin Chen and Hao Zhang. Learning implicit fields for generative shape modeling. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2
- [8] Laurent Chevalier, Fabrice Jaillet, and Atila Baskurt. Segmentation and superquadric modeling of 3D objects. *Journal of WSCG*, 11(10), 2003. 4
- [9] Christopher B. Choy, Danfei Xu, JunYoung Gwak, Kevin Chen, and Silvio Savarese. 3D-R2N2: A unified approach for single and multi-view 3D object reconstruction. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2016. 1
- [10] Erwin Coumans and Yunfei Bai. PyBullet, a Python module for physics simulation for games, robotics and machine learning. <http://pybullet.org>, 2016–2019. 3
- [11] Angela Dai and Matthias Nießner. Scan2Mesh: From unstructured range scans to 3D meshes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 1
- [12] Angela Dai, Daniel Ritchie, Martin Bokeloh, Scott Reed, Jürgen Sturm, and Matthias Nießner. ScanComplete: Large-scale scene completion and semantic segmentation for 3D scans. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 1
- [13] Angela Dai, Charles Ruizhongtai Qi, and Matthias Nießner. Shape completion using 3D-Encoder-Predictor CNNs and shape synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 1
- [14] Boyang Deng, Kyle Genova, Soroosh Yazdani, Sofien Bouaziz, Geoffrey Hinton, and Andrea Tagliasacchi. CvxNet: Learnable convex decomposition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2
- [15] Michael Firman, Oisín Mac Aodha, Simon Julier, and Gabriel J. Brostow. Structured prediction of unobserved voxels from a single depth image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 1
- [16] Thibault Groueix, Matthew Fisher, Vladimir G. Kim, Bryan C. Russell, and Mathieu Aubry. A papier-mâché approach to learning 3D surface generation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 1
- [17] Lei Han, Tian Zheng, Lan Xu, and Lu Fang. OccuSeg: Occupancy-aware 3D instance segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2
- [18] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask R-CNN. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2017. 2
- [19] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 3
- [20] Ji Hou, Angela Dai, and Matthias Nießner. 3D-SIS: 3D semantic instance segmentation of RGB-D scans. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2
- [21] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 3
- [22] Stephen James, Marc Freese, and Andrew J Davison. Pyrep: Bringing v-rep to deep robot learning. *arXiv preprint arXiv:1906.11176*, 2019. 8
- [23] Stephen James, Paul Wohlhart, Mrinal Kalakrishnan, Dmitry Kalashnikov, Alex Irpan, Julian Ibarz, Sergey Levine, Raia Hadsell, and Konstantinos Bousmalis. Sim-to-real via sim-to-sim: Data-efficient robotic grasping via randomized-to-canonical adaptation networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 7
- [24] Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2014. 3
- [25] Jeffrey Mahler, Jacky Liang, Sherdil Niyaz, Michael Laskey, Richard Doan, Xinyu Liu, Juan Aparicio Ojea, and Ken Goldberg. Dex-net 2.0: Deep learning to plan robust grasps with synthetic point clouds and analytic grasp metrics. *Proceedings of Robotics: Science and Systems (RSS)*, 2017. 7
- [26] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. NeRF: Representing scenes as neural radiance fields for view synthesis. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 1
- [27] Kaichun Mo, Paul Guerrero, Li Yi, Hao Su, Peter Wonka, Niloy J. Mitra, and Leonidas J. Guibas. StructureNet: Hierarchical graph networks for 3D shape generation. In *SIGGRAPH Asia*, 2019. 2

- [28] R. Mur-Artal and J. D. Tardós. ORB-SLAM2: An Open-Source SLAM System for Monocular, Stereo, and RGB-D Cameras. *IEEE Transactions on Robotics (T-RO)*, 33(5):1255–1262, 2017. 6
- [29] Richard A. Newcombe, Shahram Izadi, Otmar Hilliges, David Molyneaux, David Kim, Andrew J. Davison, Pushmeet Kohli, Jamie Shotton, Steve Hodges, and Andrew Fitzgibbon. KinectFusion: Real-time dense surface mapping and tracking. In *Proceedings of the International Symposium on Mixed and Augmented Reality (ISMAR)*, 2011. 1
- [30] Andrea Nicastro, Ronald Clark, and Stefan Leutenegger. X-Section: Cross-section prediction for enhanced RGB-D fusion. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2019. 1
- [31] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 1
- [32] Despoina Paschalidou, Ali Osman Ulusoy, and Andreas Geiger. Superquadrics revisited: Learning 3D shape parsing beyond cuboids. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2
- [33] Lerrel Pinto and Abhinav Gupta. Supersizing self-supervision: Learning to grasp from 50K tries and 700 robot hours. *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2016. 7
