

EFFICIENT LARGE-SCALE AIRBORNE LIDAR DATA CLASSIFICATION VIA FULLY CONVOLUTIONAL NETWORK

E. Maset^{a,*}, B. Padova^b, and A. Fusiello^a

^a DPIA, University of Udine, Via delle Scienze, 206 – Udine, Italy – (eleonora.maset, andrea.fusiello)@uniud.it

^b Helica s.r.l., Via fratelli Solari, 10 – Amaro, Italy – b.padova@helica.it

Commission III, WG III/5

KEY WORDS: LiDAR, Classification, Large-scale dataset, Deep Learning, Fully Convolutional Network

ABSTRACT:

Nowadays, we are witnessing an increasing availability of large-scale airborne LiDAR (Light Detection and Ranging) data, that greatly improve our knowledge of urban areas and natural environment. In order to extract useful information from these massive point clouds, appropriate data processing is required, including point cloud classification. In this paper we present a deep learning method to efficiently perform the classification of large-scale LiDAR data, ensuring a good trade-off between speed and accuracy. The algorithm employs the projection of the point cloud into a two-dimensional image, where every pixel stores height, intensity, and echo information of the point falling in the pixel. The image is then segmented by a Fully Convolutional Network (FCN), assigning a label to each pixel and, consequently, to the corresponding point. In particular, the proposed approach is applied to process a dataset of 7700 km² that covers the entire Friuli Venezia Giulia region (Italy), allowing to distinguish among five classes (*ground, vegetation, roof, overground* and *power line*), with an overall accuracy of 92.9%.

1. INTRODUCTION

In the last years, governments and other institutions worldwide have been promoting the survey of large areas of the national territory to be employed, e.g., for natural hazard management, urban planning and facilities monitoring. In this context, airborne LiDAR (Light Detection and Ranging) technology represents a suitable survey platform to obtain high resolution data at wide scale, requiring, however, efficient algorithms to handle and process the large amount of acquired data.

In the LiDAR data processing pipeline, classification is one of the most important and time consuming stage, necessary for the subsequent generation of cartographic products. Ground points must be extracted, e.g., to create digital terrain models (DTMs), whereas the identification of vegetation is essential to evaluate its density or, in the field of power lines monitoring, to automatically calculate the distance from the conductors, just to name a few applications.

When dealing with airborne LiDAR (ALS) large-scale datasets, the processing time becomes an essential factor to take into account. In this work, we show how an algorithm based on Convolutional Neural Networks (CNNs) was profitably employed to classify a dataset of 7700 km², that covers the entire Friuli Venezia Giulia region (Italy). The proposed approach is applied to distinguish among five classes, namely: *ground, vegetation, roof, overground* (e.g., cars, walls and chimneys) and *power line*, achieving an overall accuracy of 92.9% with a classification time of 11 minutes per km². It is worth noting that half of the surveyed region is characterized by alpine areas (higher than 600 m a.s.l.), particularly challenging for the classification task, as observed in (Winiwarter et al., 2019). However, our method achieved a good classification also in mountainous environments.

* Corresponding author

The paper is organized as follows. Section 2 reports a review on the existing methods for point cloud classification, while Sec. 3 describes in detail the proposed approach. In Sec. 4 the dataset is presented, together with the achieved results. Finally, Sec. 5 draws the conclusion.

2. STATE OF THE ART

Point cloud classification has always been a hot research topic in the LiDAR data processing field, with several applications in, e.g., land cover classification, vegetation studies in forestry and agriculture, and road infrastructure management (Wang et al., 2020). Classification methods usually rely both on geometric information (i.e., the 3D coordinates of the surveyed points and their distribution in a neighboring region), as well as on the intensity of the backscattered pulse (Scaioni et al., 2018). Furthermore, thanks to the recent availability of laser scanners that are able to digitize the entire waveform of the reflected signal, several algorithms exploit also full-waveform data and the features derived from them (Maset et al., 2015).

