



Semantic segmentation on Swiss3DCities: A benchmark study on aerial photogrammetric 3D pointcloud dataset

Gülcan Can^{a,*}, Dario Mantegazza^b, Gabriele Abbate^b, Sébastien Chappuis^a,
Alessandro Giusti^b

^a Nomoko AG, Badenerstrasse 790, 8048 Zurich, Switzerland

^b The Dalle Molle Institute for Artificial Intelligence (IDSIA), Via la Santa 1, 6962 Lugano-Viganello, Switzerland

ARTICLE INFO

Article history:

Received 28 December 2020

Accepted 5 June 2021

Available online 8 July 2021

Edited by : Prof. S. Sarkar

MSC:

68R01

51-11

68T10

68T45

68T07

Keywords:

Photogrammetry

Pointcloud

3D Data

Semantic segmentation

Deep learning

Model generalization

ABSTRACT

We introduce a new outdoor urban 3D pointcloud dataset, covering a total area of 2.7km², sampled from three Swiss cities with different characteristics. The dataset is manually annotated for semantic segmentation with per-point labels, and is built using photogrammetry from images acquired by multirotors equipped with high-resolution cameras. In contrast to datasets acquired with ground LiDAR sensors, the resulting point clouds are uniformly dense and complete, and are useful to disparate applications, including autonomous driving, gaming and smart city planning. As a benchmark, we report quantitative results of PointNet++, an established point-based deep 3D semantic segmentation model; on this model, we additionally study the impact of using different cities for model generalization.

© 2021 Elsevier B.V. All rights reserved.

1. Introduction

Many recent achievements of deep learning depend on the availability of large labeled training datasets [10,27], such as ImageNet [9] for image classification and MS COCO [18] for image segmentation. In this work, we propose a new dataset of dense urban 3D pointclouds, spanning 2.7km², acquired using photogrammetry from three cities in Switzerland (Zurich, Zug and Davos). The entire dataset is manually annotated with dense labels, which associate a point to one of five categories: terrain, construction, vegetation, vehicle, and urban asset.

The main goal of the dataset is to train semantic segmentation algorithms for urban environments. Semantic segmentation consists in partitioning the data into multiple sets of points, such that each set represents only objects of a given type. The problem is rel-

evant for many real-world applications, such as autonomous driving, content generation for games [15], augmented reality applications, and city planning [35].

Most existing datasets [3,25,30] for outdoor 3D semantic segmentation are motivated by real-time autonomous driving applications, and are therefore acquired at low resolution by street-level Light Detection and Ranging (LiDAR) sensors; this yields incomplete point clouds (for example, areas far from roads, such as roofs, are either not acquired or acquired with very low resolution) which are unsuitable for applications such as city planning, urban augmented or virtual reality (AR/VR), or gaming. In contrast, we acquire high-resolution photographs from unmanned aerial vehicles (UAV) flying on a grid pattern over the area of interest, then reconstruct the 3D shape using photogrammetry; this allows us to densely acquire most outdoor surfaces. Similar approaches have been previously adopted for several applications, including automatic urban area mapping [20], damage detection [19], and cultural heritage site mapping for digital preservation [22]. Compared to 3D models built by satellite-borne cameras, this approach yields models with higher-resolution geometry and texture.

* Corresponding author.

E-mail addresses: gulcan@nomoko.world (G. Can), dario.mantegazza@idsia.ch (D. Mantegazza), gabriele.abbate@idsia.ch (G. Abbate), sebastien@nomoko.world (S. Chappuis), alessandro@idsia.ch (A. Giusti).

Table 1

Comparison of our datasets with the common 3D pointcloud datasets in the literature. The spatial size denotes either the covered area or the driven length (in the case of mobile LiDAR captures) for a dataset.

Paper	Data Acquisition	Spatial Size	City	Cat.	RGB	Points
SemanticKITTI[3]	Mobile LiDAR	$39.2 \times 10^3 m$	Karlsruhe	25	No	4549M
Paris-Lille3D[25]	Mobile LiDAR	$1.94 \times 10^3 m$	Paris, Lille	9	No	143M
Toronto3D[30]	Mobile LiDAR	$1 \times 10^3 m$	Toronto	8	Yes	78.3M
Semantic3D[12]	Static LiDAR	-	St. Gallen	8	Yes	4009M
ISPRS[21]	Aerial LiDAR	-	Vaihingen	9	No	1.2M
DublinCity[38]	Aerial LiDAR	$2 \times 10^6 m^2$	Dublin	13	No	260M
DALES[31]	Aerial LiDAR	$10 \times 10^6 m^2$	Surrey (BC)	8	No	505M
SenSatUrban[16]	UAV Photogrammetry	$6 \times 10^6 m^2$	Birmingham, Cambridge	13	Yes	2847M
Swiss3DCities (this paper)	UAV Photogrammetry	$2.7 \times 10^6 m^2$	Zurich, Zug, Davos	5	Yes	226M (Sparse: 7.5M, Dense: 3147M)

