

Clutter Detection and Removal in 3D Scenes with View-Consistent Inpainting

Fangyin Wei

Thomas Funkhouser
Princeton University

Szymon Rusinkiewicz

Abstract

Removing clutter from scenes is essential in many applications, ranging from privacy-concerned content filtering to data augmentation. In this work, we present an automatic system that removes clutter from 3D scenes and inpaints with coherent geometry and texture. We propose techniques for its two key components: 3D segmentation based on shared properties and 3D inpainting, both of which are important problems. The definition of 3D scene clutter (frequently-moving objects) is not well captured by commonly-studied object categories in computer vision. To tackle the lack of well-defined clutter annotations, we group noisy fine-grained labels, leverage virtual rendering, and impose an instance-level area-sensitive loss. Once clutter is removed, we inpaint geometry and texture in the resulting holes by merging inpainted RGB-D images. This requires novel voting and pruning strategies that guarantee multi-view consistency across individually inpainted images for mesh reconstruction. Experiments on ScanNet and Matterport3D dataset show that our method outperforms baselines for clutter segmentation and 3D inpainting, both visually and quantitatively. Project page: <https://weify627.github.io/clutter/>.

1. Introduction

With the proliferation of RGB-D cameras, we can imagine a day when people can quickly scan their rooms with a phone and upload the reconstructed 3D model to a website for subleasing. It is more desirable to show an attractive, clean room without pillows and kitchenware scattered all around. However, people prefer not to spend hours cleaning them up before scanning. This is when a tool for automatic scene clutter removal would come in handy. Such a scene editing tool can be applied to a wide range of tasks from automatic content and privacy filtering to data augmentation (by adding new objects to cleaned scenes).

In this paper, we investigate how to reconstruct a scene without clutter, starting from a set of indoor RGB-D scans. This task consists of two steps: clutter segmentation and 3D inpainting, both of which present new challenges.

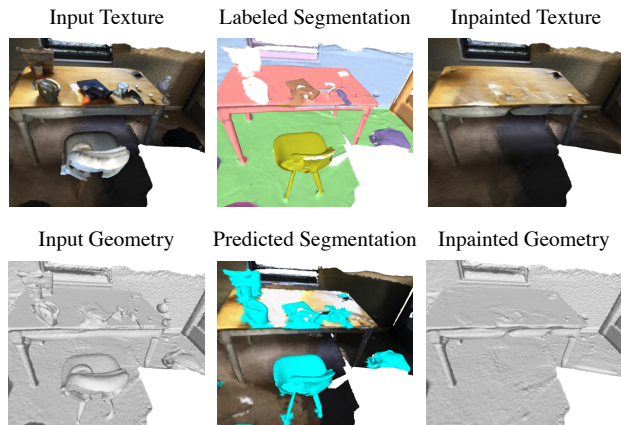


Figure 1. **Clutter detection and removal in 3D Scenes with view-consistent inpainting.** Beginning from an input scene (left), our first task is to predict clutter regions. Traditional semantic segmentation datasets and methods [9] (top center) were not designed for this problem, and contain missing or incorrect labels (especially for small objects) as well as semantic categories that only sometimes correspond to clutter. In contrast, our method (bottom center) correctly identifies clutter segments, which we then remove and inpaint with coherent texture and geometry (right).

Thanks to large-scale 3D indoor scene datasets [4, 9, 1, 40], data-driven methods have led to remarkable success in 3D scene segmentation [18, 34, 24]. However, existing annotations are insufficient for clutter segmentation for two reasons. First, ambiguity exists even in the most fine-grained datasets to date (around 550 raw categories for datasets [4, 9], 200 for benchmark [40]). For example, some objects from the category *whiteboard* are portable and belong to clutter, while others are installed on the wall and not clutter. Such intra-category ambiguity is widespread: consider *light*, *seat*, *speaker*, *decoration*, etc., all of which exist in both fixed (non-clutter) and mobile (clutter) variations. Secondly, the quality of annotations is often suboptimal, with some clutter objects being unannotated or mislabeled as non-clutter, as shown in Fig. 1, top center. Considering the mesh resolution and the time required to label a large number of small objects, it is almost impossible to annotate clutter precisely and exhaustively in a large dataset.

The task of 3D inpainting aims to synthesize a visually

and geometrically plausible completion for the missing regions of a 3D scene. While some prior work [41, 46, 31] inpaints smaller holes from incomplete scans or occlusions, we focus on larger holes created by object removal, which exacerbates the difficulty of inpainting due to complex geometry and texture. A naive solution is to run 3D hole-filling or mesh reconstruction algorithms on the remaining points. However, these algorithms only take into account local geometry and low-level texture information, whereas in many cases, semantic understanding is required for plausible completion. For example, when some clutter covering a table corner is removed, the inpainted scene is expected to recover the corner rather than just filling the hole and leaving the corner missing. Directly working on the 3D scene entails dealing with entangled semantics and complex geometry at the same time, which may result in low resolution and many artifacts in the output [12, 21]. We propose to instead inpaint the RGB-D sequences and then reconstruct the final mesh. This disentangles the problem into semantically plausible inpainting and geometry filling. The former can be solved more easily with existing advanced 2D inpainting methods and potentially more powerful than methods trained only on limited 2D data.

In this work, we endeavor to address the above challenges and propose an automatic system for scene clutter removal and 3D inpainting. To more effectively segment small clutter objects, we design an area-sensitive loss to force attention to small clutter objects which typically receive limited loss signals. We adopt a deep architecture that takes in both RGB images and 3D points as input and render virtual views as the 2D input to gain large surface coverage. Different from prior work that directly predicts inpainting in 3D after identifying objects in 3D, we project the object masks onto RGB-D images and perform image inpainting and image-guided depth completion. To enforce cross-frame consistency for a better quality of 3D reconstruction, we run novel consistency voting and pruning across frames. We loop back to image inpainting until all missing regions have been filled. The resulting collection of consistently inpainted RGB-D images can be merged into a final, clean scene.

We conduct extensive ablation studies to validate the system designs on the ScanNet [9] and Matterport3D [4] datasets. We also compare with 3D segmentation and inpainting baselines and show the effectiveness of the proposed model. In summary, our contributions are:

- An automatic system that solves the novel task of scene clutter removal and inpainting.
- A 3D segmentation method with area-sensitive loss that can better segment small objects even with noisy data.
- A 3D inpainting method with iterative view-consistent RGB-D inpainting and outperforms prior methods.