Early works mainly proposed classification algorithms based on predefined discriminant rules and simple thresholds (Rutzinger et al., 2008, Wagner et al., 2008), subsequently replaced by machine learning techniques, such as Support Vector Machine (Serna, Marcotegui, 2014) and Random Forest (Tran et al., 2018). The main limitation of these approaches lies in the need of hand-crafted features, which can be sensible to changes in the data characteristics. Moreover, these approaches usually classify each point independently, without considering the labels assigned to neighboring points (Wang et al., 2020).

In the last decade, deep learning techniques have been spreading in disciplines such as computer vision, robotics and audio processing, replacing methods based on hand-engineered features with algorithms that learn both features and classifier

end-to-end (Goodfellow et al., 2016). In particular, Convolutional Neural Networks (CNNs) and Fully Convolutional Networks (FCNs) proved to be successful tools for image classification and segmentation tasks, respectively (Szegedy et al., 2015, Garcia-Garcia et al., 2017).

Very recently, various methods based on deep learning have been applied in the remote sensing field, also for point cloud classification (Griffiths, Boehm, 2019). They can be distinguished into three main approaches: (i) classification of single points based on 2D CNNs, (ii) simultaneous classification of portions of point clouds via FCNs that operate on a 2D image, and (iii) exploitation of network architectures that allow to operate directly in the 3D space. The methods proposed in (Yang et al., 2017, Zhao et al., 2018) fall in the first category, with the 3D neighborhood features of a point that are transformed into a 2D image that is then classified by a CNN.

Leveraging on the high performances that can be obtained by the networks usually applied for image processing, (Zorzi et al., 2019) proposed to map the point cloud and the information derived from full-waveform data into an image segmented by a FCN, assigning in this way a label to each pixel and, consequently, to the point falling in the pixel. A similar approach was applied also in (Rizaldy et al., 2018): in this case, the distinction among the three classes *ground*, *vegetation* and *building* is achieved with an overall accuracy of 93%.

Deep learning architectures that operate directly on 3D data have been proposed, e.g., in (Wu et al., 2015, Tchapmi et al., 2017). The cited methods rely on a voxelization of the point cloud, i.e., data are represented by means of a 3D regular voxel grid, subsequently fed to a network that performs convolution in the 3D space. This kind of approach is usually very expensive from the computational point of view, limiting the size of the point cloud that can be taken as input by the network. PointNet (Qi et al., 2016) and its improved version PointNet++ (Qi et al., 2017) have been the first architectures able to operate on unstructured data. These networks do not require the point cloud transformation on a regular grid neither use convolution functions, but employ instead multi-layer perceptrons to extract features, both at local and at global scale. PointNet and PointNet++ proved to outperform state-of-the-art methods for common benchmark datasets such as ModelNet40 (Wu et al., 2015) and were successfully applied, e.g., to distinguish between coniferous and deciduous tree points (Briechle et al., 2019). Another point-based deep learning method that does not involve rasterization or voxelization is the one proposed by (Landrieu, Simonovsky, 2018), that pre-organizes the point cloud in the so-called *Superpoint Graph* and exploits a graph convolutional network to perform the classification task.

Methods that operate directly in the 3D space are particularly suited for point clouds representing indoor scenes and road environments, acquired by, e.g., Terrestrial Laser Scanner (TLS) and Mobile Mapping System (MMS). In the case of ALS data, instead, we will demonstrate that a good trade-off between accuracy and computing time can be achieved by exploiting the 2.5D characteristic of the data, that can be processed by custom FCNs originally developed for image segmentation tasks.

3. PROPOSED METHOD

As mentioned in Sec. 2, in order to take advantage of well-established CNN architectures, usually employed in the image processing field, we treat the ALS data classification as a

problem of image segmentation, solved with a FCN (Fig. 1). Similarly to the algorithm proposed in (Zorzi et al., 2019), our method is composed of two main stages (Fig. 2): (i) point cloud to image projection and (ii) image segmentation.

