Proceedings of the Seventh International Conference on
Railway Technology: Research, Development and Maintenance
Edited by: J. Pombo
Civil-Comp Conferences, Volume 15, Paper 17.4
Civil-Comp Press, Edinburgh, United Kingdom, 2026
ISSN: 2753-3239, doi: 10.4203/ccc.15.17.4

Automatic Object Deposits Detection Alongside Railway Tracks with LiDAR Point Cloud Sequences

M. Convert, T. Narayana Swamy, J. Sanchez and B. Salavati

**SNCF Réseau,
France**

Abstract

A railway scene involves a wide range of components including infrastructures (e.g. rails, catenary poles, ballast, etc.), environmental elements (vegetation, terrain, etc.) as well as temporary objects (stationary trains, machinery, etc.). While most of these components are either meant to be there or do not directly affect railway operations, some objects ideally should not be present on the tracks and their surroundings. These include moderately large objects like rail segments or ballast big bags, or smaller ones such as steel fishplates, nuts, and bolts which typically remain on-site following maintenance operations. While their impact may be limited, they still represent unused material that could be recycled, and can further raise safety concerns if too close to the tracks. In this paper we present an automated LiDAR-based method capable of detecting such object deposits, with a focus on the aforementioned larger elements. The precise modelization and geolocation capacities of LIDAR technology allows us to precisely locate the deposits across the entire railway network in France, and provide measurement estimates of their characteristics. The process relies on a comparison between a reference point cloud of a given railway scene in the past, and a more recent target point cloud of the same scene in which we aim to detect the deposits. The target is aligned with the reference, and any differences between the two point-clouds are tested as actual deposits or not, after filtering irrelevant sources of change such as vegetation.

Keywords: LiDAR, point cloud, remote sensing, railway maintenance, artificial intelligence, deposit detection

1 Introduction

a) SNCF Réseau

SNCF Réseau operates most of the railway network in France, managing 30 000 kilometers of tracks across the territory. Since 2018 the company has been investigating railborne LiDAR technology for topographic measurements. In recent years, the potential of LiDAR has garnered significant interest and is now increasingly being tested and integrated in various maintenance tasks. This is made possible thanks to several specialized surveillance trains equipped with LIDAR scanners circulating across France under a consistent schedule, capturing 3D models of the railway network as they pass, in the form of dense point clouds with high precision. Currently, about 180 000 kilometers of LiDAR data is being gathered every year, resulting in a historical database of the network over the last five years with an average frequency of one acquisition every eight weeks. This database is fully exploitable internally for analysis.

b) Object deposits

This paper addresses the topic of object deposit detection in the railway environment. A typical railway setting involves various elements, from the train infrastructures normally present to enable train circulation (e.g. rails, ballast or catenary poles), to natural environmental elements (e.g. vegetation and terrain). Object deposits correspond to elements found in the surroundings of train tracks, which belong to neither category. These include rail segments, unused sleepers, big bags used to contain ballast, nuts, bolts, steel fish plates, etc. which are occasionally left in place after railway maintenance without purpose. Their impact on circulation may be limited but nonetheless represent risks if they are too close to the tracks, potentially standing in the way of the loading gauge of trains. They could also facilitate malicious activity by ill-intentioned trespassers to interfere with traffic.

Figure 1.a: Rail segments left between train tracks

Figure 1.b: Ballast big bags left along train tracks

c) Our contribution

The aim of the project is to develop an automated end-to-end pipeline capable of detecting, locate and characterize object deposits found throughout the entire France railway network by exploiting the available historical LiDAR data. Beyond the practical implementation details of this system, our contributions reside in an adaptation and enhancement of the Iterative Closest Point (ICP) point cloud registration algorithm. We also formally define the concept of object deposits in a LiDAR railway setting and present how to handle them. Further, we outline the different obstacles we encountered during development and strategies for filtering and mitigating them. Finally, we provide insight into the task based on actual field expert surveys as needed.

2 Methodology

a) Data

The LiDAR point clouds used in the study have a density of more than one point per centimeter-cube. This resolution is sufficient for grasping the geometric characteristics of larger objects, but can be too coarse to recognize smaller ones like fish plates and bolts. Thus, we choose to focus on moderately large objects with a volume of at least 10 liters in this study, especially rail segments that are particularly common and are of higher concern to the operation teams. Figures 2.a and 2.b illustrate what rail segments and big bags look like in a LiDAR point cloud context.

