Interactive Object Segmentation in 3D Point Clouds

Theodora Kontogianni^{1,2}, Ekin Celikkan³, Siyu Tang¹, Konrad Schindler¹

Abstract— We propose an interactive approach for 3D instance segmentation, where users can iteratively collaborate with a deep learning model to segment objects directly in a 3D point cloud. Current methods for 3D instance segmentation are generally trained in a fully-supervised fashion, which requires large amounts of costly training labels, and does not generalize well to classes unseen during training. Few works have attempted to obtain 3D segmentation masks using human interactions. Existing methods rely on user feedback in the 2D image domain. As a consequence, users are required to constantly switch between 2D images and 3D representations, and custom architectures are employed to combine multiple input modalities. Therefore, integration with existing standard 3D models is not straightforward. The core idea of this work is to enable users to interact directly with 3D point clouds by clicking on desired 3D objects of interest (or their background) to interactively segment the scene in an open-world setting. Specifically, our method does not require training data from any target domain and can adapt to new environments where no appropriate training sets are available. Our system continuously adjusts the object segmentation based on the user feedback and achieves accurate dense 3D segmentation masks with minimal human effort (few clicks per object). Besides its potential for efficient labeling of large-scale and varied 3D datasets, our approach, where the user directly interacts with the 3D environment, enables new AR/VR and human-robot interaction applications.

I. INTRODUCTION

Just like other occurrences of supervised deep learning, 3D scene understanding tasks such as 3D semantic instance segmentation require large amounts of annotated training data. Alas, in 3D, dense manual annotation is even more time-consuming and tedious. Interactive object segmentation techniques are commonly used in the 2D image domain to collect ground-truth segmentation masks at scale [3] or on image editing applications. However, interactive object segmentation in the 3D domain is mainly under-explored.

The principle of interactive object segmentation [3, 16, 20, 22, 25, 35] is to let a user collaborate with a model when segmenting an object instance. By concentrating on one instance at a time (*e.g.*, a specific chair), the task becomes a binary foreground/background segmentation. The segmentation mask is obtained with an iterative procedure: the model produces its best guess of a mask, then the user provides feedback via corrections, which the model ingests to produce an updated mask. That process is repeated until the mask is deemed accurate enough by the user (Fig.2). The corrections usually come in the form of clicks that identify locations with incorrect labels. *Negative* clicks identify false positives, where parts of the background are labeled as

¹ETH Zürich, Switzerland

³Computer Vision Group, RWTH Aachen University, Germany



Fig. 1. We believe that high quality segmentation masks, that can for example be used as ground truth, should be at least above 90% IoU. This performance cannot yet be reached, even with fully-supervised instance segmentation methods (Tab. I). As shown above, a segmentation mask (purple) with 80% IoU still has significant erroneous regions. This motivates our approach for *interactive* 3D instance segmentation, where a user can significantly increase the segmentation quality, with just a few clicks.



Fig. 2. Interactive 3D Object Segmentation. The model predicts a segmentation mask and iteratively updates it, driven by user feedback in the form of positive clicks on false negatives, and negative clicks on false positives. The current click is visualized as enlarged circle.

foreground. *Positive* clicks identify false negatives where parts of the object have been falsely labeled as background.

Even though interactive segmentation in images is a wellestablished technology, and part of several consumer software products, little is known about its potential for labeling 3D point clouds. Despite the recent progress of 3D deep learning, and the persistent dearth of realistic large-scale training sets, interactive 3D object segmentation remains under-explored.

The field of 3D deep learning has witnessed significant progress in the past few years, particularly regarding semantic and instance segmentation in 3D point clouds [14, 17, 24, 33, 36, 38]. Nonetheless, existing instance segmentation methods are not suitable for labeling new datasets. They require lots of training data, to begin with. More importantly, they cannot generalize to classes that were not part of the training set, a situation commonly encountered when labeling new datasets. What is more, the resulting segmentation masks tend to have high, but imperfect accuracy (<90% IoU), which is impressive for an AI system, but not enough to serve as high-quality ground truth. We argue that even segmentation masks of 80% IoU are sub-optimal in representing a 3D object accurately (see Fig. 1). These masks are too inaccurate to represent ground truth. Note, though, that our proposed interactive method can be integrated with most

²ETH AI Center

existing instance segmentation algorithms to obtain improved segmentation masks with little user input.

So far, there have been few attempts to make 3D object segmentation more efficient via user interaction [28, 39]. In these methods, user interaction takes place in the 2D image domain, *i.e.*, they additionally require images whose camera pose w.r.t. the point cloud is known, and associated 2D segmentation masks. In [28] the user interacts both with 2D images and 3D point clouds (by dragging mesh vertices), forcing the user to switch between 2D and 3D views of the same object. Due to the transition between the 2D and 3D domains, both these methods employ custom network architectures, therefore their integration with existing other backbones [6, 26, 31] is not straightforward. In contrast, our method can be integrated into more powerful future backbones or lighter ones for AR/VR and mobile applications.