High resolution data yields more accurate models, but also aids the segmentation task because it contains more information to discriminate between different classes; currently, state-of-the-art models for 3D semantic segmentation rely on deep learning [4,11,36] and represent input data as voxels [37], points [23] or meshes [17]; other approaches render multiple views of the 3D scene and then rely on 2D semantic segmentation models [5,28,29], which can be trained on more abundant 2D labeled semantic segmentation datasets.

To show the potential of our dataset for training and evaluating segmentation algorithms, we consider the well-established PointNet++ model [23,24] and report its performance when using different splits for training and evaluation. In particular, the performance of machine learning models depends not only on the size of the training dataset, but also on how representative it is of the evaluation data: often, models trained on large amounts of data from a given environment fail to generalize to a different target environment. Because our dataset contains data from three cities with different characteristics, it can be used to explore this fundamental aspect.

The rest of the paper is organized into five sections. We first describe related commonly-used datasets for 3D semantic segmentation in Section 2. Then, in Section 3 we present our **main contribution**: a new pointwise labeled multi-city dataset for semantic segmentation of outdoor 3D point clouds, which we release to the research community in three versions with different point densities¹; we characterize the dataset and describe data acquisition, processing and manual labeling pipelines. In Section 4 we describe the applied deep learning model to demonstrate semantic segmentation task on our dataset. We discuss quantitative results in Section 5, where we also explore the model’s generalization ability across different cities (**secondary contribution**). Section 6 concludes the paper.

2. Related work

This section summarizes relevant pointcloud datasets with semantic segmentation labels (see Table 1). One fundamental difference among the datasets is their acquisition modality, i.e. LiDAR or photogrammetry.

2.1. LiDAR datasets

A lot of recent research efforts are related to autonomous driving applications: in particular, recognizing and segmenting roads and relevant urban elements from images or 3D point clouds acquired by the car itself. In this context, laser scanning systems,

e.g. Velodyne HDL-64E [32], are commonly used to acquire high-accuracy LiDAR pointcloud sequences from a car’s point of view. Paris-Lille[25], Semantic KITTI[3], and Toronto3D dataset[30] are among such large-scale datasets with pointwise semantic labels.

Due to the low-lying viewpoint and focus on driving-related segmentation tasks, these mobile LiDAR datasets show incomplete point clouds: e.g. the upper floors or roofs of the buildings are usually not captured. Even though these datasets serve their main scope very well, they are not suitable for other applications, such as urban planning.

Semantic3D dataset[12] is a large-scale pointcloud dataset with per-point semantic labels. This dataset is acquired via a static terrestrial laser scanning system in the north-east of Switzerland. Several points to note about this dataset are gaps due the occlusions (also known as LiDAR shadows), moving object artifacts, and varying point density based on the distance of the laser system to each surface or object in the scene.

As captured from air (either from a UAV or a helicopter), aerial LiDAR datasets such as ISPRS airborne LiDAR pointcloud dataset [21], DublinCity dataset [38] and DALES dataset [31] are also relevant in our context. One important difference of these datasets with respect to ours is that, due to the narrow divergence of laser beams, they can sometimes capture ground samples even when covered by vegetation. Compared to the DublinCity aerial LiDAR dataset[38], ours covers a moderately larger area and, in the medium-density version, has a similar point density. This shows the relevance of our contribution with respect to existing aerial LiDAR datasets.

Photogrammetric pointclouds

Sun3D [34] and Stanford Large-Scale Indoor Spaces 3D (S3DIS) [1] are commonly used pointcloud datasets acquired using Structure-from-Motion (SfM) 3D reconstruction techniques. These datasets are focused on indoor scenes, and present interesting challenges for computer vision research, such as the presence of clutter, and relevant context around different objects, that can play a role in scene understanding. Due to their limited extent, the capturing process is much less challenging than in large-scale outdoor contexts, which also need to account for variability of weather, illumination conditions, and scales of represented objects.

