

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

Lasse H. Hansen^{1,2,*} Simon B. Jensen^{1,3,*} Mark P. Philipsen¹
 Andreas Møgelmoose¹ Lars Bodum¹ Thomas B. Moeslund¹
¹Aalborg University ²IT34 ³Ambolt AI

lhe@it34.com, simon@ambolt.io {mpph, ammo, tbm}@create.aau.dk, lbo@plan.aau.dk

Abstract

Identifying and classifying underground utilities is an important task for efficient and effective urban planning and infrastructure maintenance. We present OpenTrench3D, a novel and comprehensive 3D Semantic Segmentation point cloud dataset, designed to advance research and development in underground utility surveying and mapping. OpenTrench3D covers a completely novel domain for public 3D point cloud datasets and is unique in its focus, scope, and cost-effective capturing method. The dataset consists of 310 point clouds collected across 7 distinct areas. These include 5 water utility areas and 2 district heating utility areas. The inclusion of different geographical areas and main utilities (water and district heating utilities) makes OpenTrench3D particularly valuable for inter-domain transfer learning experiments. We provide benchmark results for the dataset using three state-of-the-art semantic segmentation models, PointNeXt, PointVector and PointMetaBase. Benchmarks are conducted by training on data from water areas, fine-tuning on district heating area 1 and evaluating on district heating area 2. The dataset is publicly available¹. With OpenTrench3D, we seek to foster innovation and progress in the field of 3D semantic segmentation in applications related to detection and documentation of underground utilities as well as in transfer learning methods in general.

1. Introduction

Motivated by an ongoing and escalating demand for infrastructure development both above and beneath the ground, urban planners, engineers, and contractors depend on data regarding the placement of hidden underground utilities [32, 35]. Yet, a large portion of utility data is often inaccurate in location, outdated, and incomplete [13, 33]. This lack of accurate and current information becomes particularly

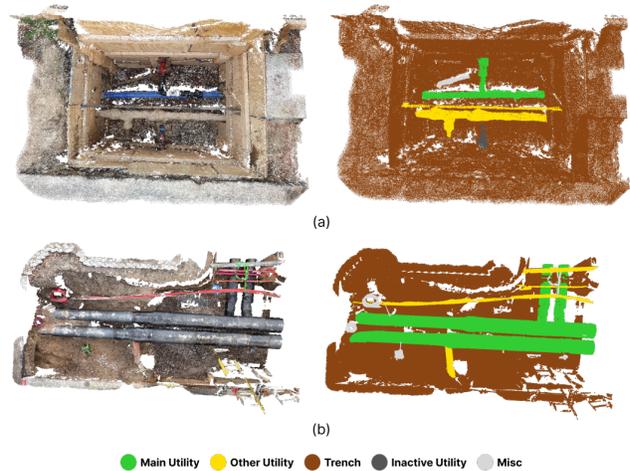


Figure 1. Point clouds from OpenTrench3D: (a) In the water project areas the *Main Utility* class is made up of water utility pipes. (b) In the district heating project areas the *Main Utility* class is made up of district heating utility pipes. Find description of dataset and classes in section 3.

critical when considering the financial impact of excavation damage to underground infrastructure. Reports highlight that such damages cost more than GBP 200 million in the UK [4] and USD 30 billion in the US [1] alone in 2019. This underscores the necessity for utility owners to recognize the importance of consistently maintaining up-to-date and accurate utility map data.

Innovation in utility surveying and mapping therefore holds significant importance. The most common method is known as open trench surveying, which occurs after the installation of new utility lines or the replacement of existing ones, involving excavation on site. However, during these surveys, only the utilities of interest to the owner are documented, while other exposed utilities are frequently overlooked, as utility owners lack short-term incentives to re-survey, validate, and potentially report updates on the assets of other utility owners [13].

*Equal contribution

¹<https://github.com/SimonBuusJensen/OpenTrench3D>

Following the recent trend in urban scene modeling, using affordable and accessible drone photogrammetry for 3D reconstruction of urban environments [5, 6, 16, 21], the adoption of close-range photogrammetric data capture via off-the-shelf smartphones for utility surveying has gained popularity [12, 38]. This is driven by the emergence of industry solutions like Pix4D. Unlike conventional open trench methods, such as direct single point measurement with GNSS receivers, this approach captures every piece of underground infrastructure visible in the 3D reconstruction. However, both the utilities of interest and other utilities must be digitized as polylines for subsequent integration into utility network systems such as GIS. The annotation process is manual and laborious. Developing methods for the classification and segmentation of utilities could not only greatly increase efficiency but also open new possibilities for streamlining the process of validating and updating the locations of existing utilities.

We introduce OpenTrench3D, to our knowledge, the first publicly available 3D point cloud dataset for the semantic segmentation of underground utilities captured in open trenches. The dataset was created from smartphone video recordings and reconstructed using photogrammetry, making it an accessible and affordable data acquisition method to replicate. It comprises 310 fully annotated point clouds, totaling approximately 528 million points across 5 label classes, following a utility owner-centric classification scheme. The dataset includes point clouds from water and district heating projects. One example of each is shown in Figure 1.