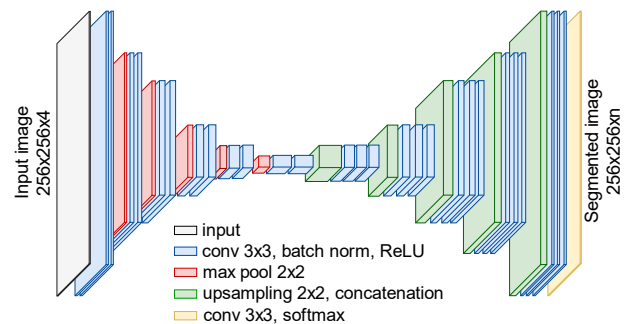


Figure 1. Architecture of the U-net model applied (best viewed in color). The segmented image has n channels, with n number of classes.

At first, the point cloud is projected into two-dimensional orthographic images, with one image channel storing the height of the point falling in the pixel. A 3D point cloud is thus represented as 2.5 data on a regular grid, which allows to efficiently take into account spatial positions and geometrical relationships between neighboring points. Moreover, other attributes recorded by the instrument are associated to three additional image channels, namely intensity, return number and total number of returns. In this way, the point cloud classification process is cast to an image segmentation problem, where the usual RGB channels are replaced by LiDAR attributes.

Because of the uneven spatial distribution of the 3D points, this projection cannot avoid collisions, that occurs when more than one point is mapped to the same pixel, unless a very small pixel size is chosen, which would have a negative effect both on computing time and classification accuracy. We cope with this by creating two different images, that are processed independently: in the first, the point with the highest altitude is assigned to the pixel, in order to enhance classification of thin objects such as power lines. In the second one, the lowest point is stored, improving the identification of the *ground* class, which is critical for the generation of DTMs. If more than two points fall in the same pixel, the ones with intermediate height inherits the label from the highest point. Thanks to this approach, we can use a pixel size of 0.10 m that is commensurate to the acquisition density of 18 points/m².

The image segmentation task is then performed by a FCN, that assigns a class label for each pixel and, consequently, to the corresponding point. In the last years, several FCN models have been proposed to solve semantic segmentation (Ciresan et al., 2012, Garcia-Garcia et al., 2017). For our application, we started from the popular U-net architecture (Ronneberger et al., 2015), already applied also in (Zorzi et al., 2019) and specifically adapted to segment the four-channel images created as previously described.

The implemented FCN is composed only of convolutional layers without any fully-connected one, allowing to operate on an input of any dimension and obtaining an output segmented image of corresponding size (Long et al., 2015).

More in detail, the first part of the network consists of a contracting path made of typical convolutional layers, whose task