The Pix4D dataset [2] comprises of aerial photogrammetric pointclouds from three outdoor scenes with different distributions of urban surfaces or objects. The authors emphasize the importance of color features apart from geometric features to classify these pointclouds into 6 semantic classes. This dataset is relatively small-scale, since it comprises of only three scenes with a total of 18.2 million points.

The SenSatUrban dataset [16] is also reconstructed via photogrammetry from aerial photographs. The photographs are taken with a UAV that follows a double-grid flight path and covers a

¹ <https://github.com/NomokoAG/Swiss3DCities>

6km² area in three cities in UK (Birmingham, Cambridge, and York). Pointwise semantic labels in 13 categories are available for these pointclouds. As an urban-focused aerial photogrammetric pointcloud, the SenSatUrban dataset is the most relevant with respect to our contribution. SenSatUrban covers an approximately twice large area, and uses 13 categories instead of five; its point density is higher than our medium-density version, but lower than our high-density version.

3. Dataset description

We describe the process used to produce our large scale aerial photogrammetry dataset, covering both acquisition of source photographs and processing to obtain 3D point clouds. We conclude the section by detailing the data characteristics.

3.1. Data acquisition

The image data is acquired via a high-resolution camera array (nadir and oblique cameras) mounted on a multirotor drone.

To capture the image data, the drone is configured to trigger the cameras simultaneously at regular intervals while following a double grid flight path [26].

Each flight acquires one or more tiles. Each tile corresponds to 412m × 412m horizontal area approximately (around 17 hectares). The Ground Sampling Distance (GSD), i.e. the inter-pixel distance measured on the ground, is measured as 1.28cm.

3.2. Data processing

After aerial image acquisition, we follow a classic photogrammetry workflow to reconstruct textured 3D models, based on RealityCapture [6,14], a commercial software.

We first estimate the global camera poses of the captured images and a georeferenced sparse point cloud of the scene using a standard SfM process (referred as "alignment" in RealityCapture). Georeferencing is achieved using Ground Control Points (GCPs) and RealityCapture's GCP annotation tool. Taking into account the drone-based acquisition described above, the GCP annotation, and the further processing of the data, we measure a total georeferencing root mean square error (RMSE) of 5.45cm horizontally and 11.60cm vertically. This implies that the data is scaled to real world units, meters in our case. For convenience, we provide all point coordinates as zero-centered per tile in our dataset.

Once the data is aligned and georeferenced, we reconstruct a dense mesh constrained to the geographic region of the tile. The raw mesh can contain up to half a billion polygons for a single tile. To get a mesh with a more manageable size, we simplify it to a maximum number of 30 million polygons, i.e. approximately 15 million vertices, with Reality Capture and texture it using the captured drone images. The output point cloud used for segmentation is composed of the vertices of such a mesh; the RGB color of each point is sampled from the mesh texture.

3.3. Manual segmentation

Our pointclouds are segmented manually into the five semantic categories: terrain, construction, vegetation, vehicles and urban assets. 3D artists complete this task using off-the-shelf 3D modeling software (such as Blender [8]); to make the process manageable, they work on each tile individually, and operate on a 1-million polygon mesh further simplified from the initial mesh. It takes between six to twelve hours for a 3D artist to manually segment each tile.

Labels are then transferred from the simplified mesh to the output point cloud. The label of each point in the output point

cloud is assigned by finding the nearest neighbor in the segmented mesh. We used an adaptive distance threshold to avoid matching outlier points. We found that this method gives satisfying results for the final segmentation of the point cloud while keeping the amount of manual work needed at a manageable level.

3.4. Dataset details

The dataset represents sixteen tiles acquired from three cities in Switzerland (see Fig. 1): six tiles from Zurich, five tiles from Zug and five tiles from Davos.

For each tile, the dataset contains pointclouds at three resolutions, i.e. approximately 500K, 15M, and 225M points per tile as shown in Fig. 2. Both 500K and 15M point density pointclouds have x,y,z, and RGB color features. For the highest density, we have only x,y,z coordinates. In the rest of this paper, we only consider the medium (15M) point density.