Since OpenTrench3D mainly consists of point clouds from water projects, it offers an excellent opportunity to study transfer learning to the smaller set of district heating point clouds. We leverage water area point clouds to develop a pretrained model, which is then fine-tuned and tested on district heating areas. This method evaluates the potential of transfer learning for semantic segmentation in open trenches across diverse utility types. Hereby, we demonstrate to what extent the amount of newly annotated point clouds of the target main utility types affects performance. We summarize the contributions of this paper as follows:

1. We present the first publicly available 3D point cloud dataset of underground utilities, OpenTrench3D, featuring 310 fully annotated point clouds and a utility owner-centric classification scheme.
2. We benchmark state-of-the-art 3D semantic segmentation deep learning methods. The results highlight the potential and effectiveness of transfer learning across utility types and evaluate the generalizability of SOTA methods.
3. We demonstrate that underground utility point clouds, acquired through cost-effective photogrammetry techniques, can be segmented using deep learning methods. Thus, making this method more accessible to industry and academia.

2. Related Work

2.1. 3D point cloud datasets

Compared to existing 3D point cloud datasets, OpenTrench3D distinguishes itself as the first dataset specifically focused on open trench underground utilities. Despite its unique focus, we consider our dataset within the broader spectrum of 3D datasets for urban and outdoor environments. In this section, we offer a brief overview and comparison of existing point cloud datasets for semantic segmentation in this domain along with a comparative summary between OpenTrench3D and other well-known datasets in Table 1. Lastly, we briefly present current approaches for segmenting underground utilities.

Urban 3D datasets can generally be categorized based on their application area into two main groups: street-level and aerial-level.

Street-level: Derived from Mobile Laser Scanning (MLS) and Terrestrial Laser Scanning (TLS), these datasets capture detailed urban features such as buildings, vehicles, and vegetation. MLS has been predominantly used in foundational datasets like SemanticKITTI[3], Paris-Lille-3D[27], and Toronto-3D[31], whereas Semantic3D [10] uses TLS. Recent additions to this category, such as Urban SGPCM [29] and SP3D [23], maintain the core attributes of their predecessors while introducing variations in scale, class diversity and environmental context

Aerial-level: These datasets are obtained through Airborne Laser Scanning (ALS) and UAV-based photogrammetry, offering a bird's-eye view of urban structures on a large scale. Early examples such as LASDU[37] and DALES[34] use ALS, but UAV photogrammetry has since become the preferred method, as exemplified by Campus3D[21] and SensatUrban[16]. Similar to street-level datasets, newer aerial datasets primarily expand upon previous offerings by introducing variations in scale, class types, and specific urban environments (e.g., the HRHD-HK[20] dataset focuses on high-rise buildings).

While there is a wealth of 3D datasets for above-ground environments, there are no publicly available datasets for underground settings, despite the critical value and importance in urban planning [32]. In the closest alternatives, we find 2D datasets like Sewer-ML [15], which contains annotations of defects in sewer inspection CCTV videos [14]; however, this data, captured from inside underground pipes, is fundamentally different. OpenTrench3D addresses this gap by offering widely different 3D scenes of urban environments, specifically focusing on the infrastructure beneath street level. Moreover, what distinguishes OpenTrench3D from existing above-ground datasets is its method of data capture. Utilizing close-range photogrammetry with everyday smartphones sets it apart from the more commonly used laser scanning techniques and UAV photogrammetry.

Table 1. A comparative overview of key urban datasets for 3D point cloud semantic segmentation. MLS: Mobile Laser Scanning system, TLS: Terrestrial Laser Scanning system, ALS: Aerial Laser Scanning system, UAV: Unmanned Aerial Vehicle and Ptg: Photogrammetry.

Dataset	Category	Capture Method	Points	Area/Length	Classes	RGB
S3DIS[2]	Indoor Scene-level	RGB-D	273M	6,020 m ²	13	Yes
ScanNet[8]	Indoor Scene-level	RGB-D	242M	34,453 m ²	20	Yes
Semantic3D[10]	Street-level	TLS	4,009M	1 km	9	Yes
Paris-rue-Madame[28]	Street-level	MLS	20M	0.16 km	17	No
SemanticKITTI[3]	Street-level	MLS	4,549M	39.2 km	28	No
Toronto-3D[31]	Street-level	MLS	78M	1 km	9	Yes
DALES[34]	Aerial-level	ALS	505M	10 km ²	9	No
Hessigheim 3D[19]	Aerial-level	ALS	126M	0.19 km ²	11	No
SensatUrban[16]	Aerial-level	UAV Ptg	2,847M	7.64 km ²	31	Yes
Swiss3DCities[5]	Aerial-level	UAV Ptg	226M	2.7 km ²	5	Yes
STPLS3D[6]	Aerial-level	UAV Ptg	150M	1.27 km ²	6	Yes
OpenTrench3D (Ours)	Underground Scene-Level	Close-range Ptg	528M	3.814 m²	5	Yes

2.2. Segmentation of underground utilities

Semantic segmentation of 3D point clouds using point-based deep neural networks has gained traction following the pioneering work of PointNet and PointNet++ by Qi et al. [24, 25]. Despite these advancements, no prior work has tackled the challenge of segmenting 3D point clouds of underground utilities with deep learning methods, likely due to the lack of available data.

A closely related and significant contribution is the work of Stranner et al. [30]. They introduced a 3D fitting algorithm designed to align synthetically generated 3D cylinders with actual pipe structures, thereby updating the existing utility line maps to reflect the as-planned versus as-built states. Designed for on-site use, their system uses live 3D reconstructions from LiDAR-equipped mobile devices but requires user initiation via an Augmented Reality interface [11], leaving classification to the user. This approach, while innovative, underscores the need for more automated classification methods in 3D underground utility segmentation.