Here, we propose an approach where the user interacts directly with the 3D point cloud. We argue that this direct approach has several advantages: (i) it obviates the need to supply images and their poses relative to the point cloud; (ii) it relieves the user of going back and forth between 2D and 3D interfaces, and (iii) it can be incorporated into any (point-based or voxel-based) 3D segmentation network. In our experiments, we rely on MinkowskiEngine [6], proven state-of-the-art architecture based on a sparse 3D voxel grid.

A main motivation for interactive segmentation is the ability to train algorithms for new domains where no appropriate training sets are available. Our method supports this, as it is general and not tied to a specific scan pattern or class nomenclature. We train on a single dataset, ScanNet [7], and evaluate on multiple others, without retraining.

In summary, we make the following contributions:

- 1) First state-of-the-art for interactive object segmentation purely on 3D point data, without extra images.
- We study the simulation of user clicks for training purposes and introduce a practical policy to synthetically generate meaningful clicks.
- 3) We experimentally show that interactive 3D object segmentation can provide accurate instance masks across various datasets, with little human effort.
- We show that the interactive 3D segmentation network generalizes well to different (indoor/outdoor) datasets and unseen object categories.
- Additionally, we have also created an interface that allows the user to pick single 3D points and used it to perform a real user study.

II. RELATED WORK

Fully-Supervised 3D Instance Segmentation. Instance segmentation on point clouds is a fundamental task in 3D scene perception and it has been thoroughly studied in the past few years [4, 9, 13, 18, 21, 34, 37]. Roughly the methods can be separated into two major categories: (1) *proposal-based* and (2) *clustering-based*. Proposal-based approaches [9, 13, 37] directly generate object proposals (inspired by Mask R-CNN [12] in 2D) and predict instance segmentation masks inside each proposal. Clustering-based methods learn

per-point features in an embedding space and then perform clustering to generate instances [4, 18, 21, 34]. Although these methods have achieved remarkable results on existing datasets, they require large amounts of labeled data for training and cannot generalize to classes that are not part of the training set. Our proposed interactive method is orthogonal to most instance segmentation methods and can be combined with them to improve segmentation masks with little user input.

Weakly-Supervised learning on 3D point clouds. Compared to fully supervised methods, weakly supervised methods try to learn the task at hand with limited training examples or labels. Recently, several methods [14, 17, 24, 33, 36, 38] started to tackle 3D semantic segmentation in a weakly supervised manner aiming to achieve the performance of their fully supervised counterparts using fewer point labels to train neural networks. Hou et al. [15] uses contrastive learning to learn meaningful feature representations that can be fine-tuned with fewer labels. Liu et al. [24] proposed a self-training approach that uses a graph propagation module that iteratively propagates the sparse labels for semantic segmentation. Xu et al. [36] performs point cloud semantic segmentation by exploiting spatial and color smoothness constraints along with gradient approximation. Chibane et al. [5] learn dense instance masks from coarse bounding box labels.

While achieving impressive results, these works focus on segmenting 3D point clouds with limited labels and sometimes require additional post-processing steps to separate object instances. Additionally, their goal is to reach the performance of fully supervised methods using fewer labels and they do not specifically target highly accurate segmentation masks. Even though weakly-supervised methods utilize fewer points, they still require individually trained models for every dataset. This is unlike interactive object segmentation which provides masks for single objects enabling fast instance annotation using only a single trained model.

Interactive 3D Object Segmentation. There is limited work on interactive frameworks for annotating 3D objects. Shen et al. [28] propose a framework for interactive 3D object annotation using interactions in both the image and the 3D domain. First, the user provides feedback on large errors using scribbles in 2D views of the 3D space. Then, the user corrects minor errors by manipulating mesh vertices. The user needs to label the same object by switching between two different modalities and that can be tiring for large-scale labeling. Additionally, the method requires a specialized architecture. [32] also relies on handcrafted features and a custom architecture to allow users to label a scene during capture. In contrast, our method allows users to label points directly in 3D space and can extend any semantic or instance segmentation network for 3D point clouds. Its results shall improve when used with more powerful future architectures.

