

Learning Digital Terrain Models From Point Clouds: ALS2DTM Dataset and Rasterization-Based GAN

Hoàng-Ân Lê , Florent Guiotte, Minh-Tan Pham , Sébastien Lefèvre , *Senior Member, IEEE*,
and Thomas Corpetti

Abstract—Despite the popularity of deep neural networks in various domains, the extraction of digital terrain models (DTMs) from airborne laser scanning (ALS) point clouds is still challenging. This might be due to the lack of the dedicated large-scale annotated dataset and the data-structure discrepancy between point clouds and DTMs. To promote data-driven DTM extraction, this article collects from open sources a large-scale dataset of ALS point clouds and corresponding DTMs with various urban, forested, and mountainous scenes. A baseline method is proposed as the first attempt to train a deep neural network to extract DTMs directly from ALS point clouds via rasterization techniques, coined DeepTerRa. Extensive studies with well-established methods are performed to benchmark the dataset and analyze the challenges in learning to extract DTM from point clouds. The experimental results show the interest of the agnostic data-driven approach, with submetric error level compared to methods designed for DTM extraction. The data and source code are available online at <https://lhoangan.github.io/deepterra/> for reproducibility and further similar research.

Index Terms—Airborne laser scanning (ALS) point cloud, dataset, deep networks, digital terrain model (DTM), generative adversarial network (GAN), rasterization.

I. INTRODUCTION

DIGITAL terrain models (DTMs), two-dimensional (2-D) representations of bare-Earth surface elevation, are critical for geographical and environmental studies, such as flood modeling [1]–[3], aboveground biomass estimation [4], infrastructure planning [5], etc. Being a subset of digital elevation models, DTMs are usually distinguished from its relative, digital surface models (DSMs), by the exclusion of objects (houses, vegetation, etc.) and representing only the bare ground surface underlying the objects.

Manuscript received June 21, 2021; revised November 29, 2021 and April 29, 2022; accepted May 29, 2022. Date of publication June 10, 2022; date of current version June 27, 2022. This work was supported by ANR SixP project under Grant ANR-19-CE02-0013. (*Corresponding author Hoàng-Ân Lê.*)

Hoàng-Ân Lê is with the Institut de Recherche en Informatique et Systèmes Aléatoires (IRISA), Université Bretagne Sud, 56000 Vannes, France, and also with the France Energies Marines, 29200 Brest, France (e-mail: hoang-an.le@irisa.fr).

Florent Guiotte, Minh-Tan Pham, and Sébastien Lefèvre are with the Institut de Recherche en Informatique et Systèmes Aléatoires (IRISA), Université Bretagne Sud, 56000 Vannes, France (e-mail: florent.guiotte@irisa.fr; minh-tan.pham@irisa.fr; sebastien.lefevre@irisa.fr).

Thomas Corpetti is with the Littoral - Environnement - Télédétection - Géomatique (LETG), UMR6554, Université Rennes 2, 35000 Rennes, France (e-mail: thomas.corpetti@univ-rennes2.fr).

Digital Object Identifier 10.1109/JSTARS.2022.3182030

DTM extraction can be achieved from satellite or aerial images [6], [7], DSMs [8], or point clouds [9], [10]. Among the possibilities, airborne laser scanning (ALS) point clouds with rich geometry information (c.f. images) and high precision due to the penetration capability through vegetation (c.f. photogrammetry) have gained attention in the recent decades. Several studies were proposed on ALS point clouds, while datasets were introduced to accommodate data-driven deep learning-based research, notably for point cloud segmentation [11]–[14].

Deep learning methods, which learn to identify salient patterns from high volume of data, have proven to be more effective over well-engineered features for several computer vision and machine learning problems [15]–[17]. For DTM extraction, different convolutional neural network (CNN)-based methods have been proposed for point cloud ground filtering [14], [18], [19], one of the two elemental steps, while the elevation interpolation problem is left open.

One hindrance for data-driven DTM extraction is the lack of dedicated large-scale datasets. While ground filtering can be formulated as a semantic segmentation problem, where one class is ground and the rest are nonground, creating ground truth DTMs require specialized knowledge and tools. To encourage deep learning-based research in this important field, we introduce a large-scale dataset of ALS point clouds and reference DTMs collected from open data sources (see Fig. 1). Several well-established methods are compared for benchmarking purposes.

Another challenge to be tackled is the discrepancy of inputs and outputs representations. 3-D ALS point clouds inherently lack topological information that 2-D image-like DTMs embrace. In this article, we propose using rasterization techniques to bridge the representation gap, which has proven to largely retain point cloud information [20]. As such, point cloud-based DTM extraction can be formulated as an image-to-image translation problem [21], and thus off-the-shelf computer vision methods can be applied. As rasterization is the idea behind rendering networks, e.g., RenderNet [22], it could be extended to an end-to-end solution.

To this end, we endeavor to predict DTM in a unified deep learning framework by concurrently tackle both the elemental problems: 1) ground filtering; and 2) elevation interpolation. The proposed baseline employs an off-the-shelf architecture and generative adversarial network (GAN) [23] to predict DTM from point cloud-based multichannel rasters. The experimental results show that simple employment of computer vision techniques approaches well-established methods specifically designed for




Fig. 1. (Top) DeepTerRa dataset with aerial laser scanning points clouds. (Bottom) Reference DTMs.

DTM extraction. Extensive analyses are performed to identify the shortcomings and potentials of our data-driven approach for DTM extraction.

The main contributions of this article are thus as follows.

- 1) The first large-scale dataset combining ALS point cloud and bare-Earth elevation (DTM), covering 52 km² with several well-established methods for benchmarking.
- 2) The first attempt, to the best of our knowledge, to formulate the DTM extraction problem in a unified deep learning framework that could potentially be developed into an end-to-end solution.
- 3) A comparative study of different rasterization strategies for DTM extraction from ALS point clouds.

II. RELATED WORK

A. DTM Extraction

DTM extraction is a long-standing problem in geospatial processing. It reconstructs from remote sensing data and the bare-Earth surface underlying objects, or land covers, both natural (trees, water, etc.) and artificial (buildings, construction, etc.). The typical process composes of two steps: 1) ground filtering; and 2) surface interpolation [8], [10]. Ground filtering, or terrain extraction, identifies the ground portions from among other entities (water, trees, buildings, etc.) and extracts them from the input geospatial data. The object-free ground portions then have their surface completed by interpolation methods, such as regularized splines [18], elevation gridding [24], or using affine-kernel embeddings [8].

Due to high topographical variation and complex structures of terrain types (e.g., flat, valley, hill, and mountainous) and objects (e.g., vegetation and construction) in geospatial data, efforts have been focused on exploring ground filtering. One strategy is to rely on terrain elevation and slope to establish rules that separate objects from ground [6], [9], [25], [26], under the assumption of gradual relief and small land covers with respect to land scales. Classification is performed by thresholding over predefined values either by chosen manually or from local adaptive filtering based on terrain slope [25].