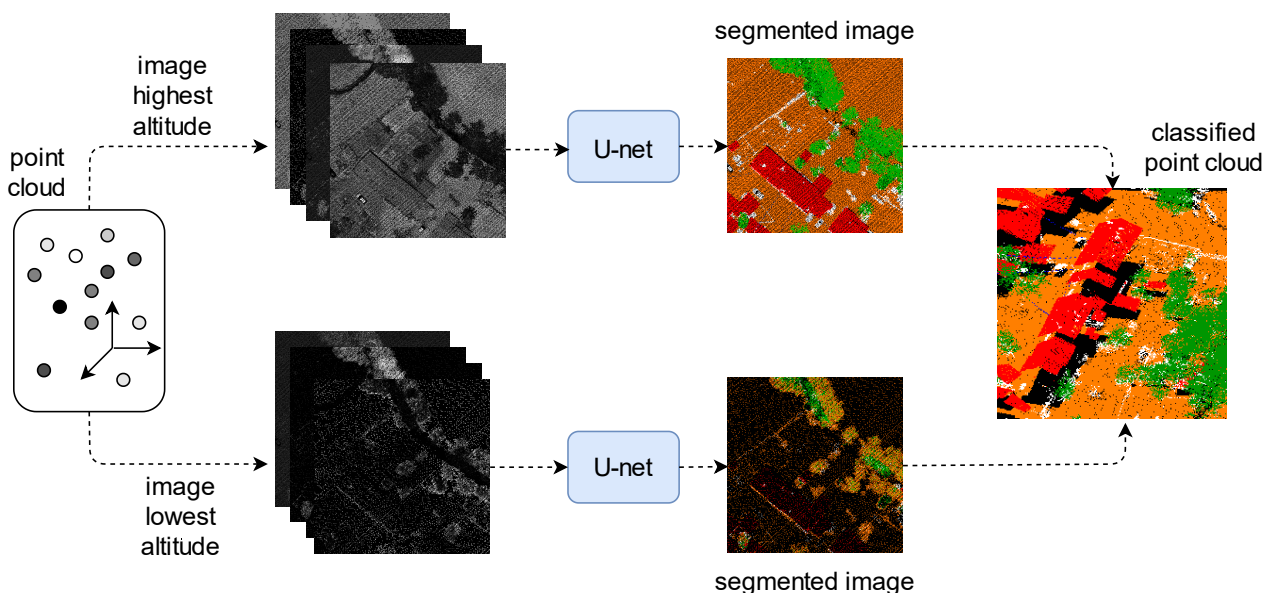


Figure 2. Workflow of the proposed method. At first, the point cloud is projected to multi-channel images. Images are then segmented by a FCN (U-net) that assigns a label to each pixel and, consequently, to each point.

is to extract low and high level features, capturing context information (as done by a custom CNN). Each layer in the contracting path performs two convolution operation with filters of size 3×3 , each followed by batch normalization and ReLU activation function. Max-pooling of size 2×2 is then applied to halve the representation size, with the first layer designed to take as input an image of dimensions 256×256 pixels and the final layer producing feature channels of size 8×8 . At each layer the number of feature maps is doubled with respect to the previous one, starting from 32 maps produced by the first layer to 1024 of the last one.

The contracting path is followed by an almost symmetrical one, known as expansive path, whose role is to enable precise localization, allowing a *per-pixel* labeling (Ronneberger et al., 2015). Each layer of the expansive path is constituted by an up-sampling of the output of the previous layer, a concatenation operation with the corresponding feature maps from the contracting path and three convolutions with filters of size 3×3 , followed by batch normalization and ReLU activation function. The final layer is characterized by a 1×1 convolution operation followed by a softmax activation function, that is used to reduce, for each pixel, the 32 components feature vector into a vector of dimension equal to the desired number n of classes. The detailed architecture is shown in Fig. 1.

Please note that the first layer of a FCN can take as input an image of fixed size (in our case, 256×256 pixels). A tiling with overlapping windows is thus adopted to process the entire dataset of arbitrary dimensions.

4. EXPERIMENTS AND RESULTS

The network was implemented in Keras (Chollet et al., 2015) and ran on a PC Intel Core i7 with 16GB RAM and a NVIDIA GeForce 1080 GPU.

4.1 Dataset

To train the network and test the performance in terms of accuracy, efficiency and speed, we employed a large-scale dataset,

Table 1. Number of points for the training/validation and test sets, divided according to the five classes.

Class	TRAIN/VAL		TEST	
	# Points	%	# Points	%
<i>ground</i>	$6.962 \cdot 10^8$	48.26	$7.738 \cdot 10^8$	37.13
<i>vegetation</i>	$5.228 \cdot 10^8$	36.24	$1.029 \cdot 10^8$	49.38
<i>roof</i>	$7.806 \cdot 10^7$	5.40	$6.033 \cdot 10^7$	2.89
<i>overground</i>	$1.447 \cdot 10^8$	10.03	$2.200 \cdot 10^8$	10.56
<i>power line</i>	$1.003 \cdot 10^6$	0.07	$8.461 \cdot 10^5$	0.04

acquired by Helica s.r.l. with a Riegl LMS-Q780 laser scanner. It covers the entire Friuli Venezia Giulia region (7700 km^2) with a mean density of 18 points per m^2 and it is divided into tiles of 0.2 km^2 . The flights were performed between December 2017 and July 2019, at an average altitude of 500 m above ground level and ensuring an overlap of 30% between adjacent flight-lines. In addition to the 3D coordinates of the points, the instrument registered also the intensity value of the reflected signal (a quantity related to the reflectance properties of the hit target), the total number of returns for each emitted laser pulse, as well as the return number associated with each echo.