III. METHOD

We start by formally defining the problem. Consider a 3D scene $P \in \mathbb{R}^{N \times C}$ where N is the number of 3D points in the scene and C the dimensionality of the features associated with each 3D point $p \in P$ (if only xyz coordinates are used:



Fig. 3. Interactive 3D Object Segmentation. We introduce interactive object segmentation of 3D point clouds where a user collaborates hand-in-hand with a trained deep model. *Left:* A user provides sparse annotations by iteratively clicking on each object one-by-one (each color indicates a new object instance). *Middle:* Given the user clicks, the model provides instance masks that the user can then further refine. *Right:* Ground-truth for reference (Colors are picked at random so they do not necessarily match the prediction colors).

C = 3). In interactive 3D object segmentation, a model allows the user to segment the desired region in the scene by successively placing positive (*foreground*) or negative (*background*) feedback in the form of clicks on a 3D point. Every time a new click is placed, the model outputs a new, updated segmentation. Once the user is satisfied with the mask of the intended object, the process concludes (Fig. 2).

A. Input representation.

The input to the interactive object segmentation network consists of a 3D scene $P \in \mathbb{R}^{N \times C}$ plus two additional channels T_{pos} and T_{neg} for the positive and negative clicks respectively. Next, we describe these additional channels.

Click Encoding Channels. To acquire an accurate segmentation mask for the object of interest, the user provides a sequence of positive and negative clicks. The positive clicks are considered on the object, and the negative clicks are on the background. These clicks are represented by their 3D coordinates in the 3D scene. S_{pos} is a set of all positive user clicks, and S_{neg} is the set of all negative ones. Benenson *et al.* [3] have done an extensive study on 2D clicks encodings and find that disks of a small radius outperform the other encodings (*e.g.* distance transform) since their effects are more locally restricted [29]. For those reasons, we encode the clicks as 3D volumes. The two additional input channels T_{pos} and T_{neg} are defined as:

$$T_{pos}(p) = \begin{cases} 1, & \|\mathbf{x}_p - \mathbf{x}_q\|_{\infty} < \epsilon \\ 0, & \text{otherwise} \end{cases}$$
(1)

where \mathbf{x}_p the 3D coordinates of any point p in the point cloud, \mathbf{x}_q the coordinates of point $q \in S_{pos}$, and ε is the volume length. We define T_{neg} similarly but with $q \in S_{neg}$. So each user click channel is a binary mask of size $N \times 1$, where N is the number of points, and the mask is one for each point inside the click-volume, and zero outside. Then we concatenate the C channels of the 3D scene with T_{pos} and T_{neg} to compose a $N \times (C + 2)$ input.

Simulating User Clicks during Training. In practice, it is impractical and not realistic to collect interactive clicks from real users during the training of the model. Therefore, we simulate user clicks for training using random sampling as is common practice in the 2D domain [16, 20, 22, 25, 35]. We sample *positive* clicks uniformly at random on the target 3D object. *Negative* clicks are randomly sampled from a volume

around the object, up to 1.4 times bigger than the object to ensure hard negative examples from nearby objects.

Simulating User Clicks during Test. To ensure that the evaluation is unbiased from real user errors and users different skills, we formalize an evaluation protocol on simulated user test clicks. This also ensures reproducible scores. At test time, we add clicks one by one based on the errors of the currently predicted mask. We imitate a user who always clicks at the center of the region with the largest error. To obtain this error region, we need to compute pair-wise distances between all falsely labeled 3D points. To reduce memory and computational load, it is crucial to perform this operation on a sparse voxelization of the 3D scene (1 point/5 cm voxel).

Online Adaptation. Interactive object segmentation can alleviate the issue of poor generalization on unknown classes and different datasets as shown in our results on different datasets (Tab. IV). Importantly, the generalization can be further improved (Tab. IV) by considering user corrections as sparse training examples to update the model during test [19].

B. Network Architecture.

In this work, we adopt the Minkowski Engine [6], an implementation of sparse convolutional networks [11]. Architectures based on the Minkowski Engine have shown impressive results in 3D semantic segmentation. The sparse convolutions allow the efficient use of 3D convolutional layers and deep architectures already staple in 2D vision.

All network weights except for the first and output layer are initialized with those provided by [6]. They are obtained by pretraining the network on the task of semantic segmentation on the ScanNetV2 dataset (Train) [7]. The input now has two additional channels for the user clicks, and the last output layer is replaced with a two-class softmax layer, which is used to produce binary segmentations (foreground/background). Additionally, the network is fine-tuned for the task of interactive object segmentation using the ScanNetV2 train dataset adapted for foreground/background segmentation: a ground truth segmentation mask is created for each object instance where a single object instance is labeled as foreground and every other point in the 3D scene as background. Groundtruth segmentation masks are created for all 3D instances in the ScanNetV2 train dataset and are used to fine-tune the model on our task, that is 3D interactive object segmentation.