3. OpenTrench3D Dataset

3.1. Data Capture

The point clouds in this dataset were captured using an end-to-end data capture and processing solution called SmartSurvey². The solution uses close-range photogrammetry, more specifically Structure from Motion (SfM) with Multi-View Stereo vision (MVS), employing video recordings from commercially available smartphones. The utility companies carried out the video capture process, while a surveying company assisted in geo-referencing the generated point clouds by carefully surveyed spray markings. For an illustrative

²<https://it34.com/en/services/smartsurvey-app-en/>

example of the process, we refer to the supplementary material. The point clouds are estimated to have an absolute geographic accuracy within ± 5 cm [12].

3.2. Dataset Description

The 3D point clouds, documenting newly installed pipes, comprise 310 point clouds from 7 areas. Five areas relate to water supply projects, while two are dedicated to district heating projects, each associated with different utility companies. These areas are named Water Area 1-5 and Heating Area 1-2, all located in different urban regions of Denmark as showed on the maps in Figure 2. The point clouds have undergone the following post-processing: voxel downsampling to achieve a resolution of 4 mm, denoising using statistical outlier removal method, and manually removing noise artifacts in specific instances. Additionally, the background environment around the trench was removed to reduce the size of the point clouds. The point clouds were manually segmented, with labels assigned to each point, using the CloudCompare software [7]. The attributes for each point cloud PLY file are as follows: column 1-3: x, y and z-coordinates in meters. Column 4-6: r, g, b-color channels with values from 0-255. Column 7: class id.

3.3. Utility Owner-centric classification scheme

When deciding on a classification scheme for our dataset, we initially sought inspiration from existing standards and data models focusing on underground utility networks. Notable examples include the CityGML Utility Network ADE [18] and UUDM [36], which on a high-level provide classifications based on the utility type such as water, gas, electricity, sewage, telecommunications, etc. At first glance, adopting a similar classification scheme based on utility types appears logical. We discovered that accurately classifying utilities



Figure 2. Map overview of the five water areas (top left and middle) and the two district heating areas (right), with zoom-in maps of selected parts of some areas. Two point cloud samples are presented from the water and heating areas, respectively, and utilities from the *Main Utility*, *Other Utility*, and *Inactive Utility* classes are highlighted for comparison.

based on appearance is challenging, especially for older pipes, that are covered in dirt. This task is complex even for experts, who also gathers utility records from all relevant utility owners and carries out on-site inspections. Gathering the extensive information required for confident classification exceeds the scope of this paper.

Given these insights, we propose a utility owner-centric classification scheme that divides underground utilities into three separate classes, as shown in figure 2, with two additional classes for background and misc. objects. These classes are naturally distinguishable from a utility owner’s perspective and will be advantageous for subsequent utility mapping tasks. The five classes are defined as follows:

Trench: The surrounding open excavation pit where the utilities are laid. The class include everything not described in the other classes.

Main Utility: Newly installed utilities, which is the main utility of interest for surveying and mapping. In our dataset, this class is representing two distinct types of utilities: water and district heating. These utility types differ in both appearance and shape. For example, newly installed water pipes predominantly feature distinct blue colors, while district heating pipes are colored black. However, within the area of one utility owner, the newly installed utilities often consist of the same materials and components, giving them a recognisable look.

Other Utility: Existing utilities uncovered during excavation belonging to other utility owners. This class represents a very

diverse set of utility structures but shares more similarities across all areas in the dataset.

Inactive Utility: Out-of-service utilities belonging to the utility owner. These are left installed in the ground when replacing older counterparts of the utilities in the *Main Utility* class. Due to the large time gap between installations, these older utilities vary in appearance—from closely resembling the newly installed utilities in the *Main Utility* class to resembling structures in the *Other Utility* class. Labelling this class proved challenging and was cross-referenced with utility map data provided by the same utility owners that also provided the point cloud data for the dataset.

Misc.: Misc. trench items such as pipe-like objects, work equipment and left-over cut pipe segments. These items are captured due to the messy nature of an excavation site and are irrelevant to the utility owner.

Table 2 shows the distribution of points for each class across all areas in the dataset. It shows a large class imbalance with the *Trench* class being, by far, the most dominant class. The *Trench* class is considered the background class. In water areas the distribution among the three utility classes is relatively balanced and the *Main Utility* class is consistently represented across areas. A notable difference between the water and heating areas is the higher number of points in the *Main Utility* class in heating areas. This aligns with the fact that district heating utilities are larger than water utilities.

Table 2. The number of point clouds in each area and the number of points per class and the average number of points per point cloud. All numbers except point clouds are stated in thousands. The total number of points in the dataset is 528 million points.

Area	Clouds	Main Utility	Other Utility	Trench	Inactive Utility	Misc	Avg. Points
Water Area 1	68	2,566 (2.6%)	3,377 (3.5%)	90,578 (92.8%)	842 (0.9%)	195 (0.2%)	1,435
Water Area 2	35	1,183 (3.2%)	1,540 (4.1%)	33,735 (90.6%)	680 (1.8%)	80 (0.2%)	1,063
Water Area 3	65	3,380 (3.0%)	4,192 (3.8%)	101,970 (91.5%)	1,860 (1.7%)	64 (0.1%)	1,715
Water Area 4	42	1,946 (3.0%)	3,816 (5.9%)	58,237 (90.4%)	358 (0.6%)	57 (0.1%)	1,534
Water Area 5	50	1,606 (2.2%)	3,403 (4.7%)	66,937 (92.5%)	377 (0.5%)	22 (0.0%)	1,447
Water Total	260	10,682 (2.8%)	16,328 (4.3%)	351,456 (91.8%)	4,118 (1.1%)	418 (0.1%)	1,473
Heat Area 1	29	20,363 (21.3%)	1,605 (1.7%)	72,927 (76.1%)	576 (0.6%)	321 (0.3%)	3,303
Heat Area 2	21	5,191 (10.5%)	535 (1.1%)	43,252 (87.7%)	18 (0.0%)	326 (0.7%)	2,349
Heat Total	50	25,554 (17.6%)	2,140 (1.5%)	116,179 (80.1%)	594 (0.4%)	646 (0.4%)	2,902

