How to Generate a Digital Terrain Model from a LiDAR Point Cloud: Complete Workflow

By **Francisco Navarrete Mandly**

Digital Terrain Model (DTM) generation from LiDAR point clouds is a fundamental process in geospatial engineering, supporting applications such as flood modeling, infrastructure planning, environmental monitoring, and terrain analysis. Although the task may seem simple at first glance, it encompasses a sequence of phases, each with its own technical complexity and decision points.

This article aims to explain the workflow used to produce a high-quality DTM from LiDAR data, detailing each processing step. We will also explore two alternative approaches to classify ground points: one based on a traditional method and one powered by Artificial Intelligence.

The screenshots shown are taken from Aplitop's Tcp PointCloud Editor, but the procedures are applicable to any point cloud processing software.

1. General Workflow

The process of generating the DTM begins with the original LiDAR point cloud, a dense collection of three-dimensional points that are typically acquired through airborne laser scanning. The entire workflow can be summarized as follows:
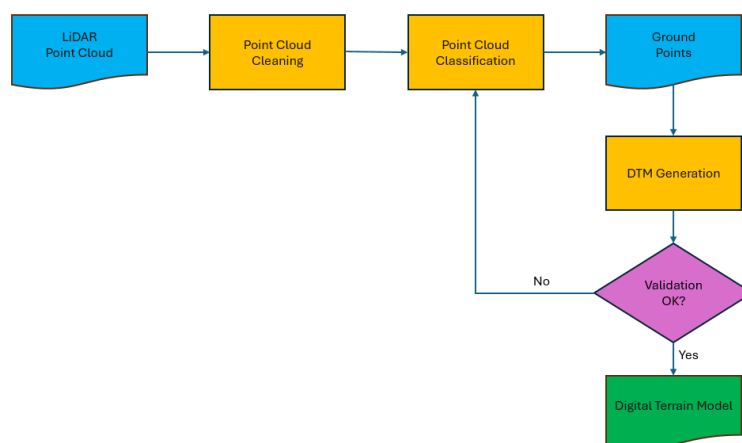

Figure 1. General Workflow

2. Point Cloud Cleaning

Raw LiDAR data inevitably contains anomalies and inconsistencies due to sensor limitations, atmospheric conditions, or the nature of the scanned environment. Therefore, point cloud cleaning is the essential first step. This process can include tasks such as:

2.1 Removing High Noise Points

Outliers with extreme or implausible elevation values are removed using statistical or geometric filters. These can be caused by reflections from birds, airplanes, or missed signals.

2.2 Elimination of Isolated Points

Isolated points, those that have no close neighbors within a defined radius, are often indicative of noise or erroneous returns. Their removal ensures a cleaner, denser dataset with better interpolation behavior.
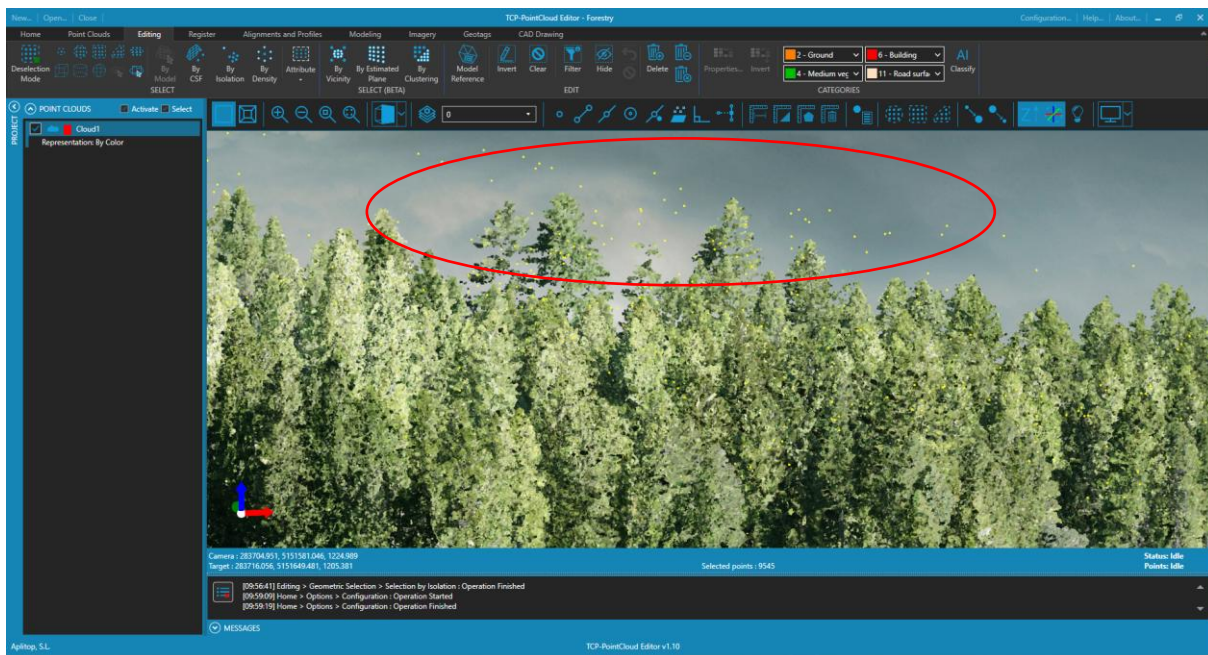