Classes and class distribution As our dataset is focused on urban areas in Switzerland, it comprises of a large amount of terrain, building, and vegetation. Even though many objects of other categories (such as vehicles or urban assets) are present in our dataset, they amount to a relatively small portion of the points because each object is relatively small. Therefore, we divide our semantic labeling to only five main categories: 1) terrain (including natural terrain, e.g. grass or soil, impervious terrain, e.g. road or sidewalk, and water areas, e.g. river or lake); 2) building; 3) urban asset (including traffic light, pole, crane, public transportation stop, trash bin, etc.); 4) vegetation (tree or bush); and 5) vehicle (car, bike, scooter, etc.). The total number of points per category are as follows: terrain - 64,491,349, construction - 104,144,973, urban asset - 2,681,512, vegetation - 52,524,820, and vehicle - 2,730,595.

4. Semantic segmentation

To provide baseline performance metrics, we report experiments using PointNet++ [24], a well-established pointcloud segmentation approach.

4.1. PointNet++

PointNet++ [24] is a deep learning model built upon the PointNet [23] model. In the PointNet++ architecture, PointNet module is used as a local feature encoder and applied in a nested fashion to learn hierarchical features. Moreover, PointNet++ uses farthest point sampling to cover more representative points during sampling.

We adopt an existing implementation [33] of PointNet++. For a given instance, the input of the model is a $N \times 6$ matrix, each row containing the x, y, z coordinates and R, G, B color of one of N input points. The output of the model is a matrix of $N \times K$ prediction probabilities, where $K = 5$ is the number of classes. Because the model is designed to handle input pointclouds up to a few thousand points ($N = 8192$ in our reference implementation), it cannot be directly applied to our large outdoor datasets; therefore, we implemented the following data pipeline. First, we partition the input data into columns with a base of $10m \times 10m$ and infinite height. During training and validation, each instance is generated by picking a column, then randomly sampling (with replacement) N points from the column. A training epoch is obtained by generating one instance per column. For every epoch, the instances are sampled again; this yields a form of data augmentation since for each column a different subset of points is sampled.

Once a model is trained, in order to segment a testing tile, we apply the model to every column separately, then merge the segmentation results. To segment a column, we randomly divide the points in the column in subsets, each containing exactly N points;



Fig. 1. A map of (a) highly-populated urban Zurich center, (c) rural and industrial Zug-Cham region, and (d) mountainous Davos region.




Fig. 2. (a) A part of the simplified mesh with 1M polygons, (b) the medium point cloud with 15M vertices, (c) the semantic groundtruth, and (d) the dense pointcloud with around 220M points for PC 1 (Davos 16_34318_-22950).

for the last subset, in case less than N points are remaining, additional points are sampled from the other subsets. Each subset defines an instance and is segmented independently using the trained model; the results are then combined.

The model is trained by minimizing the cross-entropy loss; to deal with heavy class imbalance, following in similar works [16], the loss is weighted differently for each class, according to inverse-square-root frequency. A training batch is composed by 64 instances and we train for 200 epochs; we do not use early stopping but snapshot the model which yields the minimum loss on the validation set (which is defined on tiles different than training and testing tiles, see below). Other hyper-parameters are set as in [33]. The experiments are run on a NVIDIA RTX 2080Ti GPU. The longest train and test sessions are completed in less than 6 hours.

4.2. Experimental setup

Our experimental setup focuses on the following research questions, that are more related to the characteristics of the data than to the capabilities of the specific model.

- Which categories are more challenging to segment?
- How does the model generalize across cities?
- How much can additional data help even if it is from a different city?
- Which training strategy is better for pointcloud data: an ensemble of per-city models or a single model trained on all data from multiple cities?

To answer the questions above, we train four models: three on data sampled from a single city (named *single-city* models in the following); one on data from all three cities. Then, we apply each model on three testing sets (disjoint from the training and validation sets), one per city.

Data Splits

We consider five tiles for each of the three cities. Each tile covers approximately 0.17km^2 , which yields 0.855km^2 and 70 million points per city.

For each city, the five tiles are partitioned in three tiles for training, one tile for validation and one tile for testing. Single-city models are therefore trained on three tiles and validated on one tile. The model trained on all cities is trained on nine tiles and validated on three tiles. Each of the four models is tested on three

tiles (one per city), on which we separately compute performance metrics.

Evaluation Metrics For a testing tile, a model produces five class probabilities (which sum to 1) for each point. The point is then assigned to the class that has the largest probability. From these data, we compute the following commonly-used metrics [1,3,12] to quantify segmentation performance.

Overall Accuracy is the fraction of the points for which the predicted class coincides with the ground truth class (also known as micro-averaged accuracy).