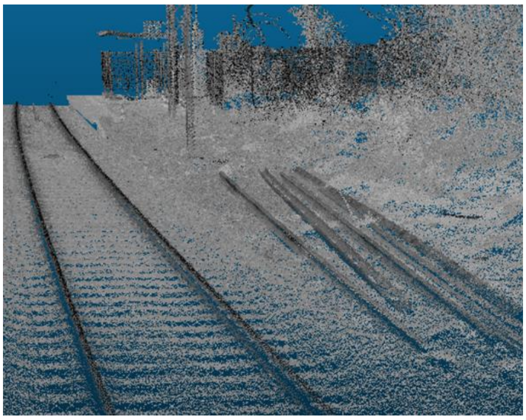

Figure 2.a: LIDAR point cloud of rail segments

Figure 2.b: LiDAR point cloud of big bags

The detection of object deposits relies on a comparison between a reference point cloud and a target point cloud of interest. The fundamental intuition that drives the proposed method is the fact the deposits we aim to detect are occasional and isolated incidents occurring at a given point in time and diverge from the nominal configuration of a scene. Thus, by assuming that such anomalies should not be observed in the usual configuration, we can compare past snapshots of a given LiDAR scene and use them as references to highlight the changes that may have occurred in a more recent target point cloud of the same scene. We therefore work on points clouds representing the same given scene at different times.

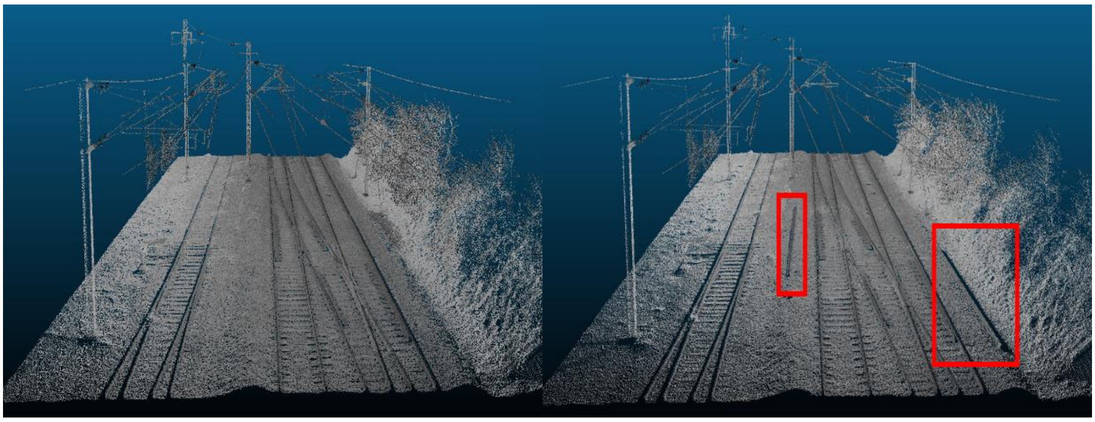

Figure 2.c: Past reference point cloud

Figure 2.d: Current target point cloud with anomalous rail segment deposits

b) Registration – change detection

The LiDAR data used for this study are generated alongside the GPS trajectory of the scanner mounted on the surveillance train. As such, the point cloud coordinates are already aligned with their actual geolocation on the railway network, i.e., 3D scans corresponding to the same railway segment appear at the same point in the coordinate system, overlapping with each other. This, however, is a coarse alignment due to the precision of the GPS which stands around 50cm on average and can go up to 3m in low reception environments like in tunnels.

To achieve a high precision change detection between two point clouds, we need to further align the past reference point cloud and the target point cloud with respect to their subtle geometries. The iterative closest point (ICP) algorithm [1] is commonly used in point cloud processing for this registration task. It aims to iteratively find a rigid transformation (rotation and translation) that minimizes the average distance between the points in the target and reference point cloud. Considering the aforementioned GPS precision is variable and fluctuates across long distances, all the subsequent operations are performed on sub-sections of the entire linear scan, each with a length of 50m to 100m. This ensures the operations (and especially the cloud alignment) is performed on locally coherent and stable geometries, which is a necessary assumption for rigid transformations to be relevant.