Figure 2. Isolated Point Filter

2.3 Filtering by Return Types

LiDAR pulses can have multiple returns, which are stored as additional point cloud attributes. If the main objective is to extract the ground, we can optionally discard the points that correspond to the first of many or intermediate.
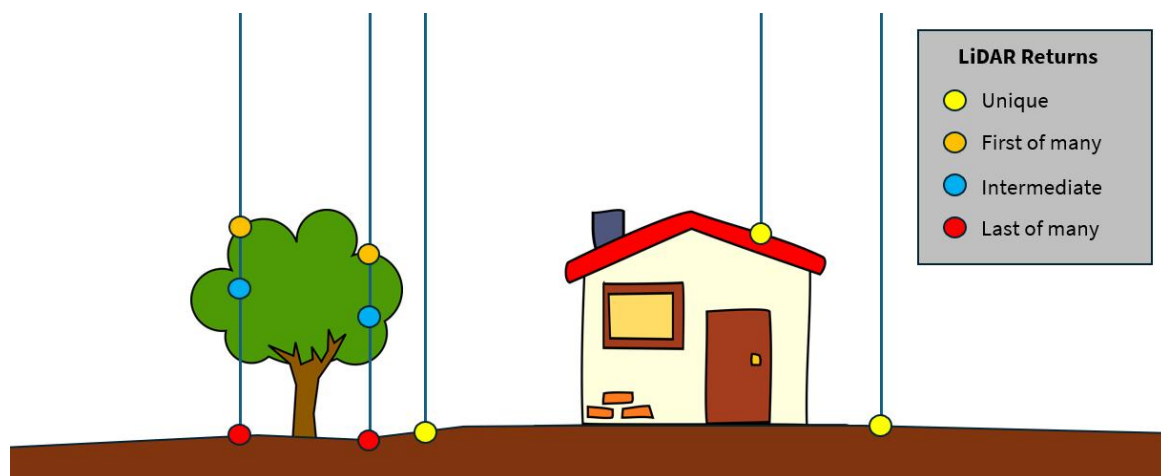

Figure 3. LiDAR Returns

3. Classification of Points

After cleaning, we obtain the filtered point cloud, ready to be classified by categories.

There are many classification techniques, typically grouped as heuristic, geometric segmentation, Machine Learning or Deep Learning-based methods. Below we will mention a traditional method and another based on models trained by Artificial Intelligence techniques.

3.1 Classification by Traditional Method

The Cloth Simulation Filter (CSF) method, introduced by Zhang et al. (2016), consists of inverting the point cloud and simulating a flexible cloth that covers it. The points of contact between the cloth and the "inverted surface" are considered the ground.

The steps in this method include selecting points that do not define the terrain, based on the curvature and slope thresholds. If we later reverse the selection, we obtain only the ground points, which are then assigned the corresponding category.