Another approach is to approximate the terrain surface and identify objects using neighboring information. The progressive morphological filter (PMF) [27] uses sliding windows of increasing size to filter the points with elevation difference. Similarly, the simple morphological filter (SMRF) [28] filters the points with linearly increasing windows and slope thresholding. Such methods still rely on thresholds selection for the size of windows and the elevation functions. One threshold-free alternative, named the skewness balancing method (SBM) [29], is based on an unsupervised statistical analysis of point cloud skewness to filter ground points. The most recent cloth simulation filtering (CSF) [30] is based on cloth simulation to filter ground points by tuning the resolution of the virtual cloth, its rigidity, and gravity parameters. Despite being simple and computationally efficient, choosing optimal threshold remains the hindrance for this strategy.

To overcome window size sensitivity and reduce parameterization, Duan *et al.* [8] exploited DSMs through a combination of rule-based classification and multiscale morphological analysis, and defined several different thresholds based on input DSM statistics. The parameters are said to be robust to various DSM types with arbitrary quality and not require tuning. Another DSM-based DTM extraction aims to preserve sharp terrain edges by using anisotropic filtering [31].

In recent years, the paradigm is shifted towards more of learning-based methods [6], [14], [18], [19]. Such methods find their emergence from the breakthroughs of deep networks [15], [16], [32], [33], which, inspired by human neural systems, model and learn parametric nonlinear functions defined by sequentially stacked layers (*hence* deep network) from massive amount of data. The CNN has proven to be effective for various computer vision tasks, achieving state-of-the-art results from high-level tasks (e.g., classification [15], [33] and segmentation [16]) to low-level (e.g., optical flow [34] and surface normal estimation [35]).

Benefiting from the advancement of computer vision, several research studies have tried to apply deep learning techniques to DTM-related problems. Gevaert *et al.* [6] and Tapper [36] extracted DTMs from satellite imagery using fully convolutional network [16] and residual network [37] architectures. Ground

filtering has also received special attention. Ye [10] and Hu and Yuan [18] extracted ground from ALS point clouds by first sampling into image patches before performing image-based classification, while Zhang *et al.* [19] applied dynamic graph convolution techniques [38], and worked directly on point clouds to benefit from the geometric structures.

On the other hand, Ayhan *et al.* [39], Kwan *et al.* [40], and Crema *et al.* [41] assumed the availability of land cover classes, i.e., from ground filtering, and compared various image inpainting methods, among which are the GAN-based [42], in completing the DTMs. GAN [23] employ adversarial losses, usually presented as a zero-sum game between two networks: 1) generators; and 2) discriminators, in which one's gain is the other's loss, to improve quality of generated samples. The framework plays as the basis for a wide range of application, such as image-to-image translation [21], image inpainting [42], or downstream tasks, as DTM extraction from RGB images [40], [43]. Different from our approach, these methods either assume imagery inputs or only focus on a single step in DTM extraction. Our method may share similar traits to an inpainting method for GAN-based constructing DTMs, the main differences are the use of rasterization to accentuate various point cloud features and indirect use of semantic information to examine its benefit. To the best of our knowledge, there is hardly any deep learning-based method that attempts to extract DTM from ALS point clouds data directly. This could be due to the discrepancy between topologically unordered point cloud inputs and the well-organized image-like DTMs, which is one of the targets this article seeks to tackle.

B. Datasets

Another explanation for the limited deep learning-based DTM extraction from ALS point clouds is the current shortage of dedicated large-scale labeled datasets. Despite the effectiveness and popularity of deep learning techniques, large-scale datasets play a critical role in their success. Existing datasets with ALS point clouds include Filtertest¹, ISPRS2012 [44], DublinCity [11], LASDU [13], and DALES [12]. However, they are all designed for urban semantic segmentation tasks (with the exception of Filtertest, which is small reference data for ground filtering), thus not equipped with DTM data (see Table I).

ALS point cloud datasets for deep learning-based ground filtering were used in [18] and [19], yet the data are either not publicly released or limited to single scene types. The most recent OpenGF dataset [14], collected from open data sources, also targets land-level ground filtering only and does not include elevation information. Other related work includes [45] and [46], which model large-scale urban scenes from MVS Meshes.

In this work, we collect and introduce a large-scale dataset, coined DeepTerRa, including ALS point clouds (1659 M points, covering 52 km²) with semantics and reference DTM data. To assist ongoing research, part of the dataset is augmenting the DALES [12] dataset with elevation data, while the rest are from different regions. To the best of our knowledge, this is the currently first and largest dataset with ALS point clouds with

TABLE I

COMPARISON OF THE DATASET PROPOSED IN THIS ARTICLE AND EXISTING ONES REPORTED THE LITERATURE. KEY FEATURES OF OUR DATASET INCLUDE REFERENCE DTM INFORMATION FOR ALS POINT CLOUDS WITH LARGER COVERAGE AND HIGHER NUMBER OF POINTS

Dataset	year	#points	coverage	DTM	semantics
Filtertest	2003	2.3M	1.1 km ²	✓	Two classes
ISPRS2012 [44]	2012	1.2M			Nine classes
DublinCity [11]	2019	260M	2 km ²		13 classes
LASDU [13]	2020	3.1M	1 km ²		Five classes
DALES [12]	2020	505M	10 km ²		Eight classes
OpenGF [14]	2021	542.1M	47.7 km ²		Three classes
DALES (ours)	2020	505M	10 km ²	✓	Eight classes
NB (ours)	2022	1153M	42 km ²	✓	11 classes

corresponding DTM data, to encourage deep learning research in DTM extraction and remote sensing.

III. DEEPTERRA

In this section, we detail the deep learning-based approach for extracting DTM from ALS point clouds. To bridge the topological gap between unordered point cloud inputs and image-like DTM targets, we employ and elevate the rasterization technique to extract information from point clouds. The rasters then serve as inputs to a CNN-based generator, which is trained with an adversarial loss (discriminator) and outputs a DTM. The pipeline overview is shown in Fig. 3.

A. Rasterization

The ALS mechanism allows capturing high scene details by projecting laser beams on a scene and recording the returning beams. Due to the penetration ability of laser, a single emitting beam may result in several returns, or echoes, especially in vegetation areas, by interacting with different objects along the height, such as foliage, leaves, branches, ground, etc. Each echo registers a point $\mathbf{p} = \{(x, y, z, I, e, \dots)\}$ with spatial coordinates x, y , and z of the interacted surface, intensity I , echo number e , etc.

Rasterization arranges an unordered set of points $P = \{\mathbf{p}_k\}$ recorded by ALS in a regular 2-D grid or raster. The raster size is determined by the spatial extension of the points on the ground plane (xy -plane) and the size of each cell into which the points are quantized. Each grid cell, or raster pixel, defines a vertical column in the 3-D space, which can be further quantized into cubic cells or voxels. In this article, a fixed size of 25-cm cubed voxels is used.

