OpenTrench3D: A Photogrammetric 3D Point Cloud Dataset for Semantic Segmentation of Underground Utilities

Supplementary Material

1. Introdution

This document presents the supplementary materials omitted from the main paper due to the space limitation.

A. Data Capture Process

The OpenTrench3D dataset is gathered using close-range photogrammetry captured using video recordings from everyday smartphones. The following is a description of the overall data capture and processing workflow used by the two utility owners that provided point clouds for OpenTrench3D. We refer to figure 1 for illustrations. This workflow highlights that utility owners on site only requires a marker and a smartphone to fulfill their role in the data capture process. The procedure is divided into three straightforward steps: (1) apply markings around the open trench, used as Ground Control Points (GCP), possibly using a spray marker; (2) carefully video record the trench from various angles, ensuring the camera is aimed down towards the utilities visible in the trench; and (3) upload the captured video through the companion application. Subsequent, the video data is then send to a server for processing into a 3D point cloud. Following the initial step, a surveying responsible is tasked with measuring the GCP markings using survey-grade instruments, such as GNSS-RTK receivers followed by uploading this data to the same system for later manual geo-referencing of the point cloud.

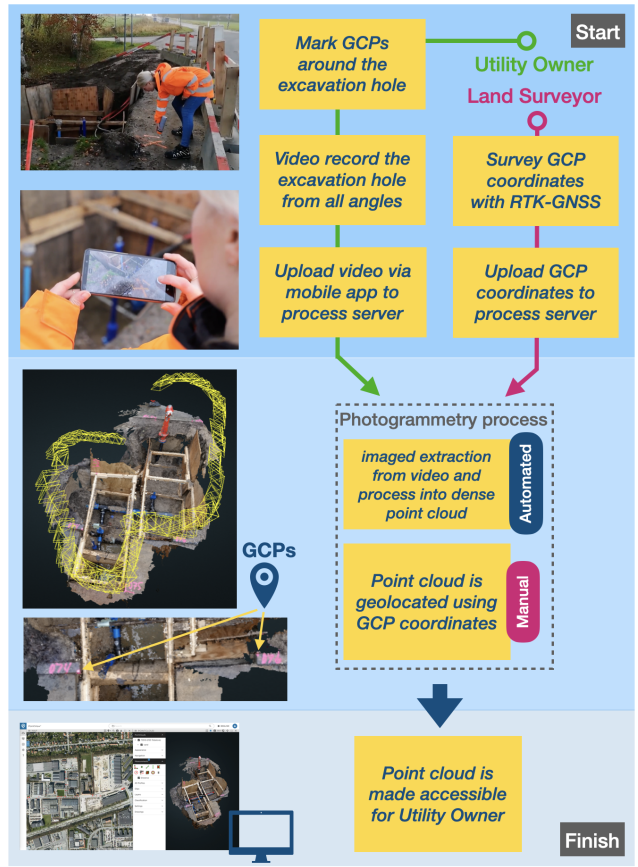

Figure 1. Workflow of the data capture process to generate 3D point cloud using photogrammetry. GCP stands for Ground Control Point.

B. Fine-tuning Evaluation on Heating Areas

In table 1 we present the supplementary results of the fine-tuning evaluation on heating areas of PointVector and PointMetaBase, similar to table 4 in the main article for PointNeXt. In the fine-tuning evaluation on heating areas, we first pre-train model weights on samples from Water Area 1-4 while using Water Area 5 as validation set for the pre-training. Secondly, we fine-tune the model weights on 1, 5, 10, 20 and all (29) samples from Heating Area 1. We both conduct fine-tuning experiments in which only weights of the segmentation head of each model are fine-tuned as well as experiments in which the weights of both the segmentation head and the decoder are fine-tuned, simultaneously. Finally, the fine-tuned models are evaluated on point clouds from Heating Area 2.

C. Qualitative Examples from 5-fold cross-validation on Water Areas

In figure 2, we provide qualitative results from running inference on various samples from Water Area 5 using a PointVector model trained on Water Area 1-4 with the *Inactive Utility* class included (colored in blue) for evaluation. The qualitative examples highlight the challenges state-of-the-art semantic segmentation methods encounter when trying to distinct the *Inactive Utility* against the *Main Utility* and *Other Utility* classes. Although, sometimes succesful, often times the methods neglect the *Inactive Utility* class, possible due to it occuring less frequently in the dataset compared to the other classes, as seen from table 2 in the main paper.