Weighted Accuracy is the macro-averaged accuracy that is multiplied with a per-class factor. For a given class c , the factor is computed as the proportion of the number of class samples N_c over the number of samples in the whole dataset N , i.e. N_c/N .

Per-class F1 score is the harmonic mean between per-class precision and recall. An F1 score of 1.0 indicates an ideal classifier.

Per-class Intersection over Union score (IoU): For a given class c , the IoU score is computed as the ratio between: the number of points that have been classified as class c AND are indeed of class c (intersection); and the number of points have been classified as class c OR are indeed of class c (union). An IoU score of 1.0 indicates an ideal classifier.

For all per-class metrics, we also report average values among all the classes and the weighted averages. For consistency, we report all the metrics as percentage values (the ratios between 0 and 1 are scaled linearly between 0 and 100).

5. Results and discussion

5.1. Overall performance metrics

On the three testing tiles, the model trained on all cities yields an overall accuracy of 82.8%, average F1 of 56.0%, and average IoU score of 45.3%.

5.2. Per-category performance

Table 2 reports, for each city and for each class, the performance of the model trained on data from all cities; namely, we report the per-class F1 score and IoU metrics.

We observe that “urban asset” and “vehicle” classes are harder to segment compared to other classes; this is expected due to their small size, and widely variable characteristics in terms of shape

Table 2

Per-class performance metrics for the model trained on all cities, when evaluated on one testing tile for each city (rows). The metrics are presented as percentage values.

	Terrain		Construction		Urban asset		Vegetation		Vehicle	
	IoU	F1	IoU	F1	IoU	F1	IoU	F1	IoU	F1
Davos	66.6	80.0	69.7	82.1	4.2	8.0	80.7	89.3	13.1	23.2
Zug	75.7	86.2	71.0	83.0	2.8	5.5	62.6	77.0	17.0	29.0
Zurich	48.3	65.2	81.2	89.6	5.0	9.4	58.3	73.6	24.0	38.8
Average	63.5	77.1	74.0	85.0	4.0	7.6	67.2	80.0	18.0	30.3

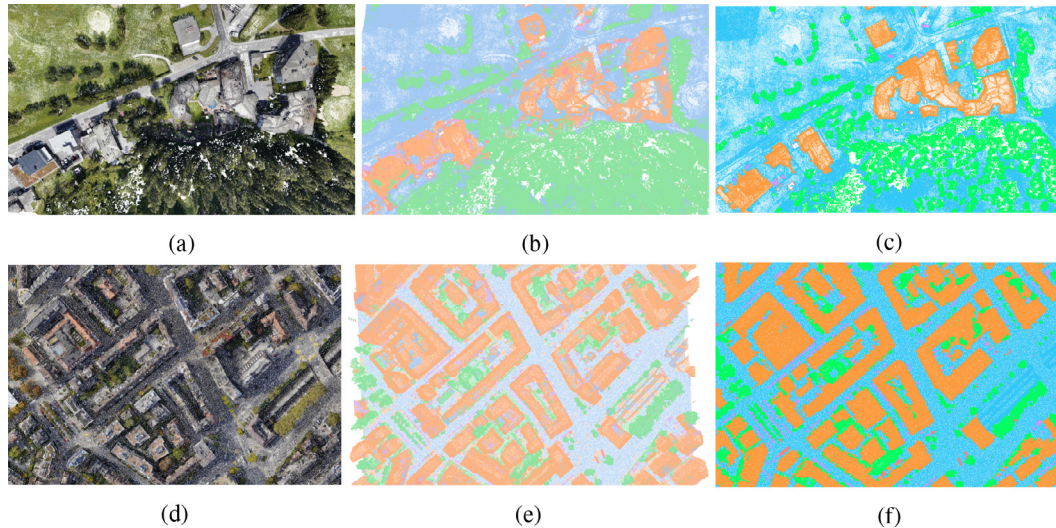


Fig. 3. (a) The RGB visualization, (b) semantic predictions with the full model, and (c) the groundtruth of a part of PC 5 from Davos; (d) the RGB visualization, (e) semantic predictions with the full model, and (f) the groundtruth of a part of PC 15 from Zurich. Legend: blue: terrain, orange: construction, green: vegetation, purple: vehicle, pink: urban asset. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

and color. IoU metrics are particularly penalized, due to the small size of each object.