Depending on the spatial resolution and scene structures, each pixel or voxel may contain zero, one, or several points. Different rasterization strategies apply different methods to make use of the multiple point information. We propose to derive several rasters to extract most of the information contained in the ALS data (see Fig. 2). In particular, two levels of rasterization are considered: 1) *voxel level*, treating all points in a cubic cell at a time; and 2) *pixel level*, the whole vertical column altogether.

Depending on information being rasterized, the resulted rasters can be categorized into two groups: 1) elevation raster; and 2) statistic raster.

¹[Online]. Available: <https://www.itc.nl/isprs/wgIII-3/filtertest>




Fig. 2. Visualization of different rasters used in this article for an exemplar of Dayton annotated LiDAR Earth scan (DALES) (top) and New Brunswick (NB) (bottom) with nadir-view point cloud (first column) and reference DTM (second column). Voxel-top is commonly known as DSM.

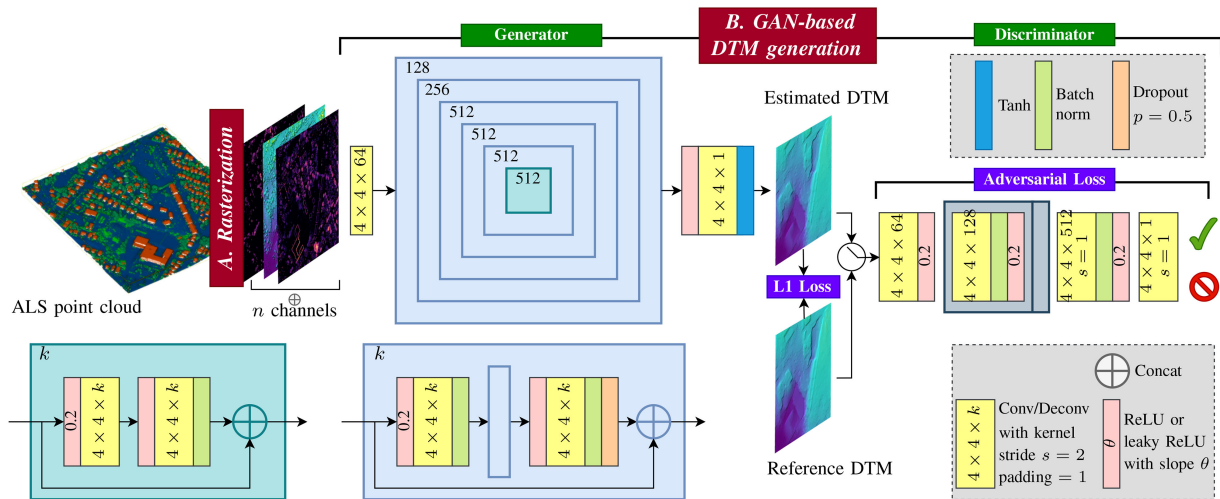


Fig. 3. Pipeline overview. An ALS point cloud is rasterized with multiple strategies, the resulted rasters are concatenated and fed to a deep neural network to generate DTMs. L1 and adversarial losses are used to train the network. Repeated blocks are color-coded. The U-Net generator is shown in a recursive manner with the recursive blocks in blue and terminating (base) block in green. Each block with a U-shape connection contains the same structure and number of layers. The only difference is the number of convolutional layers indicated by k (best viewed in color).

1) *Elevation Rasters*: The vertical coordinate z of the points are rasterized. The following three strategies are proposed.

- Pixel-mean*, which registers the mean of all points in a vertical column.
- Voxel-top* storing the mean of the upper points.
- Voxel-bottom* for the lower points.

As such, the elevation rasters contain particularly the elevation information of a point cloud, which can be used to infer DTM. In fact, the voxel-top rasters are also known as DSMs.

2) *Statistic Rasters*: The statistic rasters do not specify elevation, but the statistics of the scene structures hinting the object types and elevation. In this article, we propose the following three raster types, all at pixel-level.

- The *density* raster provides the number of points at each pixel, which is useful to assess pixel-wise reliability when processing other rasters.
- The standard deviation (*stdev*) raster describes the variation of the voxels' mean elevation in each column.
- The *echoes* raster gives the mode of the echo numbers or the most frequent echo number in a column, showing how laser beams interact with the objects encountered at each location.

As such, the raster signifies the object types, e.g., high numbers or multiple echoes, usually indicate the presence of vegetation.

As such, each raster is a single-channel image-like object representing different aspects and spatial information of the point cloud at each pixel. Hence, different from typical usage of DSM or RGB images for DTM extraction, we do not limit the usage to just one type of rasters at a time. To make the best use of various information, the generated rasters are concatenated along the depth channel before being fed to the network. Analyses on the rasters' contributions are presented in Section V.

B. GAN-Based DTM Generation

The DTM generation is formulated as an image-to-image translation problem [21]. The pairs (r, t) of generated rasters and corresponding reference DTMs are used to train a GAN. As this article does not contribute to the GAN architecture, but rather on the usage of rasterization for DTM extraction, we briefly describe the architecture, and refer interested readers to the original work for more details.

The network includes a generator \mathcal{G} , which tries to produce, from the input rasters r , a DTM $\mathcal{G}(r)$ as close as possible to




Fig. 4. DALES dataset alignment. Visualization of ground elevation difference (in meters) between *voxel-bottom* raster and provided DTM without and with offset correction; and visualization of root mean square errors (RMSEs) for different offsets during a grid search with minimum value at (3, 5). (a) Without correction. (b) Greedy grid offset search. (c) With correction.

the reference one t and a discriminator \mathcal{D} , which is trained to distinguish a generated DTM $\mathcal{G}(r)$ from the real one t . The detailed architecture of each network is shown in Fig. 3.

1) *Generator*: It is taken after the off-the-shelf U-Net architecture [47], typically for pixel-to-pixel problems. The input rasters are fed through series of several convolutional layers and nonlinearity, which reduce the spatial dimensions to learn the raster features before upsampling to the same dimensions (see Fig. 3). The generator outputs an estimated DTM of the same width and height as the input rasters.

The output DTMs are compared to the reference DTMs using an L1-loss given by

$$\mathcal{L}_G = \|t - \mathcal{G}(r)\|_1. \quad (1)$$

Euclidean losses, such as L1 or L2, are well-known for only being able to capture low-frequency information and producing blurry images [21], [48]. On the other hand, the pixel-wise comparison does not guarantee similar statistics of the generated image to the ground truth, hence might distort the ground surfaces. Thus, an adversarial loss is proposed.

2) *Discriminator*: This comprises a simple architecture adapted from [49] with conditional GAN idea from [21] and patchGAN from [50], which has proven to better capture local statistics. In the conditional GAN, the discriminator receives a DTM, generated or ground truth, together with the input rasters and learns to identify the fake from the real DTMs. The extra raster inputs allow the discriminator to condition the output based on each particular input, and hence improves its judging ability. The patchGAN loss penalizes the discriminator output at the scales of image patches, thus enforcing local statistical correctness. The adversarial loss is computed from the discriminator output as follows:

$$\mathcal{L}_D = \log \mathcal{D}(r, t) + \log(1 - \mathcal{D}(r, \mathcal{G}(t))). \quad (2)$$

The final loss is

$$\mathcal{L} = \lambda_G \mathcal{L}_G + \lambda_D \mathcal{L}_D \quad (3)$$

where λ_G and λ_D are the particular weights for each loss, and set to 100 and 1, respectively, in [21]. The optimization is alternated between \mathcal{G} and \mathcal{D} .