While applying the ICP on the entire point clouds is typical, we instead perform some preprocessing steps to refine the portions of the two point clouds to enhance the precision of the final registration. In a past study [2], a random forest model was trained to classify LiDAR points into different railway scene classes such as ground, vegetation, catenary poles, etc. It operates on a set of local geometric features on a per-point basis [3]. We used this model to first isolate the points corresponding to the

rails and the catenary poles in both the reference and target clouds, while filtering out all other classes, especially the ground and vegetation classes which are prone to evolution over time and depending on seasons. Rails and catenary poles are generally unmoving objects that correspond to good anchoring points to register two point cloud epochs. The rail class allows a registration on the XY-plane parallel to the ground, and the catenary poles bring further context in the vertical Z-axis to allow a precise vertical alignment as well as preventing erroneous “shifted” registration cases where the two clouds have their ground level properly aligned in the same XY-plane but deviate longitudinally in the direction of the rails.
Another filtering step is performed on the remaining rail and catenary pole points by discarding all points that stand above 30cm from the ground, which is measured from a digital terrain model (DTM) we generate beforehand. This is done with the aim of prioritizing precise registration of the ground as opposed to the higher portions of the clouds, as we only want to focus on the changes occurring at ground level. The ICP algorithm is finally applied on those preprocessed points and returns a transformation matrix, which is then applied to the target point cloud to register it to the reference.

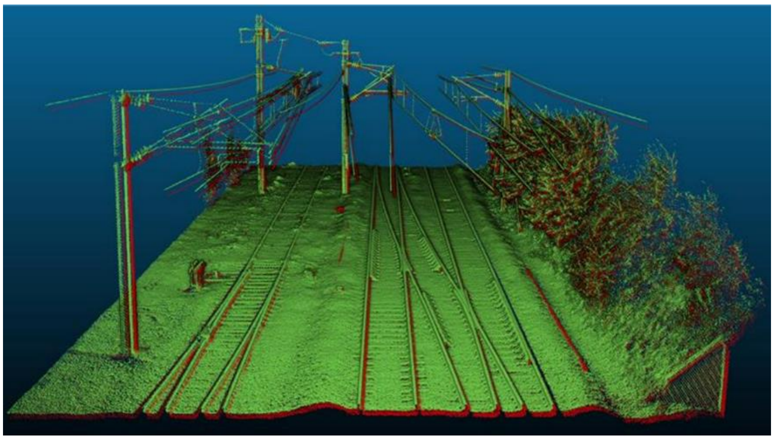

Figure 2.e: Overlap of reference point cloud (green) and target (red) before alignment

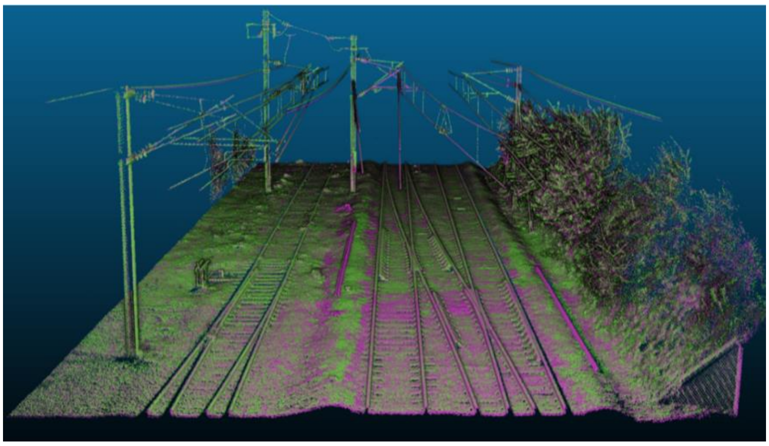

Figure 2.f: Overlap of reference point cloud (green) and target (purple) after alignment

The ICP algorithm involves a parameter that defines a threshold on the RMSE of the Euclidean distance between the two clouds during the registration process. In the classical implementations of ICP, this threshold parameter is fixed. In this paper, we propose an adaptive threshold that starts from a large value, and iteratively decreases until it reaches the final intended threshold. Two scans of the same scene acquired at different times can be shifted from each other to varying degrees depending on the precision level of the GPS system. The original ICP algorithm requires the two point clouds to have a significant overlap in the initial state to work properly. As such, using an adaptive threshold has two benefits. First, it allows us to handle cases where the two clouds are far apart initially, and better prevents the algorithm of getting trapped in local minima. Second, this method speeds up the convergence of the algorithm.