This algorithm is efficient, interpretable, and particularly effective for terrains with gradual elevation changes.
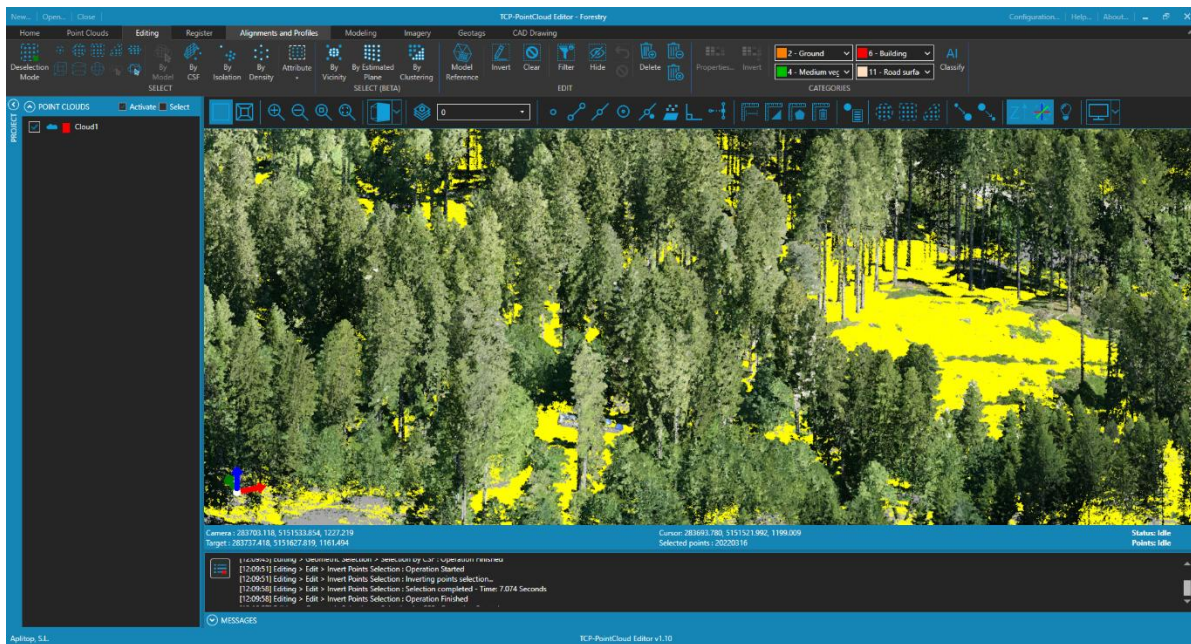

Figure 4. Ground Points Obtained After Reversing Selection by CSF

3.2 AI-Powered Classification

Artificial Intelligence (AI) offers powerful alternatives to traditional rule-based classification. This classification method uses a convolutional neural network architecture specifically designed to work directly with 3D point clouds without the need to transform them into voxels or meshes.