Implementation and Training Details. We set the input

feature dimension C = 6 (position and color). For the LiDAR data of SemanticKITTI [2], only 3D positions are provided, such that C = 3. The online adaptation experiments (Tab. IV) follow the protocol of [19] for instance adaptation (IA) with learning rate of $5 \cdot 10^{-5}$ and $\lambda = 0.1$.

IV. EXPERIMENTS

Datasets. The proposed interactive segmentation model in all our experiments is trained on a single dataset [7] and tested several others: S3DIS [1], SemanticKITTI [2], ScanNetval [7]. This follows established evaluation protocol from 2D interactive object segmentation.

In our work, we utilize the ScanNetV2 [7] train set since it is currently the largest of its kind containing approximately 1200 scenes. However, ScanNetV2 [7] is still much smaller in scale compared to the image datasets used in 2D interactive object segmentation. To compare, in 2D interactive object segmentation, a typical deep learning model [25] is using the weights from pre-training on ImageNet [8], COCO [23], and PASCAL VOC '12 [10] for the task of semantic segmentation and is fine-tuned on PASCAL VOC '12 [10] for the task of interactive object segmentation. In this work, we evaluate three well-established public 3D datasets ScanNetV2 [7] validation, S3DIS [1] and SemanticKITTI [2].

Evaluation metrics. We perform the evaluation using the two standard metrics from the 2D domain [3, 16, 20, 22, 25, 35]: (1) **Number of Clicks (NOC)@q%**, the average number of clicks needed to reach q% Intersection over Union (IoU) between predicted and ground-truth masks on every object instance (thresholded at 20 clicks). The <u>lower</u> the NOC value the <u>better</u>, and (2) **IoU@k**, the average IoU for k number of clicks per object instance. We additionally use the standard Average Precision (**AP**) metric for 3D instance segmentation.

A. Within-Domain Evaluation

At first, we evaluate our approach by training and testing on the same domain. In this scenario, we assume that we would like to extend an existing dataset by labeling more scenes. For this experiment, we use the ScanNetV2 dataset [7] with the official split of train-validation for our experiments. We fine-tune our model on interactive object segmentation on the training subset of ScanNetV2 and evaluate on the validation set. We unfortunately cannot evaluate on the test split due to hidden labels and no existing benchmark for interactive object segmentation. ScanNetV2 [7] contains the segmentation masks for 40 classes. However, the ScanNetV2 [7] benchmark evaluates only on a 20-class subset (18 for instance segmentation). We use the 20-class subset of ScanNetV2 training set to train our model, thus creating two subsets of the validation set: (1) "seen", with object instances from the 20 object categories which have been seen during training and (2) "unseen", with object instances of classes that have not been seen during training. The "seen" classes are also known as ScanNetV2 Benchmark Classes as these are the classes evaluated on the benchmark.

Results. Results are summarized in Tables I and II. No previous work uses interactive object segmentation directly

 TABLE I

 Instance segmentation scores on ScanNetV2 validation.

	Method	AP	$AP_{50\%}$	$AP_{25\%}$
	SSTNet [21]	49.4	64.3	74.0
ses	HAIS [4]	43.5	64.1	75.6
ass	HAIS [4] class agnostic	55.1	75.3	84.9
ວ	Ours (1 click per object)	20.9	38.0	67.2
Ę	Ours (2 clicks per object)	37.9	60.8	84.3
na	Ours (3 clicks per object)	48.1	73.4	90.7
ĥ	Ours (5 clicks per object)	61.8	86.0	95.7
enc	Ours (10 clicks per object)	75.5	94.8	99.0
B	Ours (20 clicks per object)	80.3	96.6	99.3
	HAIS [4] class agnostic	13.8	26.0	40.9
ses	Ours (1 click per object)	13.5	27.8	54.4
ass	Ours (2 clicks per object)	24.4	45.7	71.2
J	Ours (3 clicks per object)	33.1	58.4	81.8
en	Ours (5 clicks per object)	45.7	75.3	91.1
se	Ours (10 clicks per object)	60.3	90.2	97.9
Ūn	Ours (20 clicks per object)	60.6	88.3	97.4

TABLE II WITHIN-DOMAIN EVALUATION. INTERACTIVE 3D OBJECT SEGMENTATION ON SCANNETV2 VALIDATION

2D Domain	Pascal	Grabcut	Berkeley	MS COCO	MS COCO
(↓) NOC @ k% IoU	85%	90 %	90 %	85%	85%
DIOS w/ GC [35]	6.9	6.0	8.7	8.3	7.8
ScanNetV2 val- NOC @ k % IoU	See 80% 859	n % 90%	Unse 80% 85%	en 6 90% 80%	All % 85% 90%
Ours	8.3 10.	6 13.6	11.7 14.1	1 16.5 11.6	6 14.1 14.7

on the 3D domain. Therefore, we evaluate our method on a 3D instance segmentation setup and compare our model with state-of-the-art fully supervised methods for 3D instance segmentation (Tab. I, *top*). For a fair comparison, we also report scores for HAIS [4] in a class-agnostic setup, where the semantic class is ignored. When compared to fully-supervised instance segmentation methods (HAIS, SSTNet), our method offers segmentation masks of higher quality with only a few clicks per object. With 5 clicks per object, ours achieves 61.5 AP compared to 55.1 AP of HAIS, and with 10 clicks the score further increases to 75.5 AP.