c) Extracting relevant information

Once the two point clouds are registered, we proceed to a refining step where we specifically target the zones of interest. Here, since we focus primarily on the deposits that could directly interfere with the loading gauge of the trains and its surroundings, we use the trajectory of the LiDAR scanner to extract only the points in the area on the side of the rails, ant limit the zone of interest up to a certain distance from them. The density of the points decreases the farther we go from the source of the scan, so the detection becomes less reliable. Once the zone of interest has been isolated, we focus only on the changes on the ground by filtering all points higher than 1m from the DTM. This still leaves larger objects like big bags in the study zone. Finally, we remove false positives belonging to the vegetation class.
The changes between the two clouds are identified by means of point subtraction. A spherical zone of radius [epsilon] is computed on each point in the reference point cloud. We then iterate through the target point cloud, and each point is removed if it is contained within the spherical neighborhood of any point in the reference. Doing so filters out all points that have correspondence in the reference, while leaving the ones that likely belong to a new object that appeared since the last scan. The [epsilon] parameter defines the threshold from which a point is considered a change or not. We typically set it to the average resolution of our point clouds, i.e., 2cm.

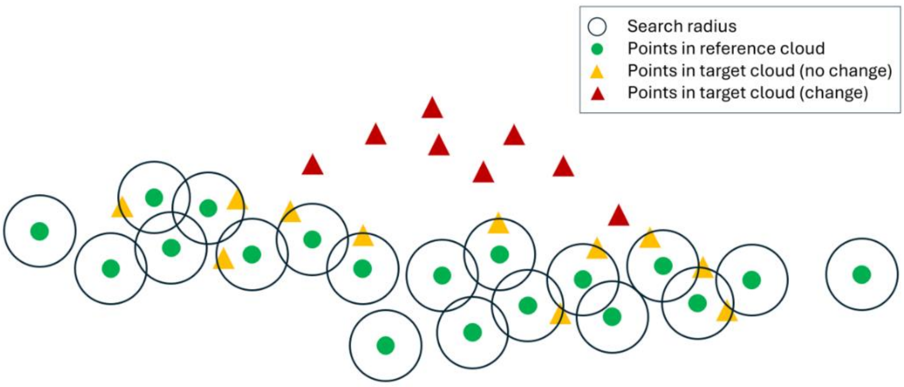

Figure 2.g: Illustration of change detection methodology (point subtraction)

Once we have highlighted the points that correspond to changes, we apply the DBSCAN clustering algorithm to discretize and segment the different objects. The clusters that are deemed too small are considered as either residual or objects beyond the scope of the study, and are discarded. A convex hull is computed for each cluster, and all clusters that have a volume greater than or equal to 10L are kept for the analysis. The final output of the system is the coordinates of the centroid computed for all remaining clusters, allowing a precise location of each object deposit found along the tracks.

3 Results

a) Quantitative analysis

The resulting system was evaluated on a total of 5 linear railway scans of several kilometers each, as displayed in Table 3.a). The system seems capable of detecting most of the deposits present in the respective scans except for Scan 2. It has however an overall tendency towards over-detection, resulting in a rather high recall but low precision. The case of Scan 2 is a good example of the limits of the system in its current state of the project. The 20 cases of false negatives missed by the system correspond to clusters of very long rail segments that were discarded due to a static geometric threshold that was defined at runtime to mitigate larger false positives that tend to appear due to cloud registration errors. This limit could be addressed in future work by refining the threshold value to use after further validation.
We also observe a high variability in the number of clusters detected across the different scans. This can be explained by the high variability in the characteristics and configurations of railway LiDAR scenes, where dynamic factors like vegetation growth and ballast displacement can introduce many artifacts after the different preprocessing steps. This could be addressed by developing a classification model that would recognize the different types of valid and invalid clusters.