In this case, a model for exterior scenes has been previously trained, differentiating between terrain, buildings, vegetation, vehicles, fences, etc., which has been applied to the entire point cloud. Non-ground categories can be temporarily hidden to assess the classification results more clearly.
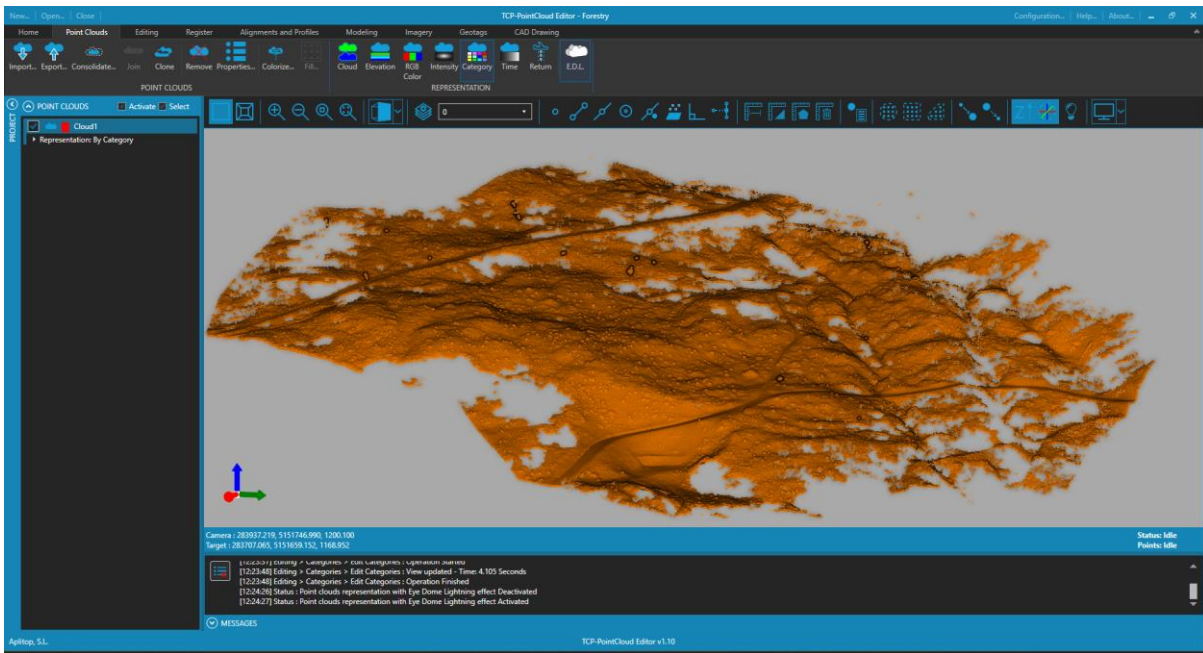
3

Figure 5. Ground Points Extracted via AI-Powered Classification

This technique has proven to be especially effective for semantic segmentation tasks in complex environments such as urban scenes or areas with dense vegetation, while maintaining high accuracy with a flexible and extensible architecture.

3.3 Manual Editing

Both approaches may result in misclassifications, especially in transition zones or occluded areas. Therefore, a manual editing stage is convenient to correct classification errors, ensuring accurate post-modeling.

If the software used offers advanced selection tools based on neighborhood or grouping, editing time can be greatly reduced.

4. Generation of the DTM

Once the terrain points have been isolated, the Digital Terrain Model is generated from them, represented as a Triangular Irregular Network (TIN) or a regular mesh.
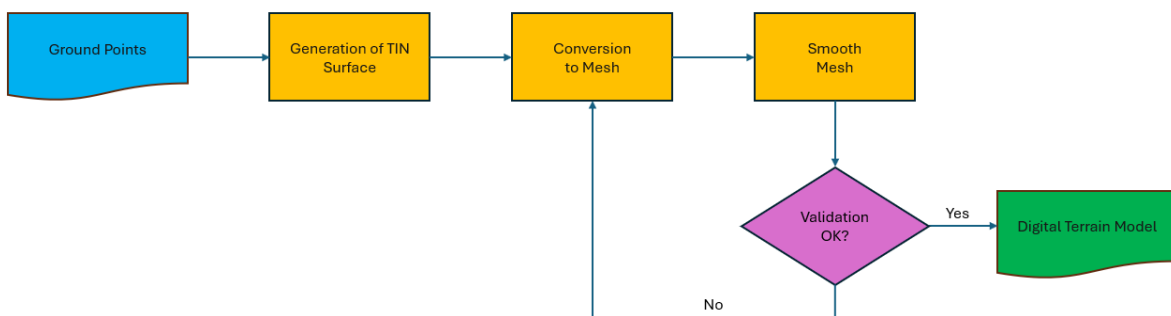

Figure 6. Flowchart for Digital Terrain Model Generation

4

4.1 TIN Surface Generation

To improve the fidelity of the result, it is recommended to initially generate a TIN surface from the ground points, carefully selecting appropriate minimum and maximum triangle edge lengths and adding breaklines, if they exist.