TABLE II

OVERVIEW OF ALS POINT CLOUDS IN DALES AND NB SUBSETS. THE POINT CLOUDS ARE ALSO ACCOMPANIED BY STANDARD LASER-SCANNING FIELDS, SUCH AS POINT-SOURCE ID, ECHO NUMBERS, SCAN ANGLE RANK, ETC

	coverage	#points	tile size	DTM res.	density	semantics
DALES	10 km ²	505M	(500 m) ²	(25 cm) ²	50.5 ppm ²	✓
NB	42 km ²	1153M	(1 km) ²	(1 m) ²	27.4 ppm ²	✓

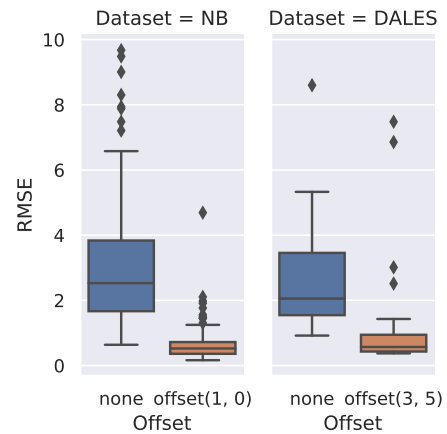


Fig. 5. Per-tile alignment evaluation. Median RMSEs of NB and DALES subset, respectively, decrease from 2.53 and 2.13 (before correction, blue boxes) to 0.52 and 0.59 (after correction, orange boxes), respectively.

IV. DATASET

To accommodate CNN training for DTM generation on ALS point clouds, we collect from open sources a large-scale dataset of ALS point clouds with reference DTM correspondences. The dataset contains two subsets: 1) DALES; and 2) NB. Table II provides an overview of the data, while general and per-tile elevation statistics of each subset are shown in Figs. 6 and 7, respectively.

1) *DALES*: The DALES dataset was presented to assist ongoing research for ALS point clouds [12]. Here, we augment it with reference DTM data, collected from the original data source of the City of Surrey².

²[Online]. Available: <https://data.surrey.ca/dataset/elevation-grid-2018>, Open Government License, City of Surrey.




Fig. 6. Elevation histogram for DALES (top) and NB (bottom) subset. The x -axis depicts elevation and y -axis the number of pixels. NB has a larger elevation range than DALES due to forested and mountainous regions while maintaining the similar shape of lower elevations in urban areas.




Fig. 7. Elevation distribution among the tiles of DALES (top) and NB (bottom) subset. The x -axis depicts the tiles and y -axis elevation. There are more variation within each tile of DALES, while NB contains higher variation among the tiles.

2) *NB*: The second subset is collected from the NB region, an open data collection, which includes LiDAR³ and corresponding DTM⁴. To vary the elevation, we sample the data tiles around the Saint John regions (urban and rural areas) and the Fundy National Park (forested and mountainous area).

The ALS point clouds are produced using the latest boresight values from the acquiring sensors (Terrascan for DALES) with preliminary quality assurance steps taken to ensure data integrity. Classification is carried out by combination of automatic and manual processes followed by manual and visual QC pass inspection. Several iterations are performed for fine-tuning and separating particular classes. The final bare Earth DTMs are produced by creating triangular irregular networks (TINs) on the ground and keypoint classes.

A quality control step is performed to enforce the alignment between the point cloud rasters and provided DTMs in each

subset. Due to possible numerical instability or different rasterizing coordinations, there are small offsets between a point cloud raster and its corresponding DTM. To quantify, the root mean square error (RMSE) is computed between a raster and provided DTM on the ground regions using ground truth semantics. Fig. 4(a) showcases an example, while Fig. 5 (blue boxes) shows the distribution of each subset (median values of 2.13 for NB and 2.53 for DALES). To correct the misalignment, a brute-force grid search is proposed. We search for both vertical and horizontal corrections in the integer range of $[0, 10]$ and pick the pairs that minimize the errors.

The error curves for DALES are shown in Fig. 4(b). The local minimums for NB and DALES are at (1, 0) and (3, 5) with median RMSE of 0.52 and 0.59, respectively. Fig. 5 shows the error statistics before and after correction. The corrected DALES example is shown in Fig. 4(c).

V. EXPERIMENTS

A. Setup

For analysis purposes, unless stated otherwise, the validation split of each subset is used. Each experiment is performed three times, and the average of the best results is reported. The RMSE (in meters) is used for evaluation, lower is better. For reporting on the test splits, the single best iteration (among the three runs) from the validation split is used to show generalizability and avoid overfitting.

The network is trained, validated, and tested separately on the DALES and NB datasets. As given in Table II, to mitigate the coverage difference of DALES and NB, the original NB point clouds are split into four quarters, of which two are chosen randomly for training, one for validation, and one for testing, resulting in 84, 42, and 42 input tiles of 500 m². The DALES subset contains 29 500-m tiles for training and 11 for testing, provided by the authors [12]. The voxel size for both sets is set to 25-cm cubed, resulting in 2000-pixel rasters, which are downsampled 1:4 before inputting to the network.

In all the experiments, unless stated otherwise, the hyperparameters are set following the default set up from [21], i.e., $\lambda_{\mathcal{C}} = 100$ and $\lambda_{\mathcal{D}} = 1$, Adam optimizer with momentum parameters $\beta_1 = 0.5$ and $\beta_2 = 0.999$, and learning rate of 0.0002. The optimization is alternated between \mathcal{G} and \mathcal{D} . The networks are trained on a cluster node with 8-core Intel Xeon E5-2620 @ 2.10 GHz, 64 GB RAM, and RTX-2080 Ti/11 GB VRAM.

B. Elevation Rasters as DTM

The pixel-mean, voxel-top, and voxel-bottom rasters particularly provide elevation information; hence can be used directly as prediction for DTMs. Knowing their performance gives more insight into the rasterization strategies and the network processing them. In this experiment, we show their performance on both validation and test splits. The results are given in Table III.

From the definitions, it is trivial that voxel-bottom rasters are the closest to DTMs, while voxel-top produces the highest errors and pixel-mean stays in between. The rasters errors also provide insights of the object types and heights in each set, hinting

³[Online]. Available: <http://geonb.snb.ca/li/>, Open Government License, NB.

⁴[Online]. Available: <http://geonb.snb.ca/nbdem/>, Open Government License, NB.