Sample	Length (km)	Target Timestamp	Reference Timestamp	Detected Clusters	True Positives	False Positives	False Negatives
Scan 1	9.189	07/11/24	23/06/23	8	3	5	0
Scan 2	7.628	27/11/24	05/09/23	123	1	122	0
Scan 3	9.637	19/11/23	12/03/22	91	3	87	1
Scan 4	9.959	26/09/23	20/12/22	18	8	10	20
Scan 5	5.258	09/04/23	12/03/22	47	5	42	2

Table 3.a: Detailed overview of evaluation results

Sample	Precision	Recall
Scan 1	0.375	1.000
Scan 2	0.008	1.000
Scan 3	0.033	0.750
Scan 4	0.444	0.286
Scan 5	0.106	0.714

Table 3.b: Precision and recall for each scan example

b) Qualitative results

When checking the actual point clouds from the previously discussed results, we observe an overall good quality in the outputs of the system (c.f. Fig. 3. a-c) The true positives detected are highlighted precisely, and the system produces reliable estimations of the deposits' geometric characteristics.

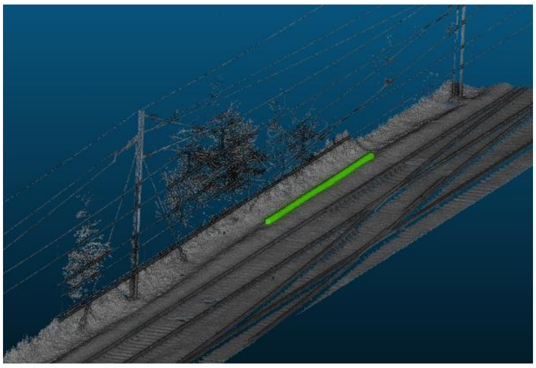

Figure 3.a: Rail segment detection

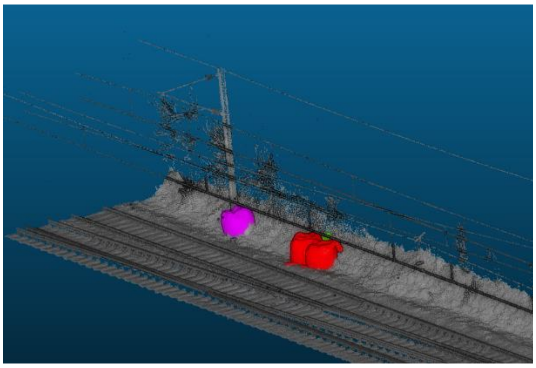

Figure 3.b: Big bag detection

Except for the false negative cases in Scan 2 which were due to large clusters exceeding the threshold, all rail segments deposits were properly detected. Beyond the rails which are the most common deposits, the system also highlighted various relevant objects such as isolated sleepers, blocks and maintenance equipment.

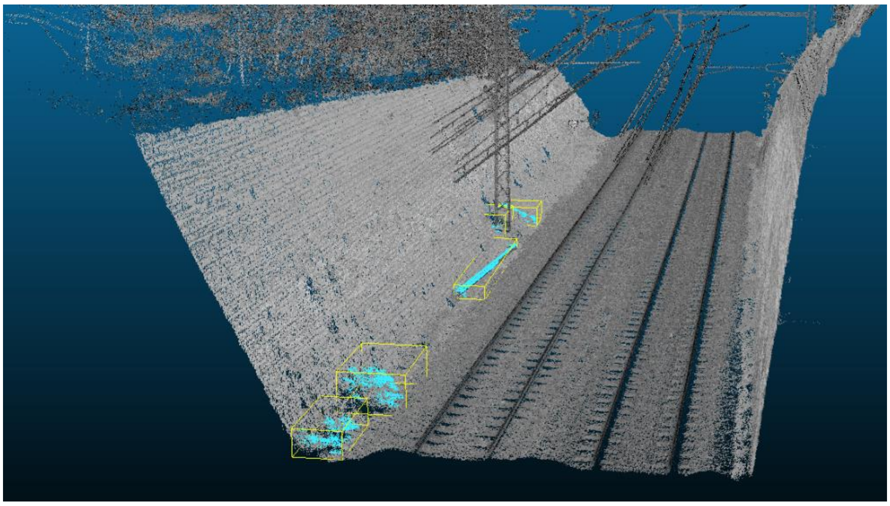

Figure 3.c) Detection of miscellaneous objects of larger size

The many false positives found in Scan 2 and Scan 3 have various causes. Some problems occurred on rigid and unmoving objects where the point cloud registration was not precise enough and the alignment distance exceeded the threshold set in the change detection step (Fig. 3.d and 3.e). Other cases were found on valid changes but that did not correspond to actual deposits (ground and vegetation). There were also occasional cases of unexpected changes caused by temporary objects like cars or passing trains that happened to be there at the time of the scan (Fig. 3.f and 3.g).