If peaks are detected on the surface, it may be due to points that have been misclassified, which can be corrected through partial deletion or reclassification.
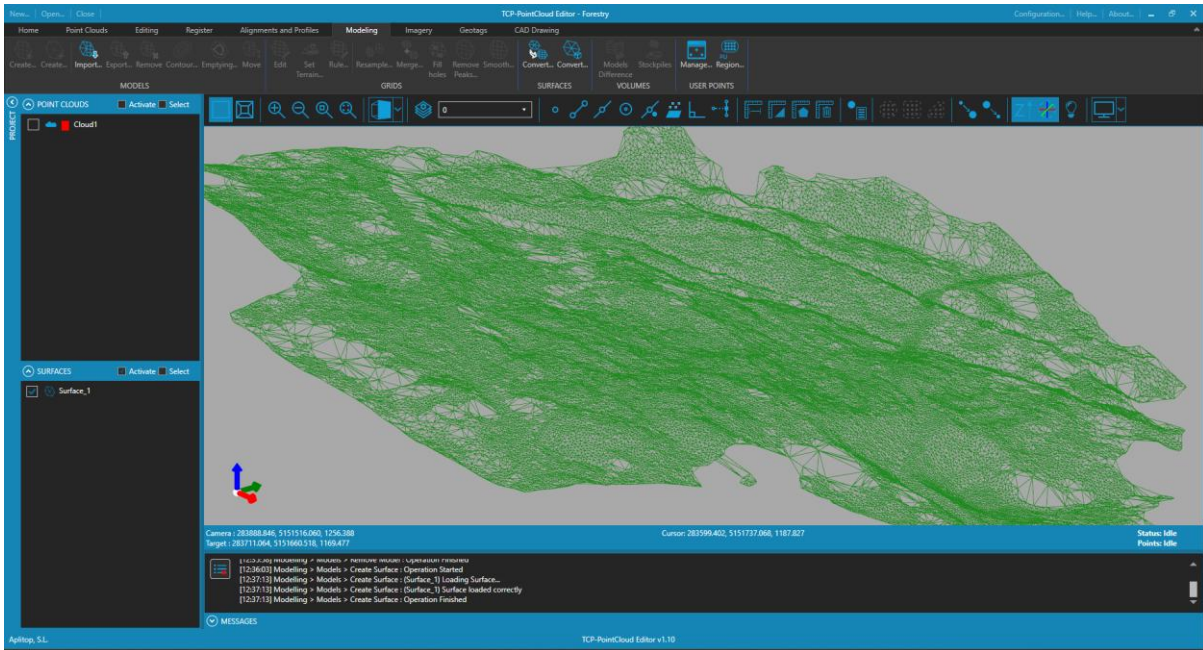

Figure 7. TIN Generated from the Ground Points

4.2 Conversion to Mesh

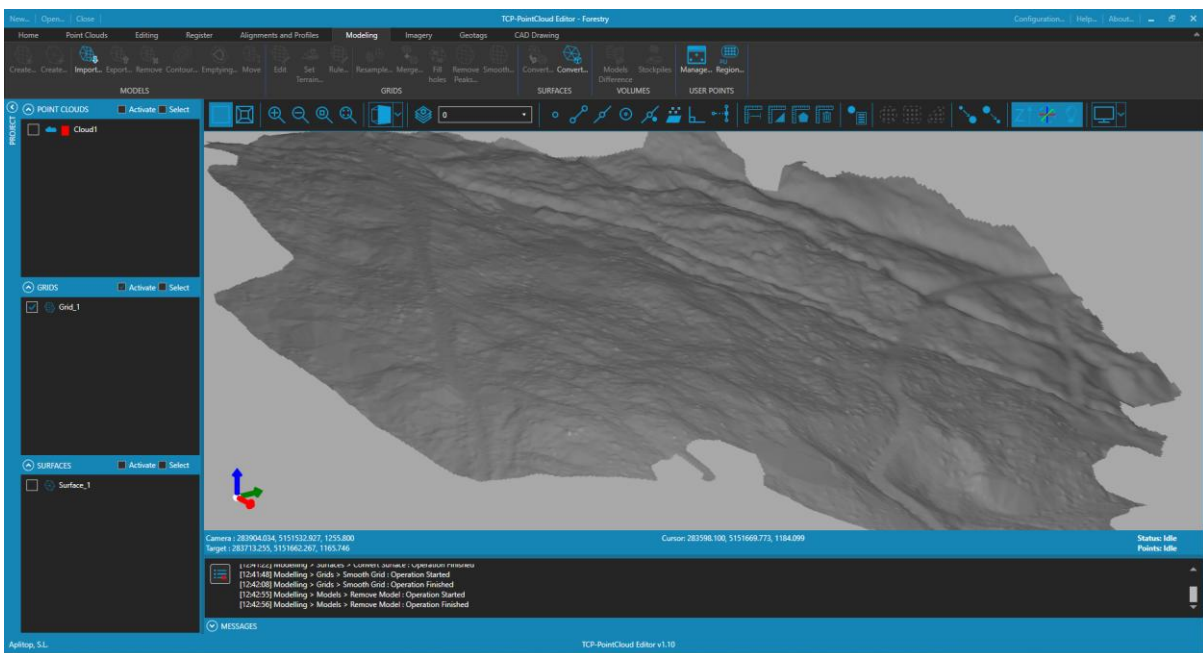The TIN is then converted into a regular mesh, with an appropriate cell size.