The difference is even more pronounced when we evaluate on the additional "*unseen*" classes of ScanNetV2 [7] that are not part of the training set (Tab. I, *bottom*). ScanNetV2 [7] is a dataset with 40 labeled classes but usually only the 20class subset is evaluated. We train our model only on the 20-class benchmark subset in contrast to HAIS [4] which trains on all classes. With just 2 additional clicks our model doubles the AP. This makes our method ideal for labeling new datasets with previously unseen semantic classes. In summary, our method is orthogonal to instance segmentation methods. We show that it can be used to improve the quality of segmentation masks at the cost of minimal user effort.

Our evaluation on interactive 3D object segmentation is presented in Tab. II. To provide a frame of reference (since there are no previous methods on interactive 3D object segmentation) we report the scores (Tab. II) from the first deep learning paper [35] on 2D interactive object segmentation. Our

 TABLE III

 COMPARISON WITH CONTRASTIVE SCENE CONTEXT (CSC) [14].

Method	Points per Scene	Points per Objects	Objects per Scene	AP _{50%}
CSC [14]	200 (train)	-	-	48.9
Ours	60 (test)	2	≈ 30	60.8
Ours	90 (test)	3	≈ 30	73.4

scores are comparable with its 2D counterpart even though the 2D method uses networks pre-trained on much bigger datasets and some of the evaluation datasets contain only one object instance per image (*e.g.*, GrabCut [27]) compared to the 30 objects of an average indoor 3D scene. The 2D method also employs additional GraphCut [35] refining.

We additionally report our scores compared to the method of Shen *et al.* [28] that segments objects using 2D scribbles in Tab. VI. [28] reports scores on selected classes of the Pix3D [30] dataset. Annotators labeled 95 randomly-selected objects to be used for fine-tuning a 3D reconstruction model pre-trained on synthetic data. Additionally, it is not easy to define the size and form of the scribbles used. The labeled objects were un-occluded and un-truncated so their statistics are not directly comparable with our evaluation on the densely populated datasets of ScanNetV2 (val) and S3DIS. However, the comparison might show a trend for common object classes.

Comparison with Weakly-Supervised. We compare our method with the weakly supervised method of [14] in Tab. III. Weakly-supervised methods aim to achieve fully supervised accuracy with less labeled data, while our method on the other hand aims to supersede the performance of existing fully supervised methods (as shown in Fig. 1, labeled datasets require segmentation mask of the highest quality) using limited additional human input. Unlike weakly-supervised methods, our approach does not require re-training the model (even with reduced number of points) when we want the segmentation masks of classes not existing in the training set or when used in different datasets. Our method performs significantly better than [14] with less input from the user. Importantly, unlike weakly-supervised methods, our method does not require re-training when new user inputs appear.

B. Out-of-Domain Evaluation.

In this section, we evaluate the performance of our method on datasets different from our training dataset. We would like to see the performance of our model both in relatively small distribution shifts between training and test using two indoor datasets. (ScanNet \rightarrow S3DIS). Additionally, we evaluate on large distribution shifts (ScanNet \rightarrow SemanticKITTI) by using our model trained on an indoor dataset and testing on an outdoor dataset containing different objects and recorded with

TABLE IV

OUT-OF-DOMAIN EVALUATION. INTERACTIVE 3D OBJECT SEGMENTATION SCORES ON S3DIS AND SEMANTICKITTI.

 TABLE V

 Per-class interactive 3D object segmentation (S3DIS A5 [1])

NOC@90														
Method	mean	Ceiling	Floor	Wall	Clutter	Column	Window	Door	Table	Chair	Bookcase	Sofa	Board	Beam
Ours Ours w/ [19]	$\begin{array}{c} 11.8\\ 11.5 \end{array}$	$4.6 \\ 4.5$	$2.1 \\ 2.2$	$\begin{array}{c} 11.4 \\ 11.0 \end{array}$	$\begin{array}{c} 15.0 \\ 14.8 \end{array}$	$\begin{array}{c} 11.9\\ 11.4 \end{array}$	$\begin{array}{c} 11.5 \\ 10.5 \end{array}$	$\begin{array}{c} 13.3\\ 12.9 \end{array}$	$\begin{array}{c} 11.8\\ 11.3 \end{array}$	$5.3 \\ 5.2$	$\begin{array}{c} 10.6 \\ 10.2 \end{array}$	$5.7 \\ 5.0$	$\begin{array}{c} 16.0 \\ 14.7 \end{array}$	$\begin{array}{c} 17.7\\ 14.0 \end{array}$

a different depth sensor like SemanticKITTI [2]). First, we use our model trained on ScanNetV2 training set *without any fine-tuning* on these datasets. We additionally report separate scores when online adaptation is used.