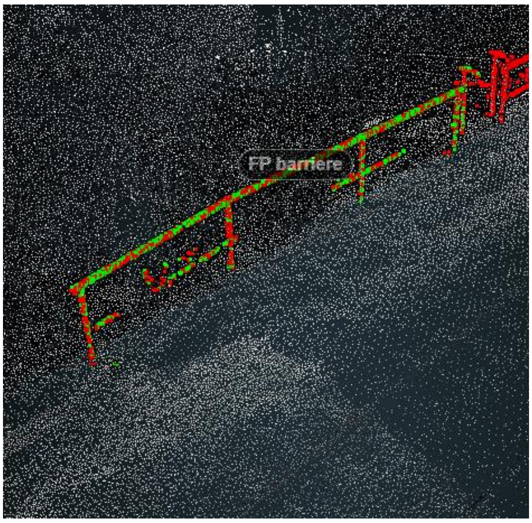

Figure 3.d: Registration error on fences

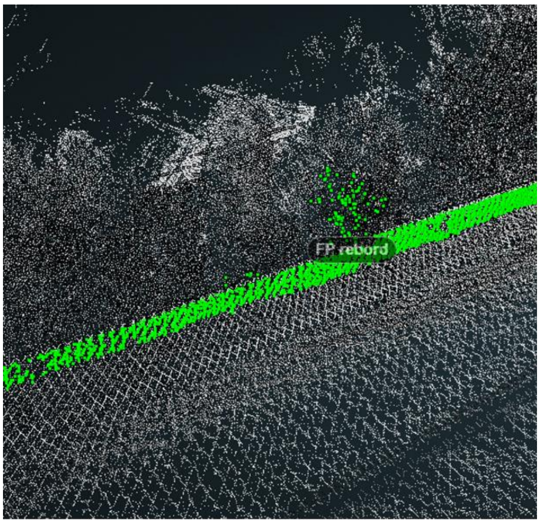

Figure 3.e: Registration error on pavement borders

Figure 3.f: Temporary cars

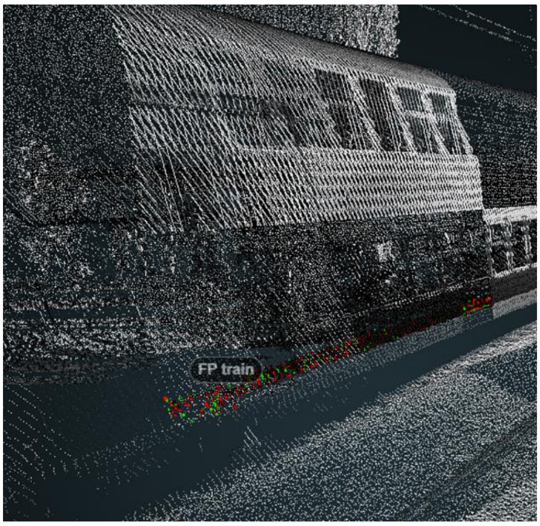

Figure 3.g: Temporary passing trains

4 Conclusions and Contributions

In this paper, we proposed a change detection method capable of identifying object deposits in railway LiDAR data. The system in its current state still requires further tuning of the different parameters involved in each processing step, as well as object deposit classification to refine the final results. It nonetheless represents a first step toward automatic surveillance for railway infrastructure.

References

[1] Besl, P. J., & McKay, N. D. (1992, April). Method for registration of 3-D shapes. In Sensor fusion IV: control paradigms and data structures (Vol. 1611, pp. 586-606). Spie.

[2] A.-S. Onody, M. Convert, F. Viguier, N. Abdallah, N. Ben Drihem, R. Guerand, J. Sanchez, B. Salavati, "Automated Detection of Vegetation Close to Catenary Wires using LiDAR Data", in J. Pombo, (Editor), "Proceedings of the Sixth International Conference on Railway Technology: Research, Development and Maintenance", Civil-Comp Press, Edinburgh, UK, Online volume: CCC 7, Paper 4.2, 2024, doi:10.4203/ccc.7.4.

[3] Weinmann, M., Jutzi, B., & Mallet, C. (2013). Feature relevance assessment for the semantic interpretation of 3D point cloud data. ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences, 2, 313-318.