Fig. 3 illustrates qualitative results of the full model for two test regions (a rural region from PC 5 in Davos and an urban region from PC 15 in Zurich). The following confusion cases are observed among categories: 1) urban asset and other categories (especially construction), 2) vegetation and terrain, and 3) vehicle and construction categories.

We hypothesize that a data pipeline that emphasizes the relative height information and favors the small categories in a stronger fashion than our current setting (e.g. a cube-based sampling rather than column-based sampling) and a stronger model than PointNet++ might help decreasing these confusion cases. As our goal is to report a baseline model on our novel dataset, we keep these model explorations for future work.

5.3. Model generalization across cities

We analyze the model generalization in a cross-city experiment setting. Fig. 4 reports performance metrics for different models, evaluated separately for each of the three cities. As seen in the first two rows and last column, the performance of the M1 and M2 models, which were trained on rural or industrial areas in Davos and Zug, decreases when they are tested on the urban Zurich test tile. Similarly, the Zurich model (M3), which is trained with the urban pointclouds, i.e. high-rise large buildings, performs worse on the rural Davos test tile (see third row, first column) than the other test tiles. This trend is observed further in the ensemble results. This point emphasizes the importance of area characteristics while learning semantics.

Fig. 5 summarizes the same data by reporting the average performance of models depending on whether they are trained on the same or different cities. For example, the performance of models




Fig. 4. The weighted accuracy (left) and the average IoU (right) for different models and model ensembles (rows) evaluated on different cities (columns). We consider: three models trained on data from each city independently; one model trained on data from all three cities; three ensembles of pairs of single-city models; one ensemble of three single-city models.




Fig. 5. The weighted accuracy (left) and average IoU (right) for the models and ensembles (bars) trained on the same or different cities.

trained on a city different than the testing city (first bar) is computed as the average of six performance values: two values predicted by the Davos model (M1) on the Zug and Zurich test tiles; two values predicted by the Zug model (M2) on the Davos and Zurich test tiles; two values predicted by the Zurich model (M3) on Davos and Zug test tiles.

Comparing the first and second bar of Fig. 5, we observe that the models trained on data from the same city have significantly better performance (average weighted accuracy 84.9%) than the models trained on data from a different city (average weighted accuracy 82.3%), despite the fact that the amount of training data is the same in different cities, and that areas used for training are always disjoint from areas used for testing.

5.4. Impact of data scale

Comparing the second and third bar of Fig. 5, we observe that the performance of the model M4 trained on data from all three cities (average weighted accuracy 87.6%) is better than the performance of the model trained just on data from the same city (average weighted accuracy 84.9%). This quantifies the impact of tripling the amount of training data, even though the additional data comes from two different cities.

5.5. Model ensembling

We consider an alternative approach to training a single model on three cities; instead, we consider the three single-city models (M1, M2, and M3), apply each model independently to each test tile, then average their predictions; in particular, for each given point in the test tile, we obtain three class probability vectors as outputs of each of the three models; we compute the element-by-element average of the three vectors, which yields a single class probability vector whose 5 elements also sum to 1. This approach is known as model ensembling [7,13] and is used frequently in machine learning.

Comparing the third and sixth bar of Fig. 5, we observe that the ensemble of the three single-city models outperforms the single model trained on the three cities. The ensembling approach is appealing, since training each model on a single-city dataset is simple and flexible: by averaging their results, we minimize the consequences of overfitting and more generally counteract model variance; on the other hand, the computational cost for inference is tripled, as three models have to be evaluated for each input.

Model ensembling experiments also allow us to quantify the performance gains from acquiring additional training data; in particular, by comparing the first and fourth bar of Fig. 5, we can observe the benefits of building an ensemble by adding a model trained on a different city; the sixth bar shows additional improvements when adding a third model to the ensemble.

6. Conclusion

This paper introduces a novel urban pointcloud dataset with pointwise semantic groundtruth. The dataset is constructed via photogrammetry on UAV-acquired high-resolution images of three Swiss cities. The dataset reports three pointcloud densities: a sparse pointcloud with RGB colors and semantic labels, a regular density pointcloud with RGB colors and semantic labels, and a dense pointcloud with only x,y,z coordinates.

The paper describes the acquisition and processing of the dataset, then illustrates several experiments on a semantic segmentation task with a prominent point-based deep learning benchmark model (PointNet++ [24]). These experiments highlight: 1) the importance of the amount of training data; 2) the advantage of using training data from the same city on which the model is evaluated; 3) the viability of simple model ensembling approaches.