Results. Results are summarized in Tables IV and V. As shown in Tab. IV (row 1) even though our interactive segmentation network was trained only on the ScanNet dataset it can still generalize easily to other indoor datasets like S3DIS. It even requires fewer clicks to achieve segmentation masks of the same quality. The more our test dataset distribution differs from the training, the more clicks are required to achieve a good segmentation mask as shown in the results on semanticKITTI (Tab. IV). The SemanticKITTI dataset requires more clicks for two reasons: (1) its sparsity restricts the needed context from the click masks and (2) its characteristics (sparsity, no color, different sensor capturing the data) are very different from the training dataset. To that end, we incorporate an online adaptation method that updates the model weights during test time using sparse user annotations [19]. Tab. IV (row 2-right) shows that online adaptation helps significantly to adapt to new dataset distributions, requiring 4 clicks less. The adaptation does not help significantly on the S3DIS dataset Tab. IV (row 2left) since ScanNet and S3DIS do not differ as much as ScanNet and SemanticKITTI. However, if we look at the evaluation metrics per class (the method is class agnostic, we only aggregated the results per individual class), online adaptation helps significantly to segment objects belonging to the classes of board and beam (Tab. V), i.e. two classes that do not exist in this form in ScanNet.

C. Qualitative Results.

We show visualizations on full scenes in Fig. 3 and of single objects on all three datasets in Fig. 4. Our method predicts accurate segmentations in heavily cluttered environments with very limited user clicks.

D. User Study

To evaluate our method with real human annotators, we also conducted a small-scale experiment replacing the simulated annotator with three real human annotators in the loop. To that end, we implemented an interactive annotation tool that runs our model at the backend. The annotators engaged in the user study had no prior experience with annotating. We TABLE VI

3 D	CLICKS	COMPARED	TO 2D	SCRIBBLES.

(trained on ScanNet) S3DIS Area 5 [1]		Sema	nticKIT	FI [2]											
NOC @ k % IoU	80 %	85%	90%	80 %	85%	90 %		Annotations	Dataset	Bed	Bookc.	Desk	Sofa	Table	Wardr.
-							SIM [28] 2D scribbles	Pix3D	15	18	10	10	19	9
Ours	6.8	8.9	11.8	19.4	19.5	19.6	Ours	3D clicks	S3DIS	-	10.6	-	5.7	11.8	-
Ours (online adapt.)	6.5	8.6	11.5	15.1	15.4	15.9	Ours	3D clicks	ScanNet val.	16.3	15.7	18.5	12.5	11.8	-



Fig. 4. **Qualitative results** with clicks and corresponding segmentation masks. Our approach accurately segments objects from different datasets with as little as **one user click** (*top row, right*). ScanNet* (*second row*) shows classes that have not been used in training: *television* (*left*) and *towel* (*right*).

explained to them how the tool works and allowed them to label a few examples to familiarize themselves with the tool before recording their behavior. The annotators see the full scene as does the simulated annotator without cropping or using bounding boxes. If they deemed the segmentation satisfactory, they could submit the result as final at any time.

First, we asked the annotators to label 10 random objects from the challenging class *picture* using our interactive segmentation tool and we compared the real users results with the corresponding results from the simulated annotator (*Task 1*). Then we asked the annotators to randomly segment 20 objects from ScanNetV2 val. set and compared them with the same objects annotated by the simulated annotator (*Task 2*). The results are presented in Tab. VII. Task 1 is especially challenging as it can be seen from the example figure of our user interface (right). The geometry of the points belonging to the class is very similar to the one of the wall nearby and only color can help distinguish the separate objects. Despite that, human annotators performed slightly better than the simulated

TABLE VII

USER STUDY WITH HUMAN ANNOTATORS. HUMAN ANNOTATORS LABELED POINTS OF THE CLASS "PICTURE" AND OF RANDOM OBJECTS SHOWING THAT OUR SIMULATION STRATEGY IS REPRESENTATIVE OF REAL HUMAN INPUT AND PRODUCES SIMILAR RESULTS.