2. Related Work

3D Semantic Segmentation Many prior works have studied how to segment 3D point clouds into semantic classes. This work focuses on identifying specific types of object categories [9, 4, 1], and thus generally works well when label sets are small [18, 34, 24], but not so well when there are many possible classes and few examples are available for most in the training set [40]. In this paper, we address the more abstract tasks of identifying clutter *vs.* not, an attribute that does not directly align with object category labels.

RGB Image Inpainting Traditional single-view methods [14, 13, 2, 3] such as patch-based or diffusion-based work often achieve realistic results but lack high-level image understanding. Since the seminal work [36] that introduced adversarial training and encoder-decoder structure into this task, various techniques have been proposed. These include multi-stage generation [33, 39], gated convolution [51], recurrent methods [26, 52], *etc.* More recent work begins to focus on more challenging settings such as pluralistic inpainting [54, 45] and large hole filling [53, 27, 43]. Specifically, LaMa [43] uses fast Fourier Convolutions and extremely large masks for training. It achieves high-fidelity results with relatively low training and inference cost. Single image 3D photography aims to synthesize novel views using image-based rendering techniques [20, 42, 23]. In some sense, it can be regarded as filling in small holes in other views from a single image.

Depth, Multi-View, and Mesh Inpainting Recent deep learning-based depth-completion methods have made great success on both supervised [35, 17, 5, 50] and unsupervised [29, 48, 28] learning. In particular, NLSPN [35] performs non-local depth propagation using deformable convolutions [55]. Multi-view inpainting is commonly used as a technique for image-based rendering [44, 25, 49, 38, 32]. These methods and our method both involve 3D-aware image inpainting. However, we aim to obtain inpainted mesh, whereas multi-view inpainting aims to render new views, *i.e.*, the cross-frame consistency can be optimized as one goal but not guaranteed. The word of “inpainting” has been overloaded for shape completion [16, 41, 46, 31]. For instance, the example-based method [41] aims to fill in small holes from incomplete scans. These methods only focus on geometry. Surface reconstruction algorithms [37, 22] can only create surface between close enough points, therefore cannot generate new surfaces to fill big holes.

3D Scene Inpainting 3D scene inpainting aims to fill in plausible geometry (and texture) for missing parts of a 3D scene. Earlier work only completes geometries [11, 10], and recently 3D inpainting of both geometry and texture

has gained more interests [12, 21, 30, 47]. TSDF-based methods [12, 21] use one forward pass to predict the scene completion. They are good at recovering dominant structures (*e.g.* flat walls and floors) but often struggle with more challenging structures. Our mesh-based method instead iteratively refines the scene by disentangling geometry and texture. Concurrent to our method, two NeRF-based methods [30, 47] also employ iterative refinement for scene object removal. One key difference is how multi-view inconsistency is addressed. We design pruning and voting techniques to select consistent views; Weder *et al.* [47] learn a per-image confidence to re-weight the NeRF losses; SPIn-NeRF [30] does not select view and does not enforce consistency. Another difference is how depth is inpainted. While [30, 47] inpaint depths only from masked depths, we perform RGB-guided depth inpainting and better align the inpainted depths with inpainted RGB images.

3. Method

Given inputs of an RGB-D sequence and its reconstructed mesh of an indoor scene, our method outputs a 3D scene without clutter (as shown in Fig. 2). The first step is to predict the 3D clutter segmentation (Sec. 3.1). We then project each clutter segment into each camera view to mask out clutter regions in the color and depth images. This is followed by color image inpainting, the result of which guides depth completion (Sec. 3.2.1). Once all RGB-D frames are inpainted, we run consistency pruning and voting. If there remain missing regions from clutter removal, the consistency-checked inpaintings are sent back to image inpainting and depth completion to fill the remaining holes (Sec. 3.2.2). The loop continues until all missing regions are filled. As a final step, we run Poisson surface reconstruction [22] on the consistency-guaranteed inpaintings to obtain the final inpainted scene mesh (Sec. 3.2.3).

3.1. Clutter Segmentation

Our focus is to propose a new type of scene segmentation based on shared properties (*e.g.* across clutter) rather than common benchmark semantic categories. Therefore, a concrete clutter definition is crucial to test our method although the precise form of the definition can be flexible (discussed in supplementary). Once clutter is defined, we group all raw categories (549 for ScanNet [9]) into clutter/non-clutter/mixed. To reduce ambiguity, we exclude mixed categories and only use other categories as coarse supervision. Note that due to the heterogeneity of clutter, most clutter categories have fewer than 20 instances each and the total instance count of clutter is smaller than that of non-clutter.

3.1.1 Area-Sensitive Loss

3D segmentation models typically exploit Cross Entropy loss to supervise 3D predictions. However, directly applying the standard Cross Entropy loss does not train a model for our task with a performance similar to [18]. A unique feature of clutter compared to general objects is that a larger portion of clutter objects are small in size. However, a standard per-vertex loss focuses too much on larger instances. To draw the model’s attention back to small objects, we modify the standard Cross Entropy loss into the form:

$$L_{CE(3d)}^* = \frac{1}{N_v} \sum_{c \in C} \sum_{i \in O_c} \left(w_i \sum_{j \in V_i} y_{(c,j)} \log(p'_{(c,j)}) \right),$$

where N_v is the total number of vertices in the scene, C is the set of all classes (clutter and non-clutter), O_c is the set of instances of class c in the current scene, V_i is the set of vertices of the object i , $y_{(c,j)}$ and $p'_{(c,j)}$ are ground truth labels and predicted class probabilities for vertex j , respectively. In particular, w_i is our introduced weight to balance instance size and is formulated as:

$$w_i = \left(\frac{\text{median}(N)}{n_i} \right)^k, n_i \in N,$$

where N is a list of the vertex counts for all objects in the scene, n_i is the vertex count for object i , and $k \geq 0$ is a modulating factor to control how strongly the weight is applied. A smaller n_i (fewer vertices in object instance i) increases the weight w_i and loss for object i . A larger k increases w_i and loss penalty for smaller object instances. In particular, when $k = 0$ and $w_i = 1$, $L_{CE(3d)}^*$ takes the same form as Cross Entropy loss.