As future work, we plan to compare additional recent deep-learning models for the semantic segmentation task on this dataset. Moreover, we plan to study the effects of semi-supervised and self-supervised learning methods on unstructured pointclouds.

As we make this dataset available to the research community, we hope that it will be useful for further analysis of model generalization, domain-gap studies with respect to LiDAR datasets, and various robotics applications such as traversability, and ultimately advance the state of the art in the field.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This work is supported partially by Nomoko AG and the Swiss Confederation through Innosuisse research project 31889.1 IP-ICT and through the NCCR Robotics. The authors would like to thank the collaborators in the SUPSI-ISIN institute and the colleagues in Nomoko AG for their contributions, specifically to Juan Vinuales and Mario Sanchez Gallardo for the drone flight operations; to Alexandre Ferreira do Carmo, Hugo Filipe Queiros da Cunha, Vincent Schmid, and Simon Scherer for supporting the photogrammetry steps and for manual segmentation of the pointclouds; to Sonia Batllori and Matthias Grass for supporting the pointcloud processing.

References

- [1] I. Armeni, O. Sener, A.R. Zamir, H. Jiang, I. Brilakis, M. Fischer, S. Savarese, 3D semantic parsing of large-scale indoor spaces, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2016, pp. 1534–1543.
- [2] C. Becker, E. Rosinskaya, N. Häni, E. d'Angelo, C. Strelcha, Classification of aerial photogrammetric 3D point clouds, *Photogrammetric Engineering & Remote Sensing* 84 (5) (2018) 287–295.
- [3] J. Behley, M. Garbade, A. Milioto, J. Quenzel, S. Behnke, C. Stachniss, J. Gall, SemanticKITTI: A dataset for semantic scene understanding of lidar sequences, in: Proceedings of the IEEE International Conference on Computer Vision, 2019, pp. 9297–9307.
- [4] S.A. Bello, S. Yu, C. Wang, J.M. Adam, J. Li, Deep learning on 3d point clouds, *Remote Sens (Basel)* 12 (11) (2020) 1729.
- [5] A. Boulch, J. Guerry, B. Le Saux, N. Audebert, Snapnet: 3d point cloud semantic labeling with 2D deep segmentation networks, *Computers & Graphics* 71 (2018) 189–198.
- [6] CapturingReality, Reality Capture, Drienova 3, 821 01 Bratislava, Slovakia, 2016. <https://www.capturingreality.com>.
- [7] R.T. Clemen, Combining forecasts: a review and annotated bibliography, *Int J Forecast* 5 (4) (1989) 559–583.
- [8] B.O. Community, Blender - a 3D modelling and rendering package, Blender Foundation, Stichting Blender Foundation, Amsterdam, 2018. <http://www.blender.org>.
- [9] J. Deng, W. Dong, R. Socher, L. Li, Kai Li, Li Fei-Fei, Imagenet: A large-scale hierarchical image database, in: 2009 IEEE Conference on Computer Vision and Pattern Recognition, 2009, pp. 248–255, doi:10.1109/CVPR.2009.5206848.
- [10] I. Goodfellow, Y. Bengio, A. Courville, Deep learning, volume 1, MIT press Cambridge, 2016.
- [11] Y. Guo, H. Wang, Q. Hu, H. Liu, L. Liu, M. Bennamoun, Deep learning for 3d point clouds: a survey, *IEEE Trans Pattern Anal Mach Intell* (2019).
- [12] T. Hackel, N. Savinov, L. Ladicky, J.D. Wegner, K. Schindler, M. Pollefeys, Semantic3d.net: A new large-scale point cloud classification benchmark, in: ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences, volume IV-1-W1, 2017, pp. 91–98.
- [13] S. Hashem, Optimal linear combinations of neural networks, *Neural networks* 10 (4) (1997) 599–614.
- [14] J. Heller, M. Havlena, M. Jancosek, A. Torii, T. Pajdla, 3d reconstruction from photographs by cmp sfm web service, in: 2015 14th IAPR International Conference on Machine Vision Applications (MVA), 2015, pp. 30–34, doi:10.1109/MVA.2015.7153126.
- [15] M. Hendriks, S. Meijer, J. Van Der Velden, A. Iosup, Procedural content generation for games: a survey, *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)* 9 (1) (2013) 1–22.
- [16] Q. Hu, B. Yang, S. Khalid, W. Xiao, N. Trigoni, A. Markham, Towards semantic segmentation of urban-scale 3d point clouds: a dataset, benchmarks and challenges, arXiv preprint arXiv:2009.03137 (2020).
- [17] T. Le, G. Bui, Y. Duan, A multi-view recurrent neural network for 3d mesh segmentation, *Computers & Graphics* 66 (2017) 103–112.
- [18] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, C.L. Zitnick, Microsoft coco: Common objects in context, in: European conference on computer vision, Springer, 2014, pp. 740–755.