- [34] Charles R. Qi, Xinlei Chen, Or Litany, and Leonidas J. Guibas. ImVoteNet: Boosting 3D object detection in point clouds with image votes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2
- [35] Charles R. Qi, Or Litany, Kaiming He, and Leonidas J. Guibas. Deep Hough voting for 3D object detection in point clouds. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2019. 2
- [36] Eric Rohmer, Surya PN Singh, and Marc Freese. V-rep: A versatile and scalable robot simulation framework. In *Proceedings of the IEEE/RSJ Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2013. 8
- [37] Radu Bogdan Rusu and Steve Cousins. 3D is here: Point Cloud Library (PCL). In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2011. 6
- [38] Daeyun Shin, Zhile Ren, Erik B. Sudderth, and Charles C. Fowlkes. 3D scene reconstruction with multi-layer depth and epipolar transformers. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2019. 1
- [39] Shuran Song, Fisher Yu, Andy Zeng, Angel X. Chang, Manolis Savva, and Thomas Funkhouser. Semantic scene completion from a single depth image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 1, 5
- [40] David Stutz and Andreas Geiger. Learning 3D shape completion from laser scan data with weak supervision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 1
- [41] Edgar Sucar, Kentaro Wada, and Andrew J. Davison. NodeSLAM: Neural object descriptors for multi-view shape reconstruction. In *Proceedings of the International Conference on 3D Vision (3DV)*, 2016. 2
- [42] Minhuk Sung, Vladimir G. Kim, Roland Angst, and Leonidas Guibas. Data-driven structural priors for shape completion. *ACM Transactions on Graphics (TOG)*, 34(6), 2015. 2
- [43] Shubham Tulsiani, Hao Su, Leonidas J. Guibas, Alexei A. Efros, and Jitendra Malik. Learning shape abstractions by assembling volumetric primitives. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 2
- [44] Kentaro Wada, Edgar Sucar, Stephen James, Daniel Lenton, and Andrew J. Davison. MoreFusion: Multi-object reasoning for 6D pose estimation from volumetric fusion. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 8
- [45] Jiajun Wu, Yifan Wang, Tianfan Xue, Xingyuan Sun, William T. Freeman, and Joshua B. Tenenbaum. MarrNet: 3D shape reconstruction via 2.5D sketches. In *Neural Information Processing Systems (NeurIPS)*, 2017. 1
- [46] Jiajun Wu, Chengkai Zhang, Xiuming Zhang, Zhoutong Zhang, William T. Freeman, and Joshua B. Tenenbaum. Learning shape priors for single-view 3D completion and reconstruction. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018. 2
- [47] Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3D ShapeNets: A deep representation for volumetric shapes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015. 3
- [48] Christopher Xie, Yu Xiang, Arsalan Mousavian, and Dieter Fox. Unseen object instance segmentation for robotic environments. *arXiv preprint arXiv:2007.08073*, 2020. 2
- [49] Haozhe Xie, Hongxun Yao, Xiaoshuai Sun, Shangchen Zhou, and Shengping Zhang. Pix2Vox: Context-aware 3D reconstruction from single and multi-view images. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2019. 1
- [50] Qiangeng Xu, Weiyue Wang, Duygu Ceylan, Radomir Mech, and Ulrich Neumann. DISN: Deep implicit surface network for high-quality single-view 3D reconstruction. In *Neural Information Processing Systems (NeurIPS)*, 2019. 3
- [51] Bo Yang, Stefano Rosa, Andrew Markham, Niki Trigoni, and Hongkai Wen. Dense 3D object reconstruction from a single depth view. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 41(12):2820–2834, 2019. 1
- [52] Bo Yang, Hongkai Wen, Sen Wang, Ronald Clark, Andrew Markham, and Niki Trigoni. 3D object reconstruction from a single depth view with adversarial learning. In *Proceedings of the International Conference on Computer Vision Workshops (ICCVW)*, 2017. 1
- [53] Yuan Yao, Nico Schertler, Enrique Rosales, Helge Rhodin, Leonid Sigal, and Alla Sheffer. Front2Back: Single view 3D shape reconstruction via front to back prediction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 1

- [54] Xiuming Zhang, Zhoutong Zhang, Chengkai Zhang, Josh Tenenbaum, Bill Freeman, and Jiajun Wu. Learning to reconstruct shapes from unseen classes. In *Neural Information Processing Systems (NeurIPS)*, 2018. [2](#)
- [55] Shengjia Zhao, Jiaming Song, and Stefano Ermon. InfoVAE: Information maximizing variational autoencoders. *arXiv preprint arXiv:1706.02262*, 2017. [3](#)
- [56] Yichao Zhou, Shichen Liu, and Yi Ma. Learning to detect 3D reflection symmetry for single-view reconstruction. *arXiv preprint arXiv:2006.10042*, 2020. [2](#)