It is worth pointing out that our area-sensitive loss may look superficially similar to median frequency loss [15] or balanced Cross Entropy [19, 8]. However, we are fundamentally different. While other losses are aimed to address the class imbalance, our loss is on an instance level to better segment small objects – object instances with small surface area are weighted more regardless of the total surface area for the class. In particular, not all clutter instances are small (there are some easy and big-sized categories), and not all non-clutter instances are big. We expect our loss to be helpful for small objects across classes. A class-balanced loss in our binary segmentation task would simply be weighting one class uniformly larger. We found empirically that this does not improve the results of our task.

3.1.2 Implementation and Virtual View Rendering

Our focus is to design effective losses for heterogeneous small objects. We opt for BpNet [18] as the backbone 3D segmentation model. It contains 2D and 3D sub-networks connected by bidirectional feature projection, which takes

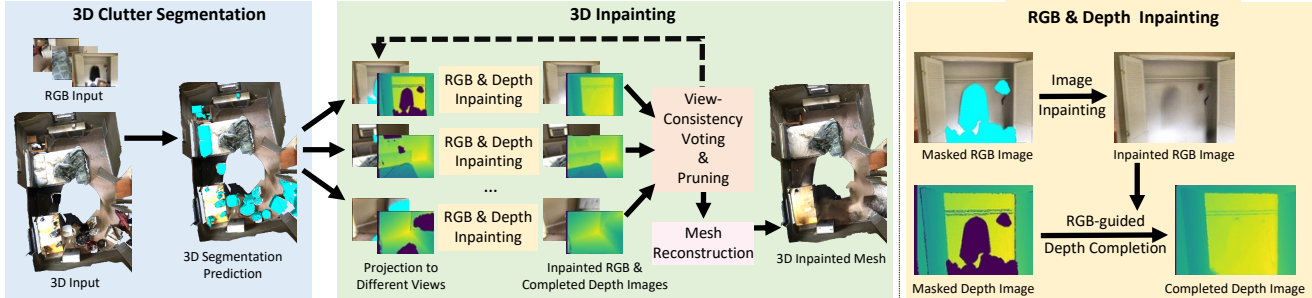


Figure 2. **Method Overview.** Given an RGB-D reconstruction and the RGB frames, we first predict 3D clutter segmentation. The prediction is then projected in different camera views to mask RGB and depth images. Next, we inpaint RGB image which we use to guide depth completion. Afterwards, we run consistency pruning and voting for all RGB-D frames. Any inpainting with remaining missing regions is sent back to RGB and depth inpainting again. The loop continues until all missing regions are filled. As last, we obtain the final inpainted scene mesh through mesh reconstruction.

as input color images and point clouds, respectively. The loss function $L_{seg} = L_{CE(3d)}^* + \lambda L_{CE(2d)}^*$ is a weighted sum of area-sensitive losses on 3D and 2D predictions, balanced by a weight λ . We follow [18] and input 3 views at a time. However, three originally captured images only cover a small surface of the entire scene, and increasing the number of input views drastically consumes the memory. Inspired by [24], we replace the originally captured images with virtual renderings for which we can set camera parameters to increase observed context.

3.2. 3D Inpainting

3D inpainting is a recently emerging task [12, 21], and one key challenge in 3D inpainting is the intricate interplay between geometry and semantics (or appearance). To address this issue, our key idea is to leverage intermediate 2D representations that disentangles textured 3D into RGB images (appearance) and geometry, which can be more easily inpainted and then used for multi-view 3D reconstruction. Notably, this novel 2D approach to the 3D inpainting problem deviates from prior work that relies on 3D convolutions for direct inpainting in 3D [12, 21]. In the following, we will first briefly overview the proposed pipeline before introducing new techniques to tackle the challenges.

3.2.1 Single-Frame Inpainting

We first project the clutter predictions from a 3D mesh onto each view to create a 2D mask for RGB and depth images. One benefit of our proposed 2D approach is to leverage the powerful 2D synthesis models. This also potentially improves generalization ability especially when faced with limited 3D data (compared to image data).

Color Image Inpainting We use LaMa [43] with released pre-trained weights to inpaint each color image. The model trained on diverse data generalizes well to indoor captured

images. However, in cases where multiple completions are plausible (especially for big masks), the model tends to fill in different contents for the same region from different views. This also leads to inconsistent depth completion and will be addressed in Sec. 3.2.2.

Depth Completion Existing depth completion methods [35, 17, 5, 50, 29, 48, 28] try to predict complete depth map from uniformly and sparsely sampled pixels, *e.g.*, by using local propagation. However, we need to complete dense maps with (potentially big) holes without local information. We adopt NLSPN [35] that performs non-local depth propagation for image-guided depth completion. Since there’s no annotated data for our task, we design a method that creates realistic holes on originally clean depth maps to generate supervision:

We assume m_1 is a coarse clutter mask of a captured depth map d_1 . We use the clutter mask m_2 from a second view to mask out regions on d_1 . In the meantime, we make sure that the regions originally masked as clutter in m_1 are not masked out. By copying and pasting m_2 from another view, we make sure that the masks during training have realistic boundaries of clutter objects. By not masking clutter from m_1 , we guarantee that the depth completion model does not need to hallucinate clutter objects out of the hole areas.

3.2.2 View-Consistency Inspired Refinement

The tremendous advances in image inpainting and depth completion have enabled high-quality inpainting of individual frames. However, due to the pluralistic nature of image inpainting and the accuracy limits of each model, inconsistency occurs across frames (Fig. 6). Such inconsistency will cause many noises in the mesh reconstruction. In this section, we discuss the unique characteristics of the task and

explain how we enforce consistency based on the observations.

Single-frame Consistency Pruning Inherent in the task of object removal, the inpainted depth d_{pre} should always be greater than the original captured depth d_{cap} at the same region. Here, the region includes each pixel p and each connected mask region i . Based on this observation, we have

$$d_{con}^1(p) = \begin{cases} d_{pre}(p) & \text{if } d_{cap}(p) < d_{pre}(p), \\ 0 & \text{else;} \end{cases}$$

$$d_{con}^2(i) = \begin{cases} d_{con}^1(i) & \text{if } d_{cap}(i) < d_{pre}(i), \\ 0 & \text{else,} \end{cases}$$

where d_{con}^1 and d_{con}^2 are the depth outputs from consistency checks for pixels and connected regions, respectively. Since our focus is to fill holes from object removal rather than the cracks in captured depth images caused by moving sensors, we mask out any pixel in d_{pre} that is empty in d_{cap} . We also drop inpainted regions whose area is over half the entire image, because inpaintings from limited input pixels are highly unreliable.