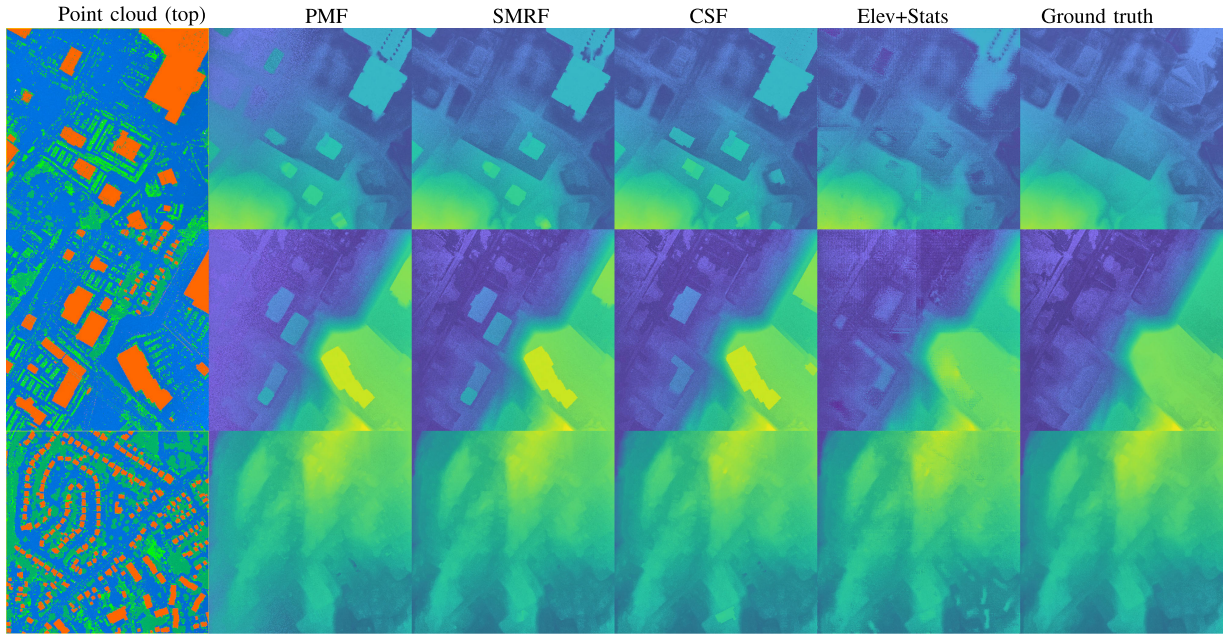


Fig. 8. RMSE (in meters) on the test split. The deep learning-based method with elevation and statistics rasters (Elev+Stats) removes most of the objects and shows better performance on large-scale buildings. The point clouds are shown with semantic labels: ground (blue), buildings (red), and vegetation (green).

TABLE III

RMSE (IN METERS) FOR USING ELEVATION RASTERS DIRECTLY AS PREDICTED DTM (NO OTHER OPERATIONS ARE PERFORMED) ON VALIDATION AND TEST SPLIT

	DALES		NB	
	val	test	val	test
pixel-mean	4.40	4.90	4.92	4.75
voxel-top	5.80	6.31	7.00	6.76
voxel-bottom	3.16	3.68	2.03	2.01

the object semantics. As laser beams can penetrate well in vegetation areas, but are blocked by solid objects, e.g., buildings, voxel-bottom errors depict the presence of buildings and large constructions. The results show that there are more and/or higher buildings in the DALES subset than in NB. On the other hand, voxel-top rasters pick the highest object points, thus show that NB might comprise more, and/or higher, trees. They agree with the fact that the DALES subset is focused on urban areas and NB contains more forested regions. Such information may be used to quantitatively identify scene category without the availability of semantic labels.

C. Combination of Rasters

In this experiment, the elevation rasters (pixel-mean, voxel-top, and voxel-bottom) are employed for DTM extraction. Several strategies are possible, including using each elevation raster individually, in combination, together with one or all statistic rasters (density, stdev, and echoes). The results for different combination strategies are given in Table IV.

In general, adding all three statistics rasters improves the network performance by a large margin. Hence, this could be the go-to solution. Regarding the elevation rasters to be used, the current results show the dependency to the type of scenes, as NB

TABLE IV

RMSE (IN METERS) OF PREDICTED DTMs FROM RASTERS COMBINATION ON THE DALES AND NB VALIDATION SETS. THE ELEVATION RASTERS (ROWS, LAST FOR ALL) ARE CONCATENATED WITH STATISTIC RASTERS (COLUMNS, FIRST FOR NONE AND LAST FOR ALL). ADDING STATISTIC RASTERS IMPROVES THE NETWORK PERFORMANCE

	-	+density	+echoes	+stdev	+3stat
DALES					
pixel-mean	1.75	1.60	1.54	1.42	1.24
voxel-top	1.64	1.61	1.62	1.54	1.12
voxel-bottom	1.72	1.60	1.56	1.63	1.27
3elev	1.39	1.31	1.27	1.45	1.23
NB					
pixel-mean	2.30	2.43	2.01	1.83	1.46
voxel-top	2.70	2.82	2.18	1.85	1.29
voxel-bottom	1.97	1.81	1.55	1.85	1.22
3elev	2.13	2.20	1.86	2.06	1.71

contains a dominantly large number of vegetation (mountainous and forested regions), the voxel-bottom rasters provide a close representation to DTMs while it is closer to voxel-top rasters for large constructions in urban areas from the DALES subset. Thus, the combination of all three elevation rasters dominates in DALES, while voxel-bottom dominates in NB. The performance gap is diminished when adding statistics information. For later experiments, only voxel-top and voxel-bottom are used with statistics rasters.

D. Loss Coefficients

In this experiment, we show how the losses contribute to the network performance. We vary the GAN loss coefficient λ_D with respect to two values of the L1 loss coefficient λ_G . As given in Table VII, although the two coefficients are not completely linearly correlated, the network performs poorly when there is no GAN loss contribution $\lambda_D = 0$, showing its necessity.

TABLE V

RMSE (IN METERS) OF PREDICTED DTMs FROM RASTERS AND SEMANTIC COMBINATION ON THE DALES AND NB VALIDATION SETS. THE ELEVATIONS (ROWS) ARE CONCATENATED WITH ONE-CHANNEL (SEM1) OR TWO-CHANNEL (SEM2) BINARY. VOXEL-BOTTOM COMBINED WITH TWO-CHANNEL SEMANTIC AND STATISTIC RASTERS SHOWS SUPERIOR PERFORMANCE

	-	+sem1	+sem2	+3stat		
				-	+sem1	+sem2
DALES						
voxel-top	1.64	1.59	1.20	1.12	1.06	0.88
voxel-bottom	1.72	1.17	0.88	1.27	0.97	0.91
NB						
voxel-top	2.70	2.03	2.02	1.29	1.25	1.22
voxel-bottom	1.97	1.76	1.19	1.22	1.19	1.06

E. Semantics

DTM extraction methods generally rely on ground filtering labels. In this experiment, we explore the effects of semantic information to the pipeline. To maintain consistency with the previous setups, we concatenate semantic maps to the rasters as extra channels and leave the exhaustive study of efficient usage of semantics for future work. This is a less direct use of semantic information compared to the more popular way in which objects' regions are removed and inpainted [39]–[41]. The semantic maps are generated by rasterizing the mode of point labels in each vertical column, i.e., each pixel receives the most frequent label in the vertical column.