	ScanNet	V2-picture	ScanNet	V2-random	-
	IoU@5	IoU@10	IoU@5	IoU@10	10 al
Human	49.6	57.0	78.1	82.0	
Simulator	32.0	54.5	79.8	84.8	

clicks especially in the low clicks regime. Labeling random objects also performs similarly to the simulated clicks.

We note that although our model is trained with simulated clicks, it shows robustness to human input even though they follow different labeling strategies and that our simulation of interactive object segmentation (Sec. III) is representative of real human input and both provide similar output.

E. Inference Speed.

After user interaction, inference with the updated segmentation mask takes **0.06 sec**, making the method suitable for real-time point cloud annotation. Training is performed only once for 20 epochs for ≈ 14 hours. (All experiments on a single NVIDIA Titan RTX)

V. CONCLUSION

In this work, we introduced, for the first time, a method for interactive object segmentation on 3D points clouds. Our approach is a simple yet powerful framework that achieves highly accurate 3D segmentation masks with little human annotation effort. Our method generalizes well to different datasets not used in training and enables the annotations of novel large-scale 3D training sets including previously unseen semantic classes. In particular, we also presented practical policies necessary for simulating user clicks during training and evaluation, verified by a user study.

Acknowledgements. This work was supported by the SNF grant 200021 204840 and an ETH AI Center postdoctoral fellowship to Theodora Kontogianni.

REFERENCES

- 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 *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
- [2] J. Behley, M. Garbade, A. Milioto, J. Quenzel, S. Behnke, C. Stachniss, and J. Gall. SemanticKITTI: A Dataset for Semantic Scene Understanding of LiDAR Sequences. In *Proceedings of the International Conference on Computer Vision*, 2019.
- [3] Rodrigo Benenson, Stefan Popov, and Vittorio Ferrari. Large-scale interactive object segmentation with human annotators. In *Proceedings* of the IEEE Conference on Computer Vision and Pattern Recognition, 2019.
- [4] Shaoyu Chen, Jiemin Fang, Qian Zhang, Wenyu Liu, and Xinggang Wang. Hierarchical aggregation for 3d instance segmentation. In Proceedings of the International Conference on Computer Vision, 2021.
- [5] Julian Chibane, Francis Engelmann, Tuan Anh Tran, and Gerard Pons-Moll. Box2Mask: Weakly Supervised 3D Semantic Instance Segmentation Using Bounding Boxes. In *Proceedings of the European Conference on Computer Vision*, 2022.
- [6] Christopher Choy, JunYoung Gwak, and Silvio Savarese. 4D Spatio-Temporal ConvNets: Minkowski Convolutional Neural Networks. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2019.
- [7] Angela Dai, Angel X. Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. ScanNet: Richly-annotated 3D Reconstructions of Indoor Scenes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017.
- [8] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-fei. ImageNet: A large-scale hierarchical image database. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2009.
- [9] Francis Engelmann, Martin Bokeloh, Alireza Fathi, Bastian Leibe, and Matthias Nießner. 3D-MPA: Multi Proposal Aggregation for 3D Semantic Instance Segmentation. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2020.
- [10] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The PASCAL Visual Object Classes Challenge 2012 (VOC2012) Results. http://www.pascalnetwork.org/challenges/VOC/voc2012/workshop/index.html, 2012.
- [11] Benjamin Graham, Martin Engelcke, and Laurens van der Maaten. 3D Semantic Segmentation with Submanifold Sparse Convolutional Networks. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2018.
- [12] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In Proceedings of the International Conference on Computer Vision, 2017.
- [13] 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, 2019.
- [14] Ji Hou, Benjamin Graham, Matthias Nießner, and Saining Xie. Exploring data-efficient 3d scene understanding with contrastive scene contexts. In *Proceedings of the IEEE Conference on Computer Vision* and Pattern Recognition, 2021.
- [15] Ji Hou, Benjamin Graham, Matthias Nießner, and Saining Xie. Exploring data-efficient 3d scene understanding with contrastive scene contexts. In *Proceedings of the IEEE Conference on Computer Vision* and Pattern Recognition, 2021.
- [16] Won-Dong Jang and Chang-Su Kim. Interactive image segmentation via backpropagating refinement scheme. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019.
- [17] Li Jiang, Shaoshuai Shi, Zhuotao Tian, Xin Lai, Shu Liu, Chi-Wing Fu, and Jiaya Jia. Guided point contrastive learning for semi-supervised point cloud semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2021.
- [18] Li Jiang, Hengshuang Zhao, Shaoshuai Shi, Shu Liu, Chi-Wing Fu, and Jiaya Jia. PointGroup: Dual-set Point Grouping for 3D Instance Segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2020.
- [19] Theodora Kontogianni, Michael Gygli, Jasper Uijlings, and Vittorio Ferrari. Continuous Adaptation for Interactive Object Segmentation by Learning from Corrections. In *Proceedings of the European Conference* on Computer Vision, 2020.