Cross-frame Consistency Pruning In the same spirit as space carving, the inpainted depth $d_{pre}^{(s)}$ from source view s , when warped to a target view t , cannot occlude captured depth in that view $d_{cap}^{(t)}$. This means at each pixel p :

$$d_{con}^{3(s)}(p) = \begin{cases} d_{con}^{2(s)}(p) & \text{if } d_{cap}^{(t)}(p) < \text{Warp}^{s-t}(d_{pre}^{(s)}(p)), \\ 0 & \text{else,} \end{cases}$$

where $d_{con}^{3(s)}$ is the depth output under view s from cross-frame consistency pruning. $\text{Warp}^{s-t}(\cdot)$ is the transformation from view s to view t . The transformation can be calculated given camera intrinsics and poses for two views. This is performed for each view pair (s, t) with $s, t \in \{0, \dots, N_s - 1\}$, where N_s is the sequence length.

Cross-frame Consistency Voting So far we have pruned some regions base on consistency to captured depths. However, there remains inconsistency among inpainted regions themselves. The causes include: at the same region from different views, (1) color images are inpainted with different contents or (2) completed depths are different even color image inpaintings are the same. To address this, we warp inpainted regions of all other views to view s using the predicted depth. Then we consider all warped depth values that fall into pixel p . The depth value at pixel p for view s is only valid if the count of projected depths within a distance

threshold α from $d_{pre}^{(s)}(p)$ is more than $\beta\%$ of the total count of projected depths falling in at p . Thus we have

$$d_{con}^{4(s)}(p) = \begin{cases} d_{con}^{3(s)}(p) & \text{if } r_\alpha > \beta\%, \\ 0 & \text{else,} \end{cases}$$

$$r_\alpha = \frac{\sum_t \mathbb{I}\left(|d_{pre}^{(s)}(p) - \text{Warp}^{t-s}(d_{pre}^{(t)}(p))| < \alpha\right)}{\sum_t \mathbb{I}\left(\text{Warp}^{t-s}(d_{pre}^{(t)}(p)) > 0\right)},$$

where $t \in \{0, \dots, N_s - 1\} \setminus \{s\}$, $d_{con}^{4(s)}$ is the depth output at view s from cross-frame consistency voting, and $\mathbb{I}(\cdot)$ is a binary indicator function with $\mathbb{I}(x) = 1$ if x is true else 0.

Iterative Refinement and Efficiency These consistency checks create new holes in $d_{con}^{4(s)}$. We mask inpainted color images and depth maps at the current step with these new holes and start over 3D inpainting with the new masked RGB-D images. The single-frame inpainting and consistency refinement are iteratively performed until all holes are filled. The complexity of the cross-frame pruning or voting is $O(N_s^2)$. We implement a batched warping function in PyTorch with GPU support, which takes about 7 minutes for a 200-frame sequence on an NVIDIA TITAN GPU.

3.2.3 Mesh Reconstruction

To obtain the final inpainted 3D scene, we run Poisson surface reconstruction [22] on the inpaintings from Sec. 3.2.2. We use a maximum of 16 million points and depth values smaller than 3.5 meters from the inpaintings. The maximum depth of the tree for surface reconstruction is set to 10.

4. Experiments

4.1. Experiment Setup

Dataset We use ScanNet dataset [9] with official split. Since clutter labels from grouped raw categories are very noisy, the original test set does not provide accurate evaluation. We manually label 8 representative test scenes with diverse settings using the annotation tool from [9]. Evaluations on this clean labeled set truly reflect performance. Please refer to supplementary materials for more details.

Implementation Details We set segmentation voxel size to 5 cm, with loss balancing weight λ and modulating factor k empirically set to 0.3 and 1 respectively.

We set $\alpha = 0.05$, $\beta\% = 30\%$ for the cross-frame consistency voting.

Table 1. **Comparison of Clutter Segmentation Methods.** We evaluate IoU and compare our method with the BPNet baseline.

Method	Original Test Set (noisy)			Manual Test Set (clean)		
	IoU(NC)	IoU(C)	mIoU	IoU(NC)	IoU(C)	mIoU
BPNet [18]	0.87	0.43	0.65	0.85	0.35	0.60
Ours	0.90	0.59	0.75	0.84	0.58	0.71

4.2. Baselines and Evaluation Metric

Clutter Segmentation We compare with a baseline that trains the original multi-class BPNet [18] for segmentation of 549 raw categories and then merges predictions based on our clutter/non-clutter grouping. All the other settings are the same as our model. We evaluate 3D mean IoU of both methods on two test sets using virtual views.

3D Inpainting We compare with the hole-filling algorithm “Close Holes” from MeshLab [7, 6] and Poisson surface reconstruction (PSR) [22] with the same setting as in 3.2.3. We further compare with two recent textured scene completion methods (SPSG [12] and FF [21]). Since they were both developed on Matterport3D dataset [4]. We also retrain our method on Matterport3D and compare with their official code and model weights.

4.3. Results

4.3.1 Full Pipeline

In Fig. 3, we show qualitative results for our entire pipeline of object removal and 3D inpainting, where we compare with baselines of 3D inpainting. The middle three columns show our predicted 3D clutter segmentation, inpainted texture and geometry. On the right, we compare with the geometry from Poisson surface reconstruction (PSR) [22] and a hole-filling algorithm.

The third column shows that our segmentation method not only accurately predicts bigger and regularly-shaped categories such as coffee tables (last row) and chairs, but also performs well on objects with irregular shapes and appear at very low frequency. We also observe that the segmentation baseline often miss small objects. Detecting these small objects is nontrivial due to the large annotation noise in the training set Fig. 1. Our method’s successful detection mostly benefits from our area-sensitive loss, which we will thoroughly study in Sec. 4.4.1.