4. Evaluation methods on OpenTrench3D

4.1. 3D Semantic Segmentation Methods

We benchmark three state-of-the-art neural networks for semantic segmentation of point clouds: PointNeXt[26], PointVector[9] and PointMetaBase[22] on OpenTrench3D. These methods employ different feature extraction methodologies and achieve state-of-the-art performance on S3DIS[2] and ScanNetV2[8], which are scene-level datasets like OpenTrench3D:

PointNeXt[26]: modernises the highly influential PointNet++[25] with improved training strategies e.g., data augmentation and optimization techniques and introduces the inverted residual bottleneck block which allows for efficient and effective model scaling. We use the XL variant of PointNeXt which has 41.5M parameters.

PointVector[9]: enhances the Point Set Abstraction blocks of the PointNet series, which are responsible for aggregating features of neighboring points. This is done by introducing a Vector-oriented Point Set Abstraction block that can aggregate neighboring features through higher-dimensional vectors instead of simple scalars. We use the XL variant of PointVector which has 24.1M parameters

PointMetaBase[22]: introduces the concept of a PointMeta-building block, which is composed of a four sub-blocks: a neighbor update function, a neighbor aggregation function, a point update function and a position embedding function. PointMetaBase experiments with various settings for each sub-block to identify optimal feature-extraction capabilities while allowing for efficient processing of points. We use the XXL variant of PointMetaBase which has 19.7M parameters.

4.2. Evaluation Metrics

We use two well-established metrics, mean intersection-over-union (mIoU) and mean accuracy (mAcc), for evaluating the 3D semantic segmentation performance of the methods. Intersection-over-union for class i :

$$IoU_i = \frac{TP_i}{TP_i + FP_i + FN_i} \quad (1)$$

Mean intersection-over-union:

$$mIoU = \frac{1}{N} \sum_{i=0}^N IoU_i \quad (2)$$

Accuracy for class i :

$$Acc_i = \frac{TP_i}{TP_i + FN_i} \quad (3)$$

Mean accuracy:

$$mAcc = \frac{1}{N} \sum_{i=0}^N Acc_i \quad (4)$$

4.3. OpenTrench3D Evaluation Description

5-fold cross-validation on water areas: We conduct a 5-fold cross-validation on the water areas in two rounds. Initially, we include the *Inactive Utility* class during training and evaluation and subsequently, we ignore it. This allows us to assess the three methods' effectiveness in distinguishing among the three utility classes. Additionally, the *Misc* class is ignored from both evaluations due to its lack of relevance to the Utility Owner.

Fine-tuning evaluation on heating areas: We conduct a fine-tuning evaluation to assess the transfer learning and fine-tuning opportunities of the dataset and to investigate the generalizability and performance of the three methods on the entire dataset. Specifically, we first pre-train model weights

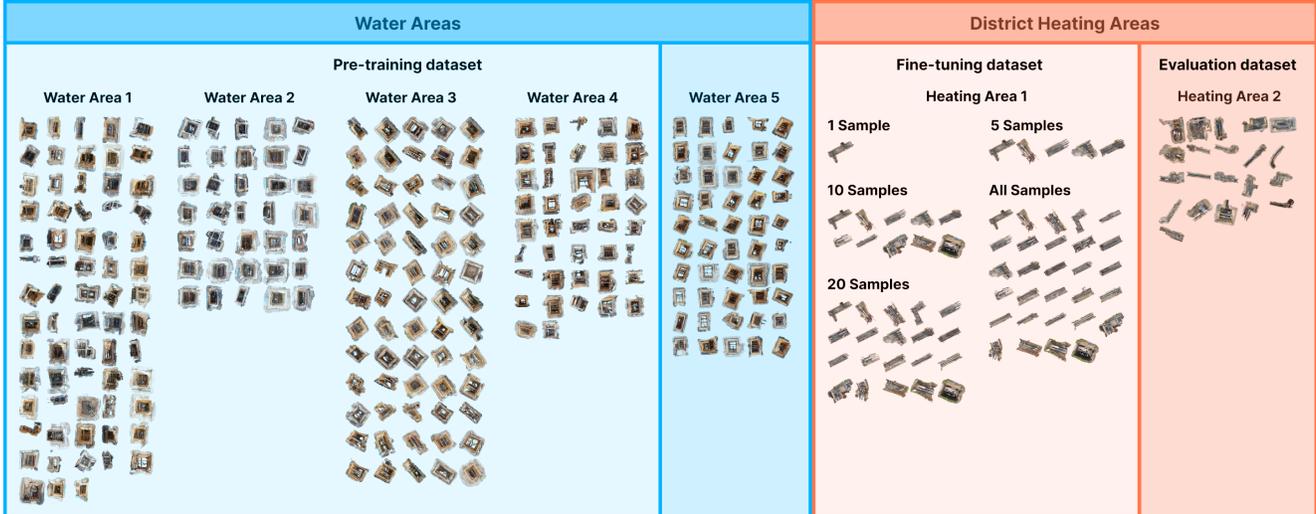


Figure 3. An overview of the dataset sub-sets of OpenTrench3D used in the 5-fold cross-validation on water areas and fine-tuning evaluation on heating areas as described in section 4.3.

by training on Water Area 1-4 while using Water Area 5 as a 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 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.