- [20] Z. Li, Q. Chen, and V. Koltun. Interactive image segmentation with latent diversity. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
- [21] Zhihao Liang, Zhihao Li, Songcen Xu, Mingkui Tan, and Kui Jia. Instance segmentation in 3d scenes using semantic superpoint tree networks. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2021.
- [22] J.H. Liew, Y. Wei, W. Xiong, S-H. Ong, and J. Feng. Regional interactive image segmentation networks. In *Proceedings of the International Conference on Computer Vision*, 2017.
- [23] T-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C.L. Zitnick. Microsoft COCO: Common objects in context. In *Proceedings of the European Conference on Computer Vision*, 2014.
- [24] Zhengzhe Liu, Xiaojuan Qi, and Chi-Wing Fu. One thing one click: A self-training approach for weakly supervised 3d semantic segmentation. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2021.
- [25] S. Mahadevan, P. Voigtlaender, and B. Leibe. Iteratively trained interactive segmentation. In *Proceedings of the British Machine Vision Conference*, 2018.
- [26] Charles R. Qi, Li Yi, Hao Su, and Leonidas J Guibas. PointNet++: Deep Hierarchical Feature Learning on Point Sets in a Metric Space. In Advances in Neural Information Processing Systems, 2017.
- [27] C. Rother, V. Kolmogorov, and A. Blake. GrabCut Interactive Foreground Extraction using Iterated Graph Cut. Proceedings of the ACM SIGGRAPH Conference on Computer Graphics, 23, 2004.
- [28] Tianchang Shen, Jun Gao, Amlan Kar, and Sanja Fidler. Interactive Annotation of 3D Object Geometry using 2D Scribbles. In Proceedings of the European Conference on Computer Vision, 2020.
- [29] Konstantin Sofiiuk, Ilya A. Petrov, and Anton Konushin. Reviving iterative training with mask guidance for interactive segmentation. *IEEE International Conference on Image Processing*, 2021.
- [30] Xingyuan Sun, Jiajun Wu, Xiuming Zhang, Zhoutong Zhang, Chengkai Zhang, Tianfan Xue, Joshua B Tenenbaum, and William T Freeman. Pix3D: Dataset and Methods for Single-Image 3D Shape Modeling. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2018.
- [31] Hugues Thomas, Charles R. Qi, Jean-Emmanuel Deschaud, Beatriz Marcotegui, François Goulette, and Leonidas J. Guibas. KPConv: Flexible and Deformable Convolution for Point Clouds. In *Proceedings* of the International Conference on Computer Vision, 2019.
- [32] Julien Valentin, Vibhav Vineet, Ming-Ming Cheng, David Kim, Jamie Shotton, Pushmeet Kohli, Matthias Nießner, Antonio Criminisi, Shahram Izadi, and Philip Torr. Semanticpaint: Interactive 3d labeling and learning at your fingertips. ACM Transactions on Graphics (TOG), 2015.
- [33] Haiyan Wang, Xuejian Rong, Liang Yang, Jinglun Feng, Jizhong Xiao, and Yingli Tian. Weakly supervised semantic segmentation in 3d graphstructured point clouds of wild scenes. arXiv preprint arXiv:2004.12498, 2020.
- [34] Xinlong Wang, Shu Liu, Xiaoyong Shen, Chunhua Shen, and Jiaya Jia. Associatively Segmenting Instances and Semantics in Point Clouds. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2019.
- [35] N. Xu, B. Price, S. Cohen, J. Yang, and T.S. Huang. Deep interactive object selection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
- [36] Xun Xu and Gim Hee Lee. Weakly supervised semantic point cloud segmentation: Towards 10x fewer labels. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2020.
- [37] Bo Yang, Jianan Wang, Ronald Clark, Qingyong Hu, Sen Wang, Andrew Markham, and Niki Trigoni. Learning Object Bounding Boxes for 3D Instance Segmentation on Point Clouds. In Advances in Neural Information Processing Systems, 2019.
- [38] Yachao Zhang, Yanyun Qu, Yuan Xie, Zonghao Li, Shanshan Zheng, and Cuihua Li. Perturbed self-distillation: Weakly supervised largescale point cloud semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2021.
- [39] Shuaifeng Zhi, Edgar Sucar, Andre Mouton, Iain Haughton, Tristan Laidlow, and Andrew J Davison. iLabel: Revealing Objects in Neural Fields. *IEEE Robotics and Automation Letters*, 2023.