The label raster is then binarized as ground/nonground and converted to one-hot vector format. Although a single-channel binary image (sem1) is, in principle, sufficient for describing binary classes, we also include the complementary channel, making two-channel semantic maps (sem2), for the sake of comparison. The results are given in Table V.

Adding semantic labels, in general, improves the network performance. Surprisingly, despite the redundant information, using two-channel complementary binary labels (+sem2) show better performance than when using one-channel labels, hinting how semantic maps should be handled for DTM extraction. This could be explained that the complementary information might not be trivial to be derived with the current network capacity, and show potential improvement when soft-labeling vectors are used, i.e., when classes are represented by floating point probabilities instead of binary numbers.

F. Comparison With State of the Art

To evaluate the data-driven approach, we compare and show the *test* results of our best validated iteration and those of the well-established ground-filtering methods designed for DTM extraction. Specifically, four baselines are used with default parameters: 1) PMF [27] (1-m cell, exponential window of maximum 33 m, distance of 15–250 cm, and slope of 45°); 2) SMRF [28] (1-m cell, max window of 18 m, elevation scalar of 1.25 m, threshold of 0.5 m, and 15% slope tolerance); 3) SBM [29]; and 4) CSF [27] (cloth resolution of 0.5 m, rigidity of 3, time step of 0.65, 500 iterations, and slope smoothing post-processing). The implementations of PMF, SMRF, and SBM are

TABLE VI

RMSE (IN METERS) OF THE PROPOSED DEEP LEARNING-BASED METHOD IN COMPARISON WITH EXISTING RULE-BASED APPROACHES ON DALES AND NB TEST SETS

	DALES	NB
Rule-based		
PMF [27]	4.27	0.81
SBM [29]	3.59	4.15
SMRF [28]	4.08	0.75
CSF [30]	1.19	0.88
Deep learning-based		
Elevation rasters	1.98	2.05
Elevation+3stat	1.27	1.31
Elevation+3stat+sem2	0.82	0.98

TABLE VII

RMSE (IN METERS) FOR VARIOUS GAN COEFFICIENTS λ_D WITH RESPECT TO TWO VALUES OF L1 COEFFICIENTS λ_G ON THE DALES VALIDATION SET. THE NETWORK PERFORMS BEST AT $(\lambda_G, \lambda_D) = (100, 1)$. THE NETWORKS PERFORM POORLY WHEN THERE IS NO GAN LOSS CONTRIBUTION, I.E., $\lambda_D = 0$

λ_D	$\lambda_G = 100$
0	2.98
0.01	2.68
0.1	2.57
1	1.72
10	1.93
100	2.48

from the point data abstraction library⁵, while CSF is provided by the original authors⁶. A TIN is computed using Delaunay triangulation on the extracted ground points before being rasterized to DTMs. The results of rule-based methods provided in Table VI are also meant for benchmarking our dataset.

As a parameter-free method, SBM relies on strong assumptions of flat ground and sacrifices its performance. PMF, SMRF, and CSF perform similarly well on the NB subset with scattered small houses and forested regions. While PMF and SMRF suffer for large buildings and constructions on the DALES subset, CSF shows more robust results even with fixed default parameters.

For comparison with deep learning-based methods, we use the best iteration on the validation set from previous experiments. The deep learning-based methods show more stable similar results on both DALES and NB subsets. Due to independence from fixed parameters, deep learning-based methods rely only on training set variations, thus can, in principle, be improved when more data are collected.

The image-to-image approach, however, is limited due to the loss of point cloud geometry of which the rule-based methods are designed to take advantage. Yet as more information is provided by the statistical rasters, the agnostic off-the-shelf network's performance approaches that of the best rule-based designed directly for ground filtering.

Qualitative results in Fig. 8 show that a network with inputs of elevation and statistical rasters are more robust to large objects. Most of the houses and large buildings that left untouched by the rule-based methods on the first-two rows are removed by

⁵[Online]. Available: <https://pdal.io/>

⁶[Online]. Available: <https://github.com/jianboqi/CSF>

the network. As the rule-based methods depend greatly of fixed parameters, they might be effective with objects of certain sizes (small objects on the third row), but failed in other cases; it is well-known that the parameters need to be carefully tuned for separate cases, showing needs for a deep learning-based and end-to-end solution.

VI. CONCLUSION

This article attempted to unify the DTM extraction process from ALS point clouds using a deep learning-based method. By formulating the problem as an image-to-image translation and using rasterization techniques, off-the-shelf computer vision-designed architectures can be used to tackle this important remote sensing task. To stimulate data-driven approaches on this problem, a large-scale dataset of ALS point clouds and reference DTMs was also introduced and evaluated.

The proposed method, despite being domain-agnostic, showed comparable results to task-dedicated methods and performed on par with state-of-the-art baselines. Experiments showed that the various point cloud-specific qualities can be gathered by simple rasterization techniques, which have proven to be helpful for predicting bare-ground elevation.

Among all types of information to be rasterized, semantic played an important role in extracting DTM due to its ability in identifying objects. Although, in this article, only one-hot vectors ground truth labels were used, the same idea could be applied to softmax smooth predicted labels. Other effective ways of exploiting semantic information for DTM extraction are encouraged for future research.

Effective may it be, rasterization, nonetheless, squashes geometry-rich 3-D point clouds into 2-D representations, hence can only depict a certain aspects of a point cloud. As such, this aggressive operation, together with the 2-D image-to-image translation approach, does not fully exploit the potential of 3-D data. The submetric performance gap with other ground filters, which rely on point cloud geometry, showed the room for development of the deep learning-based method for DTM extraction in future research.