An overview of the dataset subsets is presented in figure 3. In the fine-tuning evaluation, we ignore the *Inactive Utility* class due to its minimal presence in heating area 2. The *Misc* class is ignored in all experiments.

4.4. Training Environment

The experiments for the 5-fold evaluation on water areas and fine-tuning evaluation on heating areas are run on a high-performance computing cluster with Nvidia A40 GPUs (48GB). In each experiment, the deep learning networks are assigned 64GB RAM, 8 CPU Cores (AMD EPYC) and 1 Nvidia A40 GPU. For the development environment, we use Ubuntu 20.04, python 3.10 and torch 1.10 with CUDA 11.3.

Pre-training experiments are run for 50 epochs with a batch size of 8. For each epoch of training we loop over the dataset 30 times similar to [26]. The voxel max, which specifies the number of processed points per point cloud per batch, is set to 24,000. We use an Adamw optimizer with an initial learning rate of 0.01 and a minimum learning rate of 0.0001. Fine-tuning experiments are run for 50 epochs with a batch size of 1 when fine-tuning on 1 and 5 samples, a batch of 4 when fine-tuning on 10 samples and a batch size of 8 when fine-tuning on 20 and 29 samples. The learning

rate is set to 0.05 and minimum learning rate of 0.0001. We universally down-sample point clouds using a voxel grid size of 0.02m and set the query ball radius to 0.025m. These settings are derived from using data-driven hyperparameter-tuning [17] and an additional hyperparameter sweep.

Table 3. Results from 5-fold cross-validation on water areas as described in 4.3. The mACC, mIoU and IoU score for each class.

Method	mAcc	mIoU	Main	Other	Trench	Inactive
PointNeXt	79.7	70.6	81.6	59.7	98.1	42.8
PointVector	84.1	76.5	83.1	77.7	98.6	46.4
PointMetabase	84.5	75.8	83.6	76.6	98.6	44.5
PointNeXt	91.9	88.5	87.8	79.1	98.7	-
PointVector	93.2	90.3	90.3	81.8	98.8	-
PointMetabase	93.8	90.5	90.4	82.2	98.8	-

5. Results and Discussion

5.1. 5-fold cross-validation on Water Areas

We present the outcomes of our 5-fold cross-validation in table 3. The *Inactive Utility* class presents the most significant challenge, as expected, likely due to its visual similarities to *Main Utility* and *Other Utility* as visualized in zoomed in class examples of figure 2. In our qualitative analysis, we note that the Inactive class is often misclassified as either Main or Other Utility. For illustrative examples of this, we direct readers to the supplementary materials.

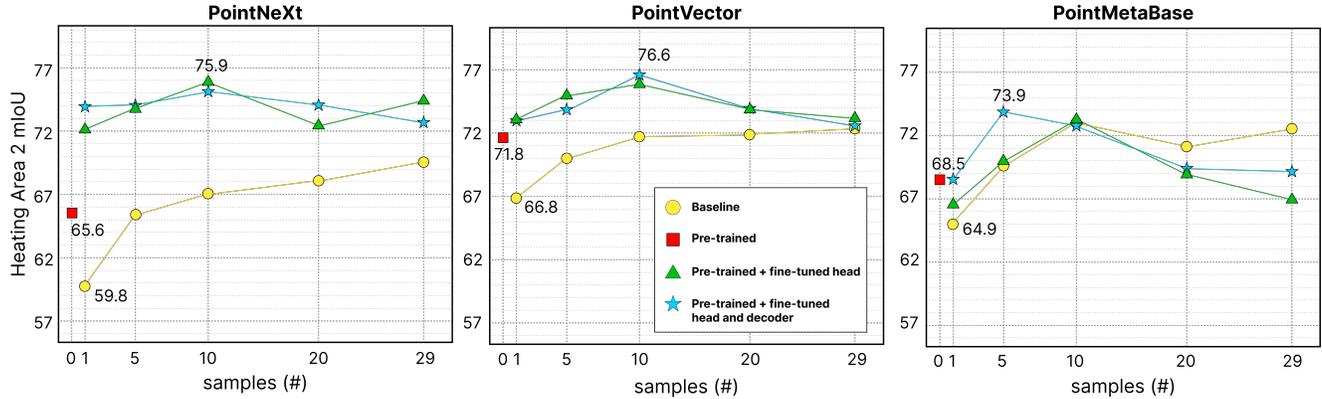


Figure 4. Results from the fine-tuning experiments described in section 3.2. For comparison, we display results from pre-trained models (red square) and results of models solely trained on samples from Heating Area 1 (Baseline).

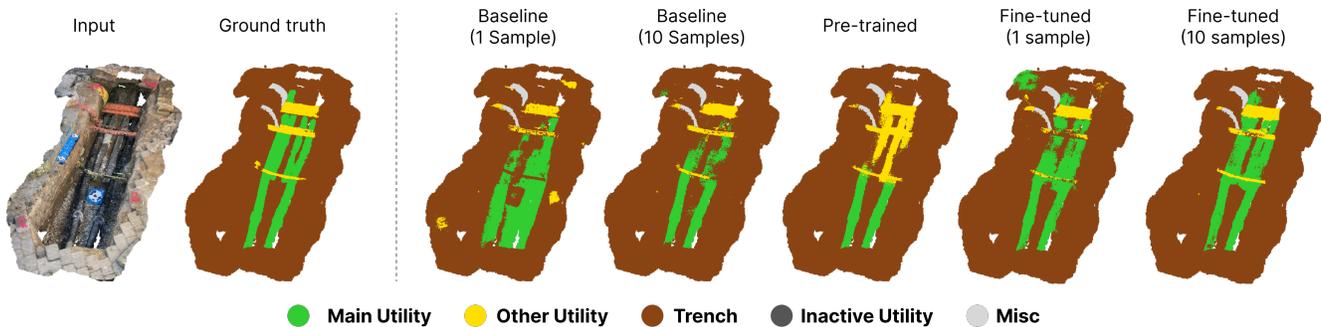


Figure 5. Qualitative results of pre-trained PointNeXt model, PointNeXt models trained on 1 and 10 samples from Heat Area 1 and finally pre-trained and fine-tuned PointNeXt models fine-tuned on 1 and 10 samples were only weights of the segmentation head are tuned. A set of additional examples are seen in the supplementary material.