The fourth column illustrates that our method can fill in the missing regions with consistent texture. In the fourth row, for example, dark colors from the clutter on the desk all disappear after inpainting. We attribute this to the powerful semantic understanding from image inpainting. However, there still exist inpainted texture artifacts, and a considerable number of them are caused by illumination. Taking the third row as an example, we can successfully segment the entire lamp but the light it casts stays, which is not ideal.

Table 2. **Comparison of 3D Inpainting Methods.** We quantitatively compare with baselines using both 2D and 3D metrics.

Method	L1(↓)	L2(↓)	PSNR(↑)	SSIM(↑)	LPIPS(↓)	CD(cm)(↓)
SPSG [6]	0.077	0.139	16.810	0.696	0.419	21.354
FF [22]	0.038	0.071	21.762	0.861	0.213	1.781
Ours	0.018	0.033	23.673	0.918	0.151	0.621

The fifth column shows that our method is better at closing holes caused by object removal and completing with geometrically plausible surfaces. For instance, in the fourth row, the desk on the left has a big missing region, caused by the chair’s occlusion. PSR fails to fill this large hole and the hole-filling algorithm fills the hole incorrectly by connecting the desktop with desk legs. In contrast, our method not only fills the entire hole but also recovers the missing part of the desktop. Similar phenomena can also be found in other scenes, *e.g.* pillow regions on the bed in the first row and on the sofa in the last row.

4.3.2 Segmentation

We report in Tab. 1 the IoU evaluations of the baseline and our method. It shows that our method outperforms the compared BPNet [18] with multi-category segmentation. This is not too surprising, since the compared model was originally applied for tasks with 20 categories [18] while there are 549 categories in our experiment. As noted by prior work [40], the large number of class categories results in severe class imbalance. In summary, the comparison illustrates that the problem of clutter segmentation cannot be solved with a straightforward fine-grained segmentation model. More analysis on the importance of each component introduced to segmentation can be found in Sec. 4.4.1.

4.3.3 3D inpainting

Qualitative Evaluation We compare with the most recent textured 3D scene inpainting methods in Fig. 4. Both methods use 3D convolutions to directly predict TSDF from incomplete input volume. They suffer from low resolution which is a typical weakness of methods using volumetric representation. In contrast, our method outputs high-resolution mesh. In the second column, SPSG [12] and FF [21] fail to complete the hole on the couch, whereas our method recovers color and geometry consistent with the input. Another advantage of our proposed 2D approach is robustness and better generalization. For example, the first column in Fig. 4 shows an outdoor scene example with large holes, which is rarely seen during training. However, since our 2D inpainting modules have seen more diverse training example, our method can also work well on uncommon examples.




Figure 3. **Results for automatic clutter removal and 3D inpainting.** Each row shows a rendered view of an input 3D scene, together with our clutter segmentation and a rendered view of the scene with inpainted color and geometry. We compare the latter to removing clutter and filling the resulting holes with either Poisson Surface Reconstruction (PSR) or a triangulation-based hole-filling algorithm.

Quantitative Evaluation To quantitatively compare our method with baselines, we create a dataset with synthetically created holes (details in the supplementary). To evaluate the inpainted mesh texture, we report L1, L2, PSNR, SSIM, and LPIPS on the renderings of inpainted meshes, following common practice in image inpainting [51, 54]. To evaluate inpainted mesh geometry, we compare the Chamfer Distance between the original surface and the reconstructed surface. We report the evaluation results in Tab. 2. We can see our proposed method outperforms baselines on all metrics. This indicates that our method produces better texture and geometry, which has been shown in Fig. 3 and Fig. 4.

4.4. Ablation Studies

4.4.1 Clutter Segmentation

In Tab. 3, we investigate the effect of each strategy introduced in Sec. 3.1. For both original and manual test sets, we report IoU for non-clutter (NC) and clutter (C) classes as well as precision and recall for clutter. We will focus on the manual test set which has more accurate clutter annotations.

Table 3. **Ablations on Clutter Segmentation.** We study the effect of our area-sensitive loss with varying k . We also compare the influence of using rendered(r ./)captured(c .) 2D views and new(n ./)original(o .) camera parameters for training. (a-e) are evaluated on $r./n$. views, and (f-h) are evaluated on $c./o$. views. Best/second-best/third-best results among (a-e) are **bold**, solid underlined, and dashed underlined, respectively.

ID	k	2D Type	Original Test Set (noisy)				Manual Test Set (clean)			
			/Camera	IoU(NC)	IoU(C)	Pre. Rec.	IoU(NC)	IoU(C)	Pre. Rec.	
(a)	0	$r./n$.	0.94	0.67	0.85	<u>0.76</u>	0.83	0.48	0.91	0.50
(b)	0.5	$r./n$.	<u>0.93</u>	<u>0.65</u>	<u>0.79</u>	<u>0.79</u>	0.85	0.55	<u>0.88</u>	0.59
(c)	1.0	$r./n$.	<u>0.90</u>	<u>0.59</u>	<u>0.67</u>	0.83	<u>0.84</u>	<u>0.58</u>	<u>0.82</u>	<u>0.66</u>
(d)	1.5	$r./n$.	0.86	0.51	<u>0.57</u>	0.83	0.83	0.59	<u>0.76</u>	0.73
(e)	2.0	$r./n$.	0.79	0.39	0.45	0.75	0.78	0.51	0.65	<u>0.70</u>
(f)	1.0	$c./o$.	<u>0.92</u>	0.62	<u>0.76</u>	0.77	0.84	0.52	<u>0.88</u>	<u>0.55</u>
(g)	1.0	$r./o$.	0.91	0.62	0.73	0.80	0.84	0.55	0.86	0.61
(h)	1.0	$r./n$.	0.91	0.61	0.70	0.83	0.85	0.59	0.85	0.66

We first study the effects of *area-sensitive loss* by varying k from 0 to 2 (a-e). All (a-e) are trained and tested on rendered views from sampled cameras. On the manual test set, we see by increasing k from 0 to 1.5, IoU(C) improves by over 0.1 while IoU(NC) only fluctuates slightly. The asymmetrical change is due to the unbalanced size distribution for the two classes. Given the difficulty of segmenting




Figure 4. **3D Inpainting Comparison.** We compare our method with two 3D scene inpainting (completion) baselines (SPSG [12] and FF [21]). The first row shows input incomplete scenes (holes created by dropping clutter objects onto the scene), followed by three methods’ outputs and the GT scene without synthetic holes at the bottom. While prior work suffers from low resolution, our approach can keep the same resolution as input and is better at completing larger holes and rarer scenes.