We selected several areas for a total of 54 km^2 characterized by different land cover types (urban, rural and forest environments) that were manually classified among five classes: *ground*, *vegetation* (height from ground $> 2\text{m}$), *roof*, *overground* and *power line*. Please note that the *overground* class contains the objects that do not fall in the other classes, including cars, fences, walls, chimneys and low vegetation (height $< 2\text{m}$). The classified dataset was then split in 25 km^2 for training/validation and 29 km^2 for testing. Table 1 shows the points distribution over the classes. One can notice that the dataset is very imbalanced, because of the different shape and size of the scanned objects: the number of points classified as ground and vegetation is much higher than the samples falling in the roof and power line classes. As specified in Sec. 4.2, the unevenly distribution over the classes is an important aspect to take into account when designing the training of the network.

Table 2. Precision (%), recall (%) and F1-score achieved in the test set for each class.

Class	Prec.	Rec.	F1-score
<i>ground</i>	95.4	95.0	0.952
<i>vegetation</i>	92.8	98.5	0.956
<i>roof</i>	93.6	95.6	0.946
<i>overground</i>	80.9	58.7	0.680
<i>power line</i>	74.9	84.0	0.792

4.2 Training

The training of the model was performed applying categorical cross-entropy as loss function, Adam optimizer (Kingma, Ba, 2014) with 0.0002 learning rate and the weight initialization approach described in (Glorot, Bengio, 2010). Using a batch size of eight images (limited by the GPU memory), the process reached convergence after 30 epochs, requiring approximately 60 hours.

As data augmentation strategy, each tile composing the training set was randomly rotated four times and, for each rotated configuration, we extracted 240 images of size 256×256 . This approach proved to be fundamental to prevent the network from learning a specific scan pattern, which would have led to inaccurate results when classifying point clouds acquired along a different flight direction, as demonstrated by our experiments. Moreover, to take into account the unbalancing of the point distribution over the classes, we ensured that respectively 30% and 35% of the training images contained pixels belonging to *roof* and *power line*, which are the under-represented classes.

4.3 Testing

The proposed method reached an overall accuracy of 92.9% on the test set, while the average per-class accuracy is 86.3%. Table 2 shows precision (i.e. the number of points correctly classified as x divided by the number points classified by the algorithm as x), recall (i.e., the number of points correctly classified as x divided by the number of points belonging to class x) and F1-score (i.e., the harmonic mean of precision and recall) for each class.

	<i>ground</i>	<i>vegetation</i>	<i>roof</i>	<i>overground</i>	<i>power line</i>
<i>ground</i>	0.95	0.03	0.00	0.02	0.00
<i>vegetation</i>	0.00	0.98	0.00	0.02	0.00
<i>roof</i>	0.02	0.00	0.96	0.02	0.00
<i>overground</i>	0.15	0.25	0.01	0.59	0.00
<i>power line</i>	0.00	0.09	0.01	0.06	0.84

Figure 3. Confusion matrix: each row represents the instances in an actual class, whereas each column represents the instances in a predicted class. Please note that values are normalized so that the sum of every row is equal to 1.

As can be noticed also from the confusion matrix represented in Fig. 3, the algorithm performs well even for a challenging class such as *power line*; on the other hand *overground* is often misclassified. This is mainly due to the fact that this class contains different objects, including low vegetation that the network usually labels as *vegetation*. However, this error can easily be corrected by simple height thresholding.

To further investigate the performance of the method, we divided the test set into two main scenarios: (i) urban and flat areas and (ii) mountainous environments, and independently evaluated the results for the two area types. Figure 4 shows the confusion matrices for the analyzed cases. The proposed algorithm allowed to achieve an overall accuracy of 95.8% in urban and flat areas, ensuring good performance also in mountainous environments (92.2%). One can notice that the accuracy of *vegetation*, *roof* and *power line* classes does not significantly change between the different scenarios, reaching high values even in forest environments and high alpine terrain. In the presence of complex topography and steep slopes, instead, low vegetation (belonging to the *overground* class) and ground are sometimes confused (a behavior highlighted also in (Winiwarter et al., 2019)), causing a decrease in the classification accuracy of *ground* and *overground* classes. Some results for different area types are presented in Fig. 5.