Other Utility emerges as the second most challenging class, whereas *Main Utility* is the easiest to classify. Remarkably, in the second evaluation where *Inactive Utility* is ignored, the IoU score of *Main Utility* exceeds 90% for both PointVector and PointMetaBase, a level of performance that is difficult to enhance further and inspire the possibilities to build application upon for automating utility mapping. Our qualitative findings reveal that errors predominantly occur in boundary areas between two classes. In our dataset, this particularly happens in areas where one of the utility classes merge into the soil e.g., the *Trench* class or when soil partially covers the utilities.

Generally, we notice a better performance in PointVector and PointMetaBase compared to PointNeXt. This follows a similar trend in benchmarks on S3DIS dataset[9, 22].

5.2. Fine-tuning evaluation on Heating Areas

We present the results from the fine-tuning evaluation on heating areas in figure 4 and table 4. Generally, we observe that fine-tuning of pre-trained weights proves effective across all three methods evaluated, particularly when fine-tuning with 1, 5 and 10 samples from Heating Area 1.

We find PointVector performs the best, achieving a mIoU score of 76.6% when fine-tuning the weights of both the segmentation head and the decoder with 10 samples from Heating Area 1, closely followed by PointNeXt with a mIoU score of 75.9%. Furthermore, PointVector consistently achieves the highest mIoU scores when pre-trained only on water areas. This supports the authors arguments that the vector-oriented Point Set Abstraction design introduces less inductive bias, which in returns improves generalization capabilities [9]. One notable results, is the fact that for both PointNeXt and PointVector, fine-tuning pre-trained weights with just a single sample from Heating Area 1 achieves a better performance compared to baseline version trained on even all 29 samples from Heating Area 1. In contrast, PointMetaBase ranks as the least effective, with its baseline model trained on 20 and 29 samples even surpassing the performance of its fine-tuned counterparts.

Unexpectedly, we observe a negative trend upon increasing the number of fine-tuning samples to 20 and 29 for all three methods. This prompts the hypothesis that further experimental investigations into the optimization of hyper-parameter settings for fine-tuning with 20 and 29 samples to

Table 4. Results from fine-tuning evaluation on heating areas as described in 4.3 for PointNeXt. The mACC, mIoU and IoU score for each class is shown. We refer to the supplementary material for corresponding tables for PointVector and PointMetaBase.

PointNeXt	Samples	mAcc	mIoU	Main	Other	Trench
Baseline	1	68.3	59.8	65.1	20.4	93.8
	5	72.8	65.6	78.7	21.3	96.9
	10	72.7	67.1	75.8	29.2	96.3
	20	76.2	68.1	75.9	32.1	96.4
	29	75.5	69.6	79.5	32.2	96.9
Pre-trained		75.0	65.6	70.6	29.2	97.0
Fine-tuned (Head)	1	80.4	72.2	76.4	43.9	96.5
	5	79.8	73.8	81.2	43.0	97.2
	10	83.7	75.9	83.6	46.5	97.6
	20	79.9	72.3	83.0	36.4	97.6
	29	84.5	74.4	81.4	44.2	97.5
Fine-tuned (Decoder+Head)	1	80.8	73.9	80.2	44.4	97.2
	5	78.9	74.0	81.3	43.4	97.2
	10	82.0	75.1	83.8	43.9	97.6
	20	83.9	74.1	82.4	42.4	97.5
	29	82.3	72.7	80.0	40.6	97.5

effectively capitalize on the additional data.

Compared to the 5-fold cross-validation on water areas, the mIoU score is significantly lower across all types of experiments, suggesting that achieving such performance likely requires more annotated data. However, this trend is less pronounced for the IoU score of the *Main Utility* class, which still manages to achieve impressively high performance in the fine-tuning evaluation. PointVector emerges as the top performer with an IoU score of 86.7%, closely followed by PointNeXt with an IoU of 83.8%, as seen in Table 4. Both scores are relatively close to their counterparts in table 3. Figure 5 clearly illustrates the significant improvement that fine-tuning brings to the *Main Utility* class. When comparing models pre-trained solely on water areas to those fine-tuned on just 1 and 10 samples, a clear improvement is evident. The fine-tuning process drastically reduces the instances of the *Main Utility* being incorrectly labeled as *Other Utility*, achieving near-perfect results for both classes. The incorrect classification of the *Main Utility* as *Other* is further observed in our qualitative analysis and was expected, considering the visual and dimensional differences between water utilities and district heating pipes, as illustrated in Figure 2.

Moreover, our qualitative analysis reveals surprising classification errors, in which the *Trench* class is incorrectly identified as the *Main Utility* class, particularly in our fine-tuned and baseline models, and significantly less so in the

model pre-trained on water areas. While the exact cause is uncertain, we hypothesize that the large difference in number of points, percentage-wise, in Heating Area 1 compared to other areas, may be a factor. Additionally, we observed that many of the heating pipes in Heating Area 1 are partially, and sometimes almost completely, covered with soil, which might contribute to the confusion.