Figure 5. **Visualizations of the impact of size-balancing parameter k on clutter detection.** Areas marked in cyan are true positives, magenta are false positives, and yellow are false negatives.

small objects with very noisy labels, the substantial increase of clutter IoU from (a) to (d) proves the effectiveness of our area-sensitive loss. To better explain the increase, we further look into the precision and recall for clutter. As shown in the last two columns of Tab. 3 (a-d), a larger k (within some range 0-1.5) significantly increases the recall with a falling precision. When k goes up to 2 (Tab. 3 (e)), all metrics start to fall due to disproportionately large weights.

The above trend in (a-d) is reflected in Fig. 5. An increased recall corresponds to the enlarged cyan (true positive) and shrunk yellow (false negative) regions. The larger

pink (false positive) regions imply a drop in precision. With $k = 1$, the model can segment the tiny clutter on the desk and fully segment the pillow on the top right, which is partially missed when $k = 0$. The lower precision is mostly caused by false positives (pink) around clutter. This is not as harmful as expected since we originally also dilate predicted masks to remove shadows. We want to note that the currently missed predictions are very hard cases. For example, the big box on the left is very similar to a small cabinet and the raw category *box* itself is a mix of clutter and non-clutter. The two small yellow missing spots on the top right are two small plug heads, and their surrounding mesh surface is very dark to recognize the clutter.

In the original test set, we can observe both IoU(C) and IOU(NC) drop with increasing k . Comparing the changes of IoU(NC) in both test sets from (a) to (d), a large drop (0.94 to 0.86) is observed in the original set while the value stays the same in the manual set (0.83 to 0.83). This is because by increasing the area-sensitive strength through k , more clutter objects are recognized, which can be mislabeled as non-clutter in the original test set. This also verifies our previous claim (Sec. 1) that existing datasets have very low-quality clutter labels. Despite the low-quality label, we also observe the trend of increased recall and declined precision similar to that of the manual test set.

We compare the effect of using rendered or captured images in Tab. 3 (c, f, g, h). All (f-h) experiments are tested on captured images. We can see that (g) trained on virtual renderings with captured camera parameters performs on par with (f) trained on original captured images. This shows that the 2D domain gap does not affect the performance. From (g) to (h), we see that the larger variations of camera pose (h) improve over using only poses on the hand-held sensor trajectories (g). Comparing (c) and (h) that use the same model trained on $r./n$. views, we see that the model trained with virtual renderings performs better when tested on original captures (h) than on its virtual renderings (c). This is probably because testing on captured images is easier since the cameras all fall into a single trajectory.

4.4.2 3D Inpainting

Despite image inpainters’ great power of inpainting semantically plausible content, consistency across views is never guaranteed. In this section, we study the effect of applying different consistency rules, visualized in Fig. 6.

For a single frame, while the inpainted content can look plausible with the remaining pixels, the inpainted image may not conform to the original RGB image. In row (a), for instance, although the inpainted RGB completes a plausible table, these regions originally occluded by chairs should instead be the floor. If we directly reconstruct a mesh from a sequence where occluded views are inpainted differently,




Figure 6. **Visualizations of the necessity and effect of consistency rules.** Original images (first column) are independently inpainted (second and third column), frequently leading to poor merged results (fourth column). Our checks result in more plausible reconstructions (fifth column). The three rows illustrate the effects of (a) single-frame pruning, (b) cross-frame pruning, and (c) cross-frame voting.

we will get a mesh in the fourth column. A simple yet effective solution is to discard frames where the inpainted regions have smaller depth values than the original frames (single-frame consistency in Sec. 3.2.2).

Not only RGB inpainter, but inaccurate depth predictions can also cause artifacts in the reconstructed mesh. Row (b) demonstrates a case where the depth completion for the space between the desk and the wall can easily create noise. The mesh in the fourth column is what we will get without correction. We know this mesh is wrong because, from a view similar to the third row, we should be able to see the walls and floors. Such a cross-frame prior leads us to cross-frame consistency pruning: any non-inpainted regions visible in one frame should also be visible in another frame. This idea can effectively remove many artifacts and return a much cleaner mesh as shown in the fifth column.

So far, we have applied all the consistency assumptions that leverage captured images. The last step is consistency check across inpainted regions in different views. For example, the dark inpainted regions under the table can easily cause an inaccurate depth estimation. This results in the cluttered artifacts in the fourth column. By applying the cross-frame voting from Sec. 3.2.2, regions with no clustered depths from different views will be dropped.

5. Conclusion

This paper studies the problem of detecting and removing clutter from 3D scenes, while providing plausible fillings. We started by defining clutter as frequently-moving objects and highlighting the issue of highly noisy annotations in existing 3D scene segmentation datasets. Then we presented a solution for clutter segmentation and 3D inpainting. To better segment clutter objects with heterogeneous shapes, we developed an instance-level area-sensitive loss that significantly improves recall, particularly for small

objects. Our 3D inpainting method is grounded in the key principle of decomposing the task into geometry and color (semantics). To achieve this, we project 3D segmentation onto RGB-D images to complete the color and depth in the 2D space. These 2D inpaintings are then used for mesh reconstruction after view-consistency voting and pruning. To our knowledge, this is the first work that investigates the combined problem of clutter segmentation and inpainting in 3D. While our method compares favorably with baseline methods, there remain several open problems. For instance, removing shadows before inpainting and approaches that solve this problem directly in 3D space are both intriguing avenues for further inquiry.

Acknowledgement We would like to thank Intel and Google for partial funding of this work.