In order to manage large datasets, a requirement of the classification algorithm is to be computationally efficient: our approach showed an inference time of only 11 minutes per km^2 (including the time for reading and writing point cloud files in LAS format), which allowed its application to the whole set of 7700 km^2 ($\approx 138 \times 10^9$ points), significantly reducing the large amount of time that is usually spent by the companies for the manual classification. The parallelization of the point cloud to image projection step could lead to further improvement in terms of time efficiency and productivity.

5. CONCLUSION

The good trade-off between speed and accuracy that characterizes end-to-end deep learning approaches makes these algorithms superior to established methods based on hand-crafted features.

In this paper we presented a deep learning approach for point cloud classification that, thanks to an overall accuracy of 92.9% and a low inference time (11 minutes per km^2), was effectively applied for the classification of a whole large-scale dataset of 7700 km^2 , covering the entire Friuli Venezia Giulia region (Italy). Half of the region is mountainous, yet the algorithm performs satisfactorily also in such challenging environments as alpine areas. The network was trained on a subset of only 25 km^2 and the reached accuracy allowed a massive reduction of the manual work that is usually spent to correct the misclassification errors produced by commercial software routines.

As a future work, we will test other network architectures, focusing on the implementation of methods whose computing time is compatible with the processing of large-scale datasets, such as the one presented in this paper. Particular attention will be dedicated to mountainous environments, that pose stimulating challenges.

ACKNOWLEDGMENT

The authors are grateful to Regione Autonoma Friuli Venezia Giulia (Italy) for allowing the publication of the results obtained on the LiDAR dataset acquired as part of the regional remote sensing project.

	ground	vegetation	roof	overground	power line
ground	0.98	0.01	0.00	0.01	0.00
vegetation	0.00	0.97	0.00	0.03	0.00
roof	0.03	0.00	0.95	0.02	0.00
overground	0.10	0.13	0.07	0.70	0.00
power line	0.00	0.07	0.02	0.07	0.84

	ground	vegetation	roof	overground	power line
ground	0.93	0.04	0.00	0.03	0.00
vegetation	0.00	0.99	0.00	0.01	0.00
roof	0.01	0.00	0.97	0.02	0.00
overground	0.16	0.26	0.01	0.57	0.00
power line	0.00	0.10	0.01	0.05	0.84

Figure 4. Confusion matrices, normalized per rows. Left: results achieved in urban and flat areas. Right: accuracy assessment in mountainous environments.

REFERENCES

Briechele, S., Krzystek, P., Vosselman, G., 2019. Semantic labeling of als point clouds for tree species mapping using the deep neural network pointnet++. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 42-2/W13.

Chollet, F. et al., 2015. Keras. <https://keras.io>.

Ciresan, D., Giusti, A., Gambardella, L. M., Schmidhuber, J., 2012. Deep neural networks segment neuronal membranes in electron microscopy images. *Advances in neural information processing systems*, 2843–2851.

Garcia-Garcia, A., Orts-Escolano, S., Oprea, S., Villena-Martinez, V., Garcia-Rodriguez, J., 2017. A review on deep learning techniques applied to semantic segmentation. *arXiv preprint arXiv:1704.06857*.

Glorot, X., Bengio, Y., 2010. Understanding the difficulty of training deep feedforward neural networks. *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, 249–256.

Goodfellow, I., Bengio, Y., Courville, A., 2016. *Deep learning*. 1, MIT press Cambridge.

Griffiths, D., Boehm, J., 2019. A review on deep learning techniques for 3D sensed data classification. *Remote Sensing*, 11(12), 1499.


Kingma, D. P., Ba, J., 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.

Landrieu, L., Simonovsky, M., 2018. Large-scale point cloud semantic segmentation with superpoint graphs. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 4558–4567.


Long, J., Shelhamer, E., Darrell, T., 2015. Fully convolutional networks for semantic segmentation. *Proceedings of the IEEE conference on computer vision and pattern recognition*, 3431–3440.

Maset, E., Carniel, R., Crosilla, F., 2015. Unsupervised classification of raw full-waveform airborne lidar data by self organizing maps. *International Conference on Image Analysis and Processing*, Springer, 62–72.


Qi, C. R., Su, H., Mo, K., Guibas, L. J., 2016. Pointnet: Deep learning on point sets for 3d classification and segmentation. *arXiv preprint arXiv:1612.00593*.