Upon examining the incorrect predictions, the error often appears obvious to the human eye. For instance, where a single continuous pipe is classified as both *Main-* and *Other Utility* classes at different segments of the pipe. We believe this calls for post-processing, possibly integrated with some prior knowledge about class appearances or human interaction related to the work by Stranner et al. [30].

Finally, from the fine-tuning evaluation experiment, it has been discerned that the batch size has a significant influence on the methods' performance, particularly when fine-tuning with few samples e.g. 5 and 10. For example, upon fine-tuning a pre-trained PointVector model on 10 samples with a batch size of 1 over 50 epochs, a mIoU of 66.6% was attained, whereas under identical conditions, when adjustment to a batch size of 4, the mIoU score was 76.6%. A difference of 10.0 percentage points. Comparable trends were observed for PointMetaBase and PointNeXt.

6. Conclusion

In this paper, we introduce OpenTrench3D, the first publicly available point cloud dataset of underground utilities from open trenches. It features 310 fully annotated point clouds consisting of a total of 528 million points categorised into 5 unique classes following a utility owner-centric classification scheme. The dataset is acquired through photogrammetric techniques, offering an accessible and cost-effective method for capturing underground utility data. Our experiments on OpenTrench3D involve: (i) a 5-fold cross-validation on water project areas, highlighting impressive performance for the main utility classes using state-of-the-art deep learning semantic segmentation methods and (ii) fine-tuning pre-trained models on district heating utilities with minimal samples, demonstrating the capability to learn segmenting entirely new utility classes effectively from just a few examples. By providing this dataset, we aim to catalyze innovation and advance the research in the domain, ultimately contributing to reducing excavation damages, optimizing urban planning and improving underground utility surveying and mapping.

7. Acknowledgement

We thank the students who assisted in data annotation, and to IT34, Ambolt AI, Innovation Fund Denmark, Digital-Lead and AI Denmark for their support and funding. Special thanks to Kalundborg and Novafos Utility Company for granting data access.

References

- [1] Common Ground Alliance. Dirt annual report of the common ground alliance, 2019. Visited online 3 Mar 2024. **1**
- [2] 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. **3, 5**
- [3] 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 *ICCV*, 2019. **2, 3**
- [4] Richard Broome, Phil Cornforth, Stephen Crossland, Nicole Metje, and Andy Rhoades. Utility strike damage report, 2020. Visited online 3 Mar 2024, https://utilitystristg.wpengine.com/wp-content/uploads/2022/12/usag_2020_utility_strike_damages_report.pdf. **1**
- [5] Gülcan Can, Dario Mantegazza, Gabriele Abbate, Sébastien Chappuis, and Alessandro Giusti. Semantic segmentation on swiss3dcities: A benchmark study on aerial photogrammetric 3d pointcloud dataset. *Pattern Recognition Letters*, 150:108–114, 2021. **2, 3**
- [6] Meida Chen, Qingyong Hu, Zifan Yu, Hugues Thomas, Andrew Feng, Yu Hou, Kyle McCullough, Fengbo Ren, and Lucio Soibelman. Stpls3d: A large-scale synthetic and real aerial photogrammetry 3d point cloud dataset. In *BMVC*, 2022. **2, 3**
- [7] CloudCompare. CloudCompare (version 2.12). [GPL software]. **3**
- [8] 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 2432–2443, 2017. **3, 5**
- [9] Xin Deng, WenYu Zhang, Qing Ding, and XinMing Zhang. Pointvector: A vector representation in point cloud analysis. In *CVPR*, pages 9455–9465, 2023. **5, 7**
- [10] Timo Hackel, N. Savinov, L. Ladicky, Jan D. Wegner, K. Schindler, and M. Pollefeys. Semantic3d.net: A new large-scale point cloud classification benchmark. In *ISPRS*, pages 91–98, 2017. **2, 3**
- [11] Lasse H. Hansen, Philipp Fleck, Marco Stranner, Dieter Schmalstieg, and Clemens Arth. Augmented reality for subsurface utility engineering, revisited. *IEEE Transactions on Visualization and Computer Graphics*, 27(11):4119–4128, 2021. **3**
- [12] Lasse H. Hansen, T. M. Pedersen, E. Kjems, and S. Wyke. Smartphone-based reality capture for subsurface utilities: Experiences from water utility companies in denmark. *ISPRS, XLVI-4/W4*:25–31, 2021. **2, 3**
- [13] L. H. Hansen, R. van Son, A. Wieser, and E. Kjems. Addressing the elephant in the underground: An argument for the integration of heterogeneous data sources for reconciliation of subsurface utility data. in *ISPRS, XLVI-4/W4*:43–48, 2021. **1**
- [14] Joakim B. Haurum and Thomas B. Moeslund. A survey on image-based automation of CCTV and SSET sewer inspections. 111:103061, 2020. **2**
- [15] Joakim B. Haurum and Thomas B. Moeslund. Sewer-ML: A multi-label sewer defect classification dataset and benchmark. In *CVPR*, pages 13451–13462, 2021. **2**
- [16] Qingyong Hu, Bo Yang, Sheikh Khalid, Wen Xiao, Niki Trigoni, and Andrew Markham. Towards semantic segmentation of urban-scale 3d point clouds: A dataset, benchmarks and challenges. In *CVPR*, 2021. **2, 3**
- [17] Simon B. Jensen, Galadrielle H.-R., Andreas Møgelmoose, and Thomas B. Moeslund. Data-driven hyperparameter tuning for point-based 3d semantic segmentation. In *ICIPCW*, pages 3696–3700, 2023. **6**
- [18] Tatjana Kutzner, Ihab Hijazi, and Thomas H. Kolbe. Semantic modelling of 3d multi-utility networks for urban analyses and simulations: The CityGML utility network ADE. *International Journal of 3-D Information Modeling*, 7(2):1–34, 2018. **3**
- [19] Michael Kölle, Dominik Laupheimer, Stefan Schmoehl, Norbert Haala, Franz Rottensteiner, Jan Dirk Wegner, and Hugo Ledoux. The hessigheim 3d (h3d) benchmark on semantic segmentation of high-resolution 3d point clouds and textured meshes from uav lidar and multi-view-stereo. *ISPRS*, 1:11, 2021. **3**
- [20] Maosu Li, Yijie Wu, Anthony G.O. Yeh, and Fan Xue. Hrhdk: A benchmark dataset of high-rise and high-density urban scenes for 3d semantic segmentation of photogrammetric point clouds. In *ICIPCW*, pages 3714–3718, 2023. **2**
- [21] Xinke Li, Chongshou Li, Zekun Tong, Andrew Lim, Junsong Yuan, Yuwei Wu, Jing Tang, and Raymond Huang. Campus3d: A photogrammetry point cloud benchmark for hierarchical understanding of outdoor scene. In *ACMMM*. Association for Computing Machinery, 2020. **2**
- [22] H. Lin, X. Zheng, L. Li, F. Chao, S. Wang, Y. Wang, Y. Tian, and R. Ji. Meta architecture for point cloud analysis. In *CVPR*, pages 17682–17691, 2023. **5, 7**
- [23] Sergey Lytkin, Vladimir Badenko, Alexander Fedotov, Konstantin Vinogradov, Anton Chervak, Yevgeny Milanov, and Dmitry Zotov. Saint petersburg 3d: Creating a large-scale hybrid mobile lidar point cloud dataset for geospatial applications. *Remote Sensing*, 15(11), 2023. **2**
- [24] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *CVPR*, pages 652–660, 2017. **3**
- [25] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In *NIPS*, 2017. **3, 5**
- [26] Guocheng Qian, Yuchen Li, Houwen Peng, Jinjie Mai, Hasan Abed Al Kader Hammoud, Mohamed Elhoseiny, and Bernard Ghanem. Pointnext: Revisiting pointnet++ with improved training and scaling strategies. In *NIPS*, 2022. **5, 6**
- [27] Xavier Roynard, Jean-Emmanuel Deschaud, and François Goulette. Paris-lille-3d: A large and high-quality ground-truth urban point cloud dataset for automatic segmentation and classification. *The International Journal of Robotics Research*, 37(6):545–557, 2018. **2**
- [28] Andrés Serna, Beatriz Marcotegui, François Goulette, and Jean-Emmanuel Deschaud. Paris-rue-Madame database: a 3D mobile laser scanner dataset for benchmarking urban de-