Figure 8. Regular Mesh Generated from the TIN

4.3 Mesh Smoothing

Finally, the mesh is smoothed to reduce abrupt transitions or noise introduced during triangulation, resulting in a seamless and topologically continuous visual model.
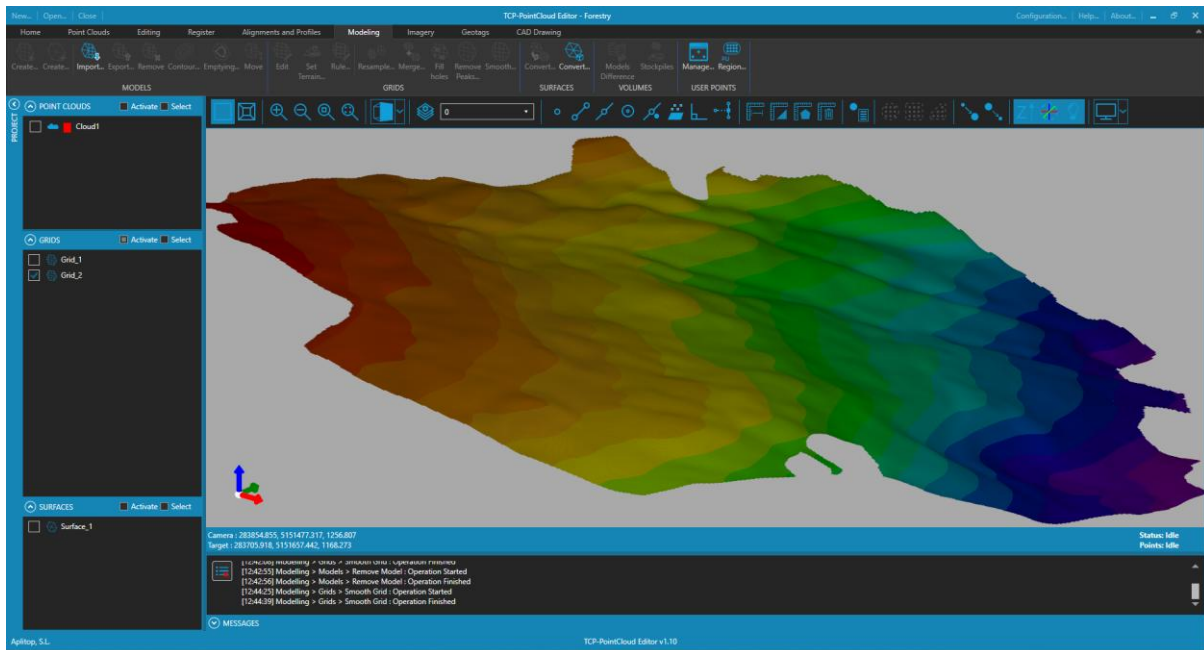

Figure 9. Smoothed Mesh Represented by Elevation

5. Validation

Before finalizing the DTM, it is highly recommended to evaluate its quality.

Among the possible options, contour lines can be generated to visually inspect anomalies. It can also be useful to generate profiles and compare elevation values of the TIN, original mesh, and smoothed mesh surface with those of the original points.