Table 1. Supplementary table for the fine-tuning experiments seen in figure 4 of the main paper. For comparison, we display the performance of pre-trained models (red square) as well as models without prior pre-training, but solely trained on 1, 5, 10, 20 and all samples from Heating Area 1 (Baseline).

		PointVector					PointMetaBase				
	Samples	mAcc	mIoU	Main	Other	Trench	mAcc	mIoU	Main	Other	Trench
Baseline	1	70.9	66.8	78.1	25.5	96.7	70.3	64.9	73.2	25.6	96.0
	5	73.8	70.0	77.4	36.0	96.6	74.6	69.4	75.4	36.4	96.5
	10	76.2	71.7	79.8	38.4	97.0	77.1	73.0	78.8	**43.2**	97.0
	20	79.5	71.8	79.8	38.6	97.1	**80.4**	71.2	79.9	36.3	97.5
	29	81.2	72.6	77.5	43.1	97.1	79.5	72.5	80.3	39.8	97.2
Pre-trained		79.3	71.8	80.3	37.6	97.6	76.7	68.5	75.7	32.4	97.4
Fine-tuned (Head)	1	76.4	73.1	83.8	37.9	97.7	72.6	66.6	76.0	27.5	96.3
	5	77.4	74.0	82.8	41.8	97.5	74.3	70.0	75.4	37.9	96.6
	10	82.5	75.9	85.8	44.0	97.9	78.2	73.3	81.4	41.0	97.4
	20	81.8	73.8	83.5	39.9	98.0	79.7	68.9	78.4	30.7	**97.7**
	29	**84.5**	73.2	84.6	37.3	97.8	79.5	66.9	72.0	31.4	97.3
Fine-tuned (Decoder+Head)	1	76.2	73.0	83.8	37.7	97.6	75.9	68.6	72.7	37.4	95.8
	5	77.5	73.9	83.3	40.7	97.6	78.8	**73.9**	**81.8**	42.6	97.4
	10	82.2	**76.6**	**86.7**	**45.2**	**98.1**	79.0	72.8	80.9	40.0	97.4
	20	81.8	74.0	84.3	39.5	**98.1**	80.0	69.4	78.9	31.4	**97.7**
	29	80.1	72.6	81.3	38.9	97.8	78.9	69.1	75.4	34.7	97.3

D. Qualitative Examples from Fine-tuning evaluation on Heating Areas

We provide additional qualitative examples from running inference on various samples from Heating Area 2 using various trained versions of the PointNeXt model in figure 3. These are supplements to the qualitative examples in figure 5 of the main paper. We showcase inference results from a PointNeXt models trained from scratch on solely 1 and 10 samples from Heat Area 1 (Baseline), a PointNeXt model trained on samples from Water Area 1-4 (pre-trained) and finally PointNeXt models pre-trained on Water Area 1-4 and further fine-tuned on 1 and 10 samples from Heating Area 1, were only weights of the segmentation head are tuned.