REFERENCES

- [1] K. Bakula, "Reduction of DTM obtained from LiDAR data for flood modeling," *Arch. Photogrammetry*, vol. 22, pp. 51–61, 2011.
- [2] N. A. Muhadi, A. F. Abdullah, S. K. Bejo, M. R. Mahadi, and A. Mijic, "The use of LiDAR-derived DEM in flood applications: A review," *Remote Sens.*, vol. 12, no. 14, 2020, Art. no. 2308.
- [3] M. Muthusamy, M. R. Casado, D. Butler, and P. Leinster, "Understanding the effects of digital elevation model resolution in urban fluvial flood modelling," *J. Hydrol.*, vol. 596, 2021, Art. no. 126088.
- [4] V. González-Jaramillo, A. Fries, and J. Bendix, "AGB estimation in a tropical mountain forest (TMF) by means of RGB and multispectral images using an unmanned aerial vehicle (UAV)," *Remote Sens.*, vol. 11, no. 12, 2019, Art. no. 1413.
- [5] Z. Xu *et al.*, "Road extraction in mountainous regions from high-resolution images based on DSDNet and terrain optimization," *Remote Sens.*, vol. 13, no. 1, 2021, Art. no. 90.
- [6] C. M. Gevaert, C. Persello, F. Nex, and G. Vosselman, "A deep learning approach to DTM extraction from imagery using rule-based training labels," *ISPRS J. Photogrammetry Remote Sens.*, vol. 142, pp. 106–123, 2018.
- [7] C. A. Rokhmana and A. R. Sastra, "DTM generation from aerial photo and Worldview-3 images by using different DSM filtering methods," in *Proc. 6th Geoinf. Sci. Symp.*, 2019, pp. 152–158.
- [8] L. Duan, M. Desbrun, A. Giraud, F. Trastour, and L. Laurere, "Large-scale DTM generation from satellite data," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. Workshop*, 2019, pp. 1442–1450.
- [9] M. Debella-Gilo, "Bare-earth extraction and DTM generation from photogrammetric point clouds including the use of an existing lower-resolution DTM," *Int. J. Remote Sens.*, vol. 37, no. 13, pp. 3104–3124, 2016.
- [10] W. Ye, "Extraction of digital terrain models from airborne laser scanning data based on transfer-learning," Master's thesis, Univ. Waterloo, Waterloo, ON, Canada, 2019.
- [11] I. Zolanvari *et al.*, "DublinCity: Annotated LiDAR point cloud and its applications," in *Proc. Brit. Mach. Vis. Conf.*, 2019, pp. 127.1–127.13.
- [12] N. Varney, V. K. Asari, and Q. Graehling, "DALES: A large-scale aerial LiDAR data set for semantic segmentation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. Workshops*, 2020, pp. 717–726.
- [13] Z. Ye *et al.*, "LASDU: A large-scale aerial LiDAR dataset for semantic labeling in dense urban areas," *ISPRS Int. J. Geo Inf.*, vol. 9, no. 7, 2020, Art. no. 450.
- [14] N. Qin, W. Tan, L. Ma, D. Zhang, and J. Li, "OpenGF: An ultra-large-scale ground filtering dataset built upon open ALS point clouds around the world," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. Workshops*, 2021, pp. 1082–1091.
- [15] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," in *Proc. Conf. Neural Inf. Process. Syst.*, 2012, pp. 1097–1105.
- [16] J. Long, E. Shelhamer, and T. Darrell, "Fully convolutional networks for semantic segmentation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2015, pp. 3431–3440.
- [17] R. Girshick, J. Donahue, T. Darrell, and J. Malik, "Rich feature hierarchies for accurate object detection and semantic segmentation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2014, pp. 580–587.
- [18] X. Hu and Y. Yuan, "Deep-learning-Based classification for DTM extraction from ALS point cloud," *Remote Sens.*, vol. 8, no. 9, 2016, Art. no. 730.
- [19] J. Zhang, X. Hu, H. Dai, and S. Qu, "DEM extraction from ALS point clouds in forest areas via graph convolution network," *Remote Sens.*, vol. 12, no. 1, 2020, Art. no. 178.
- [20] F. Guiotte, M.-T. Pham, R. Dambreville, T. Corpetti, and S. Lefèvre, "Semantic segmentation of LiDAR points clouds: Rasterization beyond digital elevation models," *IEEE Geosci. Remote Sens. Lett.*, vol. 17, no. 11, pp. 2016–2019, Nov. 2020.
- [21] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, "Image-to-image translation with conditional adversarial networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 5967–5976.
- [22] T. Nguyen-Phuoc, C. Li, S. Balaban, and Y.-L. Yang, "RenderNet: A deep convolutional network for differentiable rendering from 3D shapes," in *Proc. Conf. Neural Inf. Process. Syst.*, 2018. [Online]. Available: <https://proceedings.neurips.cc/paper/2018/hash/68d3743587f71fbaa5062152985aff40-Abstract.html>
- [23] I. Goodfellow *et al.*, "Generative adversarial nets," in *Proc. Neural Inf. Process. Syst.*, 2014, pp. 2672–2680.
- [24] M. F. Hutchinson, T. Xu, and J. A. Stein, "Recent progress in the ANU-DEM elevation gridding procedure," in *Proc. Geomorphometry*, 2011, pp. 19–22.
- [25] G. Sithole, "Filtering of laser altimetry data using a slope adaptive filter," *Int. Arch. Photogrammetry Remote Sens.*, vol. 34, no. 3, pp. 203–210, 2001.
- [26] G. Sithole and G. Vosselman, "Experimental comparison of filter algorithms for bare-earth extraction from airborne laser scanning point clouds," *ISPRS J. Photogrammetry Remote Sens.*, vol. 59, no. 1, pp. 85–101, 2004.
- [27] K. Zhang, S.-C. Chen, D. Whitman, M.-L. Shyu, J. Yan, and C. Zhang, "A progressive morphological filter for removing nonground measurements from airborne LIDAR data," *IEEE Trans. Geosci. Remote Sens.*, vol. 41, no. 4, pp. 872–882, Apr. 2003.
- [28] T. J. Pingel, K. C. Clarke, and W. A. McBride, "An improved simple morphological filter for the terrain classification of airborne LIDAR data," *ISPRS J. Photogrammetry Remote Sens.*, vol. 77, pp. 21–30, 2013.
- [29] M. Bartels and H. Wei, "Threshold-free object and ground point separation in LIDAR data," *Pattern Recognit. Lett.*, vol. 31, no. 10, pp. 1089–1099, 2010.
- [30] W. Zhang *et al.*, "An easy-to-use airborne LiDAR data filtering method based on cloth simulation," *Remote Sens.*, vol. 8, no. 6, 2016, Art. no. 501.
- [31] A. Pijl, J.-S. Bailly, D. Feurer, M. A. El Maoui, M. R. Boussema, and P. Tarolli, "TERRA: Terrain extraction from elevation rasters through repetitive anisotropic filtering," *Int. J. Appl. Earth Observ. Geoinf.*, vol. 84, 2020, Art. no. 101977.
- [32] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "ImageNet: A large-scale hierarchical image database," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2009, pp. 248–255.