References

- [1] Iro Armeni, Ozan Sener, Amir R Zamir, Helen Jiang, Ioannis Brilakis, Martin Fischer, and Silvio Savarese. 3d semantic parsing of large-scale indoor spaces. In *CVPR*, pages 1534–1543, 2016. 1, 2
- [2] Coloma Ballester, Marcelo Bertalmio, Vicent Caselles, Guillermo Sapiro, and Joan Verdera. Filling-in by joint interpolation of vector fields and gray levels. *IEEE transactions on image processing*, 10(8):1200–1211, 2001. 2
- [3] Connelly Barnes, Eli Shechtman, Adam Finkelstein, and Dan B Goldman. Patchmatch: A randomized correspondence algorithm for structural image editing. *ACM Transactions on Graphics (TOG)*, 28(3):24, 2009. 2
- [4] Angel Chang, Angela Dai, Thomas Funkhouser, Maciej Halber, Matthias Niessner, Manolis Savva, Shuran Song, Andy Zeng, and Yinda Zhang. Matterport3d: Learning from rgb-d data in indoor environments. *arXiv preprint arXiv:1709.06158*, 2017. 1, 2, 6
- [5] Xinjing Cheng, Peng Wang, Chenye Guan, and Ruigang Yang. Cspn++: Learning context and resource aware convolutional spatial propagation networks for depth completion. In *AAAI*, volume 34, pages 10615–10622, 2020. 2, 4
- [6] Paolo Cignoni, Marco Callieri, Massimiliano Corsini, Matteo Dellepiane, Fabio Ganovelli, and Guido Ranzuglia. MeshLab: an Open-Source Mesh Processing Tool. In Vittorio Scarano, Rosario De Chiara, and Ugo Erra, editors, *Eurographics Italian Chapter Conference*. The Eurographics Association, 2008. 6
- [7] Paolo Cignoni, Alessandro Muntoni, Guido Ranzuglia, and Marco Callieri. MeshLab. 6
- [8] Yin Cui, Menglin Jia, Tsung-Yi Lin, Yang Song, and Serge Belongie. Class-balanced loss based on effective number of samples. In *CVPR*, pages 9268–9277, 2019. 3
- [9] Angela Dai, Angel X Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *CVPR*, pages 5828–5839, 2017. 1, 2, 3, 5

- [10] Angela Dai, Christian Diller, and Matthias Nießner. Sg-nn: Sparse generative neural networks for self-supervised scene completion of rgb-d scans. In *CVPR*, 2020. 2
- [11] 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 *CVPR*, 2018. 2
- [12] Angela Dai, Yawar Siddiqui, Justus Thies, Julien Valentin, and Matthias Nießner. Spsg: Self-supervised photometric scene generation from rgb-d scans. In *CVPR*, pages 1747–1756, 2021. 2, 3, 4, 6, 8
- [13] Alexei A Efros and William T Freeman. Image quilting for texture synthesis and transfer. In *Proceedings of SIGGRAPH*, pages 341–346, 2001. 2
- [14] Alexei A Efros and Thomas K Leung. Texture synthesis by non-parametric sampling. In *ICCV*, volume 2, pages 1033–1038. IEEE, 1999. 2
- [15] David Eigen and Rob Fergus. Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture. In *ICCV*, pages 2650–2658, 2015. 3
- [16] Youichi Horry, Ken-Ichi Anjyo, and Kiyoshi Arai. Tour into the picture: using a spidery mesh interface to make animation from a single image. In *Proceedings of SIGGRAPH*, pages 225–232, 1997. 2
- [17] Mu Hu, Shuling Wang, Bin Li, Shiyu Ning, Li Fan, and Xiaojin Gong. Penet: Towards precise and efficient image guided depth completion. In *ICRA*, pages 13656–13662. IEEE, 2021. 2, 4
- [18] Wenbo Hu, Hengshuang Zhao, Li Jiang, Jiaya Jia, and Tien-Tsin Wong. Bidirectional projection network for cross dimensional scene understanding. In *CVPR*, 2021. 1, 2, 3, 4, 6
- [19] Shruti Jadon. A survey of loss functions for semantic segmentation. In *CIBCB*, pages 1–7. IEEE, 2020. 3
- [20] Varun Jampani, Huiwen Chang, Kyle Sargent, Abhishek Kar, Richard Tucker, Michael Krainin, Dominik Kaeser, William T Freeman, David Salesin, Brian Curless, et al. Slide: Single image 3d photography with soft layering and depth-aware inpainting. In *ICCV*, pages 12518–12527, 2021. 2
- [21] Ru-Fen Jheng, Tsung-Han Wu, Jia-Fong Yeh, and Winston H Hsu. Free-form 3d scene inpainting with dual-stream gan. *BMVC*, 2022. 2, 3, 4, 6, 8
- [22] Michael Kazhdan, Matthew Bolitho, and Hugues Hoppe. Poisson surface reconstruction. In *Proceedings of the fourth Eurographics symposium on Geometry processing*, volume 7, 2006. 2, 3, 5, 6
- [23] Johannes Kopf, Kevin Matzen, Suhib Alsisan, Ocean Quigley, Francis Ge, Yangming Chong, Josh Patterson, Jan-Michael Frahm, Shu Wu, Matthew Yu, et al. One shot 3d photography. *ACM Transactions on Graphics (TOG)*, 39(4):76–1, 2020. 2
- [24] Abhijit Kundu, Xiaoqi Yin, Alireza Fathi, David Ross, Brian Brewington, Thomas Funkhouser, and Caroline Pantofaru. Virtual multi-view fusion for 3d semantic segmentation. In *ECCV*, pages 518–535. Springer, 2020. 1, 2, 4
- [25] Feiran Li, Gustavo Alfonso Garcia Ricardez, Jun Takamatsu, and Tsukasa Ogasawara. Multi-view inpainting for rgb-d sequence. In *3DV*, pages 464–473. IEEE, 2018. 2
- [26] Jingyuan Li, Ning Wang, Lefei Zhang, Bo Du, and Dacheng Tao. Recurrent feature reasoning for image inpainting. In *CVPR*, pages 7760–7768, 2020. 2
- [27] Wenbo Li, Zhe Lin, Kun Zhou, Lu Qi, Yi Wang, and Jiaya Jia. Mat: Mask-aware transformer for large hole image inpainting. In *CVPR*, pages 10758–10768, 2022. 2
- [28] Adrian Lopez-Rodriguez, Benjamin Busam, and Krystian Mikolajczyk. Project to adapt: Domain adaptation for depth completion from noisy and sparse sensor data. In *ACCV*, 2020. 2, 4
- [29] Fangchang Ma, Guilherme Venturelli Cavalheiro, and Sertac Karaman. Self-supervised sparse-to-dense: Self-supervised depth completion from lidar and monocular camera. In *ICRA*, pages 3288–3295. IEEE, 2019. 2, 4
- [30] Ashkan Mirzaei, Tristan Aumentado-Armstrong, Konstantinos G Derpanis, Jonathan Kelly, Marcus A Brubaker, Igor Gilitschenski, and Alex Levinshtein. Spin-nerf: Multiview segmentation and perceptual inpainting with neural radiance fields. In *CVPR*, pages 20669–20679, 2023. 3
- [31] Himangi Mittal, Brian Okorn, Arpit Jangid, and David Held. Self-supervised point cloud completion via inpainting. *arXiv preprint arXiv:2111.10701*, 2021. 2
- [32] Shohei Mori, Jan Herling, Wolfgang Broll, Norihiko Kawai, Hideo Saito, Dieter Schmalstieg, and Denis Kalkofen. 3d pixmix: Image inpainting in 3d environments. In *2018 IEEE International Symposium on Mixed and Augmented Reality Adjunct (ISMAR-Adjunct)*, pages 1–2. IEEE, 2018. 2
- [33] Kamyar Nazeri, Eric Ng, Tony Joseph, Faisal Z Qureshi, and Mehran Ebrahimi. Edgeconnect: Generative image inpainting with adversarial edge learning. *arXiv preprint arXiv:1901.00212*, 2019. 2
- [34] Alexey Nekrasov, Jonas Schult, Or Litany, Bastian Leibe, and Francis Engelmann. Mix3d: Out-of-context data augmentation for 3d scenes. In *3DV*, pages 116–125. IEEE, 2021. 1, 2
- [35] Jinsun Park, Kyungdon Joo, Zhe Hu, Chi-Kuei Liu, and In So Kweon. Non-local spatial propagation network for depth completion. In *ECCV*, 2020. 2, 4
- [36] Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A Efros. Context encoders: Feature learning by inpainting. In *CVPR*, pages 2536–2544, 2016. 2
- [37] Songyou Peng, Chiyu "Max" Jiang, Yiyi Liao, Michael Niemeyer, Marc Pollefeys, and Andreas Geiger. Shape as points: A differentiable poisson solver. In *NeurIPS*, 2021. 2
- [38] Julien Philip and George Drettakis. Plane-based multi-view inpainting for image-based rendering in large scenes. In *Proceedings of the ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games*, pages 1–11, 2018. 2
- [39] Yurui Ren, Xiaoming Yu, Ruonan Zhang, Thomas H Li, Shan Liu, and Ge Li. Structureflow: Image inpainting via structure-aware appearance flow. In *ICCV*, pages 181–190, 2019. 2