- tection, segmentation and classification methods. In *International Conference on Pattern Recognition, Applications and Methods*, 2014. 3
- [29] Hamin Song, Kichun Jo, Jieun Cho, Youngrok Son, Chansoo Kim, and Kwangjin Han. A training dataset for semantic segmentation of urban point cloud map for intelligent vehicles. *ISPRS*, 187:159–170, 2022. 2
- [30] Marco Stranner, Philipp Fleck, Dieter Schmalstieg, and Clemens Arth. Instant segmentation and fitting of excavations in subsurface utility engineering. *TVCG*, pages 1–11, 2024. 3, 8
- [31] Weikai Tan, Nannan Qin, Lingfei Ma, Ying Li, Jing Du, Guorong Cai, Ke Yang, and Jonathan Li. Toronto-3D: A large-scale mobile lidar dataset for semantic segmentation of urban roadways. In *CVPRW*, pages 202–203, 2020. 2, 3
- [32] R. Van Son, S. W. Jaw, J. Yan, V. Khoo, R. Loo, S. Teo, and G. Schrotter. A framework for reliable 3d underground utility mapping for urban planning. *ISPRS*, XLII-4/W10:209–214, 2018. 1, 2
- [33] R. Van Son, S. W. Jaw, and A. Wieser. A data capture framework for improving the quality of subsurface utility information. *ISPRS*, XLII-4/W15:97–104, 2019. 1
- [34] Nina Varney, Vijayan K Asari, and Quinn Graehling. Dales: A large-scale aerial lidar data set for semantic segmentation. In *CVPRW*, pages 186–187, 2020. 2, 3
- [35] Simon Quaade Vinther, Frida Dalbjerg Kunnerup, Lars Bodum, Lasse H. Hansen, and Simon Wyke. Visualisation of 3d uncertainties for subsurface infrastructure using augmented reality. In *International 3D GeoInfo Conference*, pages 111–128. Springer, 2023. 1
- [36] Jingya Yan, Rob Van Son, and Kean Huat Soon. From underground utility survey to land administration: An underground utility 3d data model. *Journal of Land Use Policy*, 102:105267, 2021. 3
- [37] Zhen Ye, Yusheng Xu, Rong Huang, Xiaohua Tong, Xin Li, Xiangfeng Liu, Kuifeng Luan, Ludwig Hoegner, and Uwe Stilla. Lasdu: A large-scale aerial lidar dataset for semantic labeling in dense urban areas. *ISPRS*, 9(7), 2020. 2
- [38] RZM Yuen and Jan Boehm. Potential of consumer-grade cameras and photogrammetric guidelines for subsurface utility mapping. *ISPRS*, 48:243–250, 2022. 2