- [33] K. He, X. Zhang, S. Ren, and J. Sun, "Delving deep into rectifiers: Surpassing human-level performance on ImageNet classification," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2015, pp. 1026–1034.
- [34] D. Sun, X. Yang, M.-Y. Liu, and J. Kautz, "PWC-Net: CNNs for optical flow using pyramid, warping, and cost volume," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 8934–8943.
- [35] H.-A. Le, A. S. Baslamisli, T. Mensink, and T. Gevers, "Three for one and one for three: Flow, segmentation, and surface normals," in *Proc. Brit. Mach. Vis. Conf.*, 2018. [Online]. Available: <http://bmvc2018.org/contents/papers/0201.pdf>
- [36] G. Tapper, "Extraction of DTM from satellite images using neural networks," Master's thesis, Linköping Univ., Linköping, Sweden, 2016.
- [37] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 770–778.
- [38] Y. Wang, Y. Sun, Z. Liu, S. E. Sarma, M. M. Bronstein, and J. M. Solomon, "Dynamic graph CNN for learning on point clouds," *ACM Trans. Graph.*, vol. 38, no. 5, pp. 1–12, Oct. 2019.
- [39] B. Ayhan, C. Kwan, J. Larkin, L. Kwan, D. Skarlatos, and M. Vlachos, "Performance comparison of different inpainting algorithms for accurate DTM generation," *Proc. Geospatial Inform. X*, vol. 11398, pp. 144–159, 2020.
- [40] C. Kwan, D. Gribben, B. Ayhan, and J. Larkin, "Practical digital terrain model extraction using image inpainting techniques," in *Recent Advances in Image Restoration with Applications to Real World Problems*, C. Kwan, Ed. Rijeka, Croatia: IntechOpen, 2020, ch. 8.
- [41] S. Crema *et al.*, "Can inpainting improve digital terrain analysis? comparing techniques for void filling, surface reconstruction and geomorphometric analyses," *Earth Surf. Processes Landforms*, vol. 45, no. 3, pp. 736–755, Mar. 2020.
- [42] J. Yu, Z. Lin, J. Yang, X. Shen, X. Lu, and T. S. Huang, "Generative image inpainting with contextual attention," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 5505–5514.
- [43] E. Panagiotou, G. Chochlakis, L. Grammatikopoulos, and E. Charou, "Generating elevation surface from a single RGB remotely sensed image using deep learning," *Remote Sens.*, vol. 12, no. 12, Jun. 2020, Art. no. 2002.
- [44] F. Rottensteiner *et al.*, "The ISPRS benchmark on urban object classification and 3D building reconstruction," in *Proc. ISPRS Ann. Photogrammetry Remote Sens. Spatial Inf. Sci.*, 2012, pp. 293–298.
- [45] L. Zhu, S. Shen, X. Gao, and Z. Hu, "Large scale urban scene modeling from MVS meshes," in *Proc. Eur. Conf. Comput. Vis.*, 2018, pp. 614–629.
- [46] J. Han *et al.*, "Urban scene LOD vectorized modeling from photogrammetry meshes," *IEEE Trans. Image Process.*, vol. 30, pp. 7458–7471, 2021.
- [47] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *Proc. Med. Image Comput. Comput.- Assist. Interv.*, 2015, pp. 234–241.
- [48] A. B. L. Larsen, S. K. Sønderby, H. Larochelle, and O. Winther, "Autoencoding beyond pixels using a learned similarity metric," in *Proc. 33rd Int. Conf. Mach. Learn.*, 2016, pp. 1558–1566.
- [49] A. Radford, L. Metz, and S. Chintala, "Unsupervised representation learning with deep convolutional generative adversarial networks," in *Proc. 4th Int. Conf. Learn. Representations*, 2016. [Online]. Available: <https://iclr.cc/archive/www/doku.php%3Fid=iclr2016:main.html>
- [50] C. Li and M. Wand, "Precomputed real-time texture synthesis with Markovian generative adversarial networks," in *Proc. Eur. Conf. Comput. Vis.*, 2016, pp. 702–716.



Hoàng-Ân Lê received the B.Sc. degree in computer science from the University of Science, Vietnam National University, Ho Chi Minh City, Vietnam, in 2012, the M.Sc. degree in engineering from EURECOM, Télécom ParisTech, Paris, France, in 2015, and the Ph.D. degree in computer vision and machine learning from the University of Amsterdam, Amsterdam, The Netherlands, in 2021, while working on the TrimBot2020 project within the EUHorizon2020 programme.

Since then, he has been a Postdoctoral Researcher with OBELIX Team, Institut de Recherche en Informatique et Systèmes Aléatoires, University of South Brittany, Vannes, France. His research interests include computer vision and machine learning, with a focus on multimodality-multitask learning, 3-D understanding, and applications in Earth observation and remote sensing.



Florent Guiotte received the M.Eng. degree in computer science from University Rennes 1, Rennes, France, in 2016, and the Ph.D. degree in computer science from University Rennes 2, Rennes, France, in 2021.

In 2021, he was a Postdoctoral Fellow with OBELIX Team, Institut de Recherche en Informatique et Systèmes Aléatoires. Since 2022, he has been with L'Avion Jaune, Montpellier, France. His research interests include computer vision and machine learning applied to remote sensing images and


point clouds.



Minh-Tan Pham received the M.Eng. and Ph.D. degrees in information and signal processing from Telecom Bretagne, Brest, France, in 2013 and 2016, respectively.

In 2013, he was an Intern with the Department of Geomatics, Laval University, Canada, and with French Space Agency, France, in 2015. From 2016 to 2019, he was a Postdoctoral Fellow with Institut de Recherche en Informatique et Systèmes Aléatoires (IRISA) Laboratory. Since 2019, he has been an Assistant Professor with University Bretagne Sud, IUT

de Vannes, and a Researcher with OBELIX Team, IRISA. His research interests include image processing and machine learning applied to remote sensing data, with a focus on mathematical morphology, hierarchical representation and self-supervised deep networks for feature extraction, object detection, and classification tasks.



Sébastien Lefèvre (Senior Member, IEEE) received the M.Sc. degree in computer engineering from the University of Technology of Compiègne, France, in 1999, the Ph.D. degree in computer science from the University of Tours, France, in 2002, and the Habilitation degree in computer science from the University of Strasbourg, France, in 2009.

Since 2010, he has been a Full Professor of Computer Science with Université Bretagne Sud, Vannes, France.

Prof. Lefèvre is involved in numerous activities related to artificial intelligence for Earth and environment observation. He is the Founder of the OBELIX Team, Institute for Research in Computer Science and Random Systems, the Chair of the GeoData Science track of the EMJMD Copernicus Master in Digital Earth, and a Co-Chair of the ECML-PKDD MACLEAN Workshop series on machine learning for earth observation. He is an Associate Editor for the IEEE TRANSACTIONS ON GEOSCIENCE AND REMOTE SENSING.



Thomas Corpetti received the Engineering degree in electrical engineering from INSA Rennes (National Institute of Applied Sciences) and the master's degree in computer vision from University Rennes I, Rennes, France, in 1999. He then got Ph.D. and Habilitation degrees in computer vision and applied mathematics from University Rennes 1, in 2002 and 2011, respectively.

In 2004, he became a Permanent Researcher with the French National Institute for Scientific Research (CNRS), Paris, France, on the analysis of remote sensing image sequences for environmental applications. From 2009 to 2012, he was with LIAMA, Beijing, China, a Sino-French laboratory in computer sciences, automatics, and applied mathematics. He is currently with the Observatory for Universe Sciences of Rennes, France, and Littoral - Environnement - Télédétection - Géomatique, UMR6554, CNRS, as the Director of Research. His research focuses on the definition of computer vision tools for the analysis of remote sensing data (low and high resolution) for environmental applications.