(a)



(b)



(c)

Figure 5. Examples of classification results (a) in an urban scenario, (b) in a rural area characterized by vineyards (on the left) and (c) in a mountainous environment (best viewed in color). Classes: ground (orange), vegetation (green), roof (red), overground - including low vegetation - (white), power line (blue).

- Qi, C. R., Yi, L., Su, H., Guibas, L. J., 2017. PointNet++: Deep Hierarchical Feature Learning on Point Sets in a Metric Space. *arXiv preprint arXiv:1706.02413*.
- Rizaldy, A., Persello, C., Gevaert, C., Oude Elberink, S., Vosselman, G., 2018. Ground and Multi-Class Classification of Airborne Laser Scanner Point Clouds Using Fully Convolutional Networks. *Remote sensing*, 10(11), 1723.
- Ronneberger, O., Fischer, P., Brox, T., 2015. U-net: Convolutional networks for biomedical image segmentation. *International Conference on Medical image computing and computer-assisted intervention*, Springer, 234–241.
- Rutzinger, M., Höfle, B., Hollaus, M., Pfeifer, N., 2008. Object-based point cloud analysis of full-waveform airborne laser scanning data for urban vegetation classification. *Sensors*, 8(8), 4505–4528.
- Scaioni, M., Höfle, B., Baugarten Kersting, A., Barazzetti, L., Previtali, M., Wujanz, D. et al., 2018. Methods from information extraction from lidar intensity data and multispectral lidar technology. *The International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences*, 42(3), 1503–1510. doi.org/10.5194/isprs-archives-XLII-3-1503-2018.
- Serna, A., Marcotegui, B., 2014. Detection, segmentation and classification of 3D urban objects using mathematical morphology and supervised learning. *ISPRS Journal of Photogrammetry and Remote Sensing*, 93, 243–255.
- Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., Rabinovich, A., 2015. Going deeper with convolutions. *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1–9.
- Tchapmi, L., Choy, C., Armeni, I., Gwak, J., Savarese, S., 2017. Segcloud: Semantic segmentation of 3d point clouds. *2017 International Conference on 3D Vision (3DV)*, IEEE, 537–547.
- Tran, T. H. G., Ressel, C., Pfeifer, N., 2018. Integrated change detection and classification in urban areas based on airborne laser scanning point clouds. *Sensors*, 18(2), 448.
- Wagner, W., Hollaus, M., Briese, C., Ducic, V., 2008. 3D vegetation mapping using small-footprint full-waveform airborne laser scanners. *International Journal of Remote Sensing*, 29(5), 1433–1452.
- Wang, D., Wang, J., Scaioni, M., Si, Q., 2020. Coarse-to-Fine Classification of Road Infrastructure Elements from Mobile Point Clouds Using Symmetric Ensemble Point Network and Euclidean Cluster Extraction. *Sensors*, 20(1), 225. doi.org/10.3390/s20010225.
- Winiwarter, L., Mandlbürger, G., Schmohl, S., Pfeifer, N., 2019. Classification of ALS Point Clouds Using End-to-End Deep Learning. *PFJ—Journal of Photogrammetry, Remote Sensing and Geoinformation Science*, 87(3), 75–90.
- Wu, Z., Song, S., Khosla, A., Yu, F., Zhang, L., Tang, X., Xiao, J., 2015. 3d shapenets: A deep representation for volumetric shapes. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 1912–1920.
- Yang, Z., Jiang, W., Xu, B., Zhu, Q., Jiang, S., Huang, W., 2017. A Convolutional Neural Network-Based 3D Semantic Labeling Method for ALS Point Clouds. *Remote Sensing*, 9(9), 936.
- Zhao, R., Pang, M., Wang, J., 2018. Classifying airborne LiDAR point clouds via deep features learned by a multi-scale convolutional neural network. *International Journal of Geographical Information Science*, 32(5), 960–979.
- Zorzi, S., Maset, E., Fusiello, A., Crosilla, F., 2019. Full-Waveform Airborne LiDAR Data Classification Using Convolutional Neural Networks. *IEEE Transactions on Geoscience and Remote Sensing*, 57(10), 8255–8261.