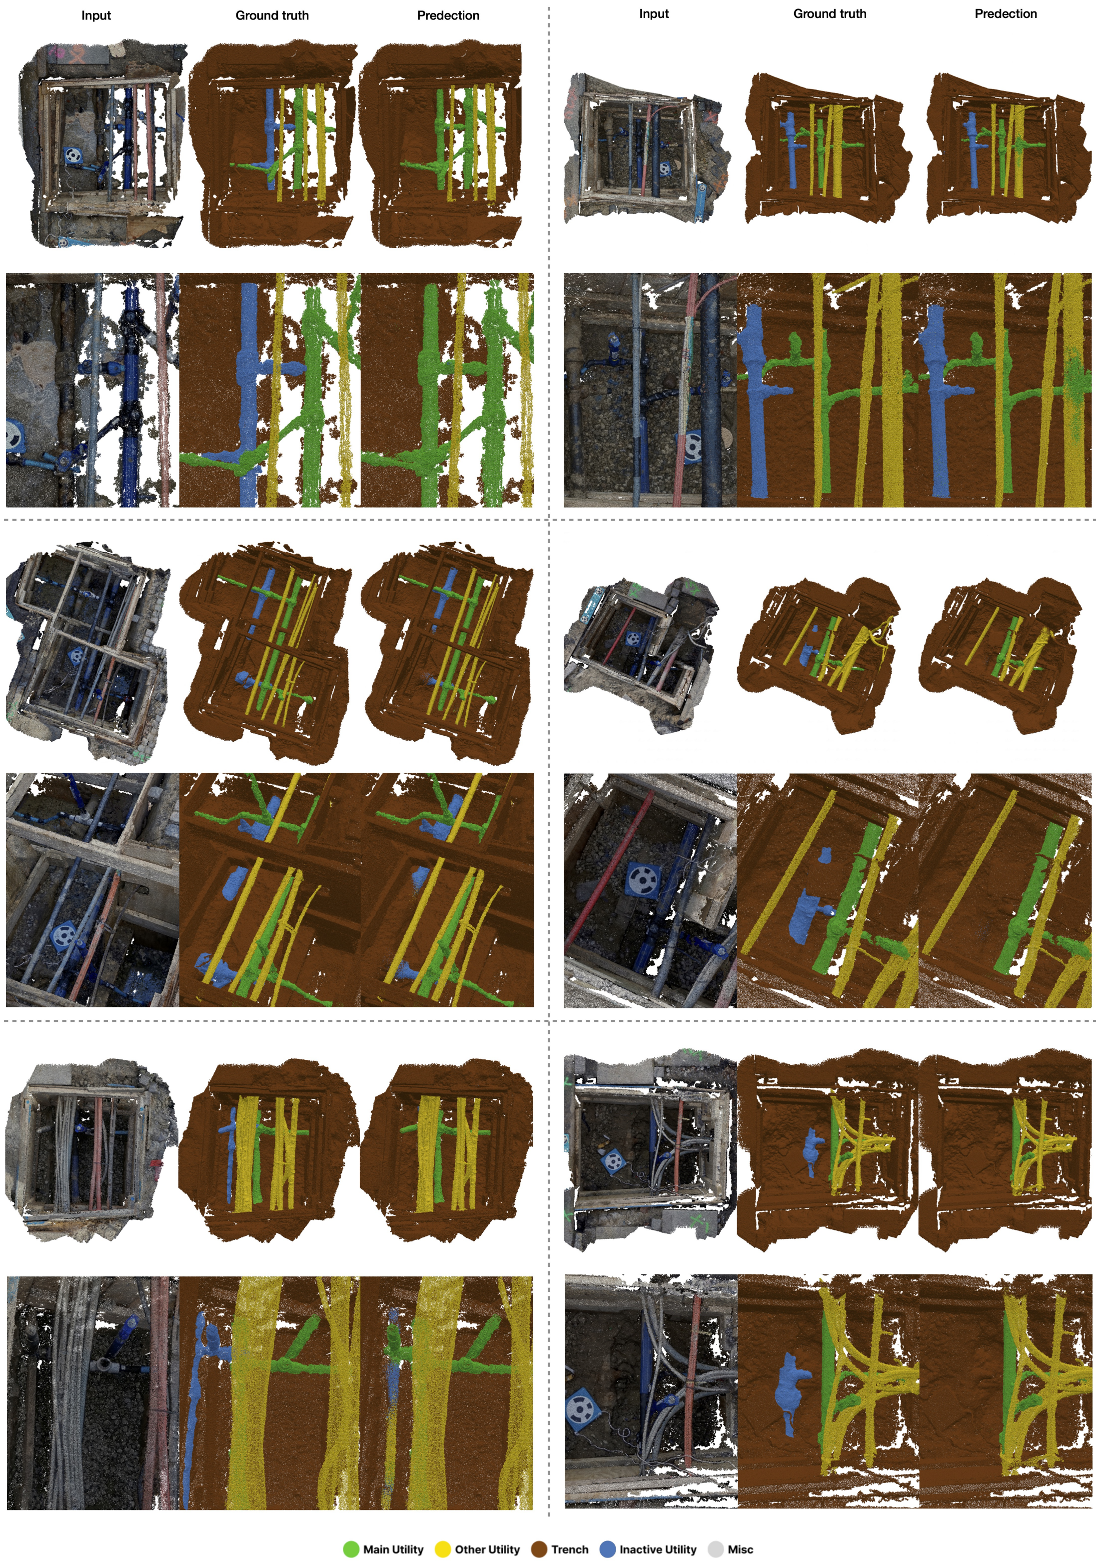

Figure 2. Qualitative Examples from running inference on samples from Water Area 2, with a trained PointVector model on Water Area 1, 3, 4 and 5. In this model, the *Inactive Utility* class was included to test against the *Main Utility* and *Other Utility* classes.

Figure 3. Qualitative Examples from running inference on samples from Heating Area 2 with 5 trained PointNeXt model versions: 2 models which are trained on 1 and 10 samples from Heating Area 1 (called Baseline), 1 model pre-trained on samples from Water Area 1-4 and 1 model pre-trained on samples from Water Area 1-4 and fine-tuned with 1 and 10 samples from Heating Area 1.