- [40] David Rozenberszki, Or Litany, and Angela Dai. Language-grounded indoor 3d semantic segmentation in the wild. In *ECCV*, 2022. 1, 2, 6
- [41] Shankar Setty and Uma Mudenagudi. Example-based 3d inpainting of point clouds using metric tensor and christoffel symbols. *Machine Vision and Applications*, 29(2):329–343, 2018. 2
- [42] Meng-Li Shih, Shih-Yang Su, Johannes Kopf, and Jia-Bin Huang. 3d photography using context-aware layered depth inpainting. In *CVPR*, pages 8028–8038, 2020. 2
- [43] Roman Suvorov, Elizaveta Logacheva, Anton Mashikhin, Anastasia Remizova, Arsenii Ashukha, Aleksei Silvestrov, Naejin Kong, Harshith Goka, Kiwoong Park, and Victor Lempitsky. Resolution-robust large mask inpainting with fourier convolutions. In *WACV*, pages 2149–2159, 2022. 2, 4
- [44] Theo Thonat, Eli Shechtman, Sylvain Paris, and George Drettakis. Multi-view inpainting for image-based scene editing and rendering. In *3DV*, pages 351–359. IEEE, 2016. 2
- [45] Ziyu Wan, Jingbo Zhang, Dongdong Chen, and Jing Liao. High-fidelity pluralistic image completion with transformers. In *ICCV*, pages 4692–4701, 2021. 2
- [46] Xinying Wang, Dikai Xu, and Fangming Gu. 3d model inpainting based on 3d deep convolutional generative adversarial network. *IEEE Access*, 8:170355–170363, 2020. 2
- [47] Silvan Weder, Guillermo Garcia-Hernando, Aron Monszpart, Marc Pollefeys, Gabriel J Brostow, Michael Firman, and Sara Vicente. Removing objects from neural radiance fields. In *CVPR*, pages 16528–16538, 2023. 3
- [48] Alex Wong and Stefano Soatto. Unsupervised depth completion with calibrated backprojection layers. In *ICCV*, pages 12747–12756, 2021. 2, 4
- [49] Hailong Yan, Wenqi Wu, Zhenghua Deng, Junjian Huang, Zhizhang Li, and Luting Zhang. Image inpainting for 3d reconstruction based on the known region boundaries. *Mathematics*, 10(15):2761, 2022. 2
- [50] Yanchao Yang, Alex Wong, and Stefano Soatto. Dense depth posterior (ddp) from single image and sparse range. In *CVPR*, pages 3353–3362, 2019. 2, 4
- [51] Jiahui Yu, Zhe Lin, Jimei Yang, Xiaohui Shen, Xin Lu, and Thomas S Huang. Free-form image inpainting with gated convolution. In *ICCV*, pages 4471–4480, 2019. 2, 7
- [52] Haoran Zhang, Zhenzhen Hu, Changzhi Luo, Wangmeng Zuo, and Meng Wang. Semantic image inpainting with progressive generative networks. In *Proceedings of the 26th ACM international conference on Multimedia*, pages 1939–1947, 2018. 2
- [53] Shengyu Zhao, Jonathan Cui, Yilun Sheng, Yue Dong, Xiao Liang, Eric I Chang, and Yan Xu. Large scale image completion via co-modulated generative adversarial networks. *arXiv preprint arXiv:2103.10428*, 2021. 2
- [54] Chuanxia Zheng, Tat-Jen Cham, and Jianfei Cai. Pluralistic image completion. In *CVPR*, pages 1438–1447, 2019. 2, 7
- [55] Xizhou Zhu, Han Hu, Stephen Lin, and Jifeng Dai. Deformable convnets v2: More deformable, better results. In *CVPR*, pages 9308–9316, 2019. 2