- [19] M.E. Mohammadi, D.P. Watson, R.L. Wood, Deep learning-based damage detection from aerial SfM point clouds, *Drones* 3 (3) (2019) 68.
- [20] F. Nex, F. Remondino, Uav for 3d mapping applications: a review, *Applied geomatics* 6 (1) (2014) 1–15.
- [21] J. Niemeyer, F. Rottensteiner, U. Soergel, Contextual classification of lidar data and building object detection in urban areas, *ISPRS J. Photogramm. Remote Sens.* 87 (2014) 152–165.
- [22] F. Poux, R. Neuville, L. Van Wersch, G.-A. Nys, R. Billen, 3D point clouds in archaeology: advances in acquisition, processing and knowledge integration applied to quasi-planar objects, *Geosciences (Basel)* 7 (4) (2017) 96.
- [23] C.R. Qi, H. Su, K. Mo, L.J. Guibas, Pointnet: Deep learning on point sets for 3d classification and segmentation, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 652–660.
- [24] C.R. Qi, L. Yi, H. Su, L.J. Guibas, Pointnet++: Deep hierarchical feature learning on point sets in a metric space, in: *Advances in Neural Information Processing Systems*, 2017, pp. 5099–5108.
- [25] X. Roynard, J.-E. Deschaud, F. Goulette, Paris-lille-3d: a large and high-quality ground-truth urban point cloud dataset for automatic segmentation and classification, *Int J Rob Res* 37 (6) (2018) 545–557.
- [26] B. Rubí, R. Pérez, B. Morcego, A survey of path following control strategies for uavs focused on quadrotors, *Journal of Intelligent & Robotic Systems* (2019) 1–25.
- [27] J. Schmidhuber, Deep learning in neural networks: an overview, *Neural networks* 61 (2015) 85–117.
- [28] H. Su, V. Jampani, D. Sun, S. Maji, E. Kalogerakis, M.-H. Yang, J. Kautz, Splatnet: Sparse lattice networks for point cloud processing, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 2530–2539.
- [29] H. Su, S. Maji, E. Kalogerakis, E. Learned-Miller, Multi-view convolutional neural networks for 3d shape recognition, in: *Proceedings of the IEEE International Conference on Computer Vision*, 2015, pp. 945–953.
- [30] W. Tan, N. Qin, L. Ma, Y. Li, J. Du, G. Cai, K. Yang, J. Li, Toronto-3d: A large-scale mobile lidar dataset for semantic segmentation of urban roadways, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 2020, pp. 202–203.
- [31] N. Varney, V.K. Asari, Q. Graehling, Dales: A large-scale aerial lidar data set for semantic segmentation, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 2020, pp. 186–187.
- [32] VelodyneLidar, Velodyne Lidar: model HDL-64e, San Jose, California, United States, 2007. <https://velodynelidar.com/products/hdl-64e/>.
- [33] E. Wijnmans, Pointnet++ pytorch, 2018, (GitHub). https://github.com/erikwijnmans/Pointnet2_PyTorch.
- [34] J. Xiao, A. Owens, A. Torralba, Sun3d: A database of big spaces reconstructed using sfm and object labels, in: *Proceedings of the IEEE International Conference on Computer Vision*, 2013, pp. 1625–1632.
- [35] B. Yang, Developing a mobile mapping system for 3D GIS and smart city planning, *Sustainability* 11 (13) (2019) 3713.
- [36] J. Zhang, X. Zhao, Z. Chen, Z. Lu, A review of deep learning-based semantic segmentation for point cloud, *IEEE Access* 7 (2019) 179118–179133.
- [37] Y. Zhou, O. Tuzel, Voxelnet: End-to-end learning for point cloud based 3d object detection, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 4490–4499.
- [38] S. Zolanvari, S. Ruano, A. Rana, A. Cummins, R.E. da Silva, M. Rahbar, A. Smolic, Dublincity: annotated lidar point cloud and its applications, *arXiv preprint arXiv:1909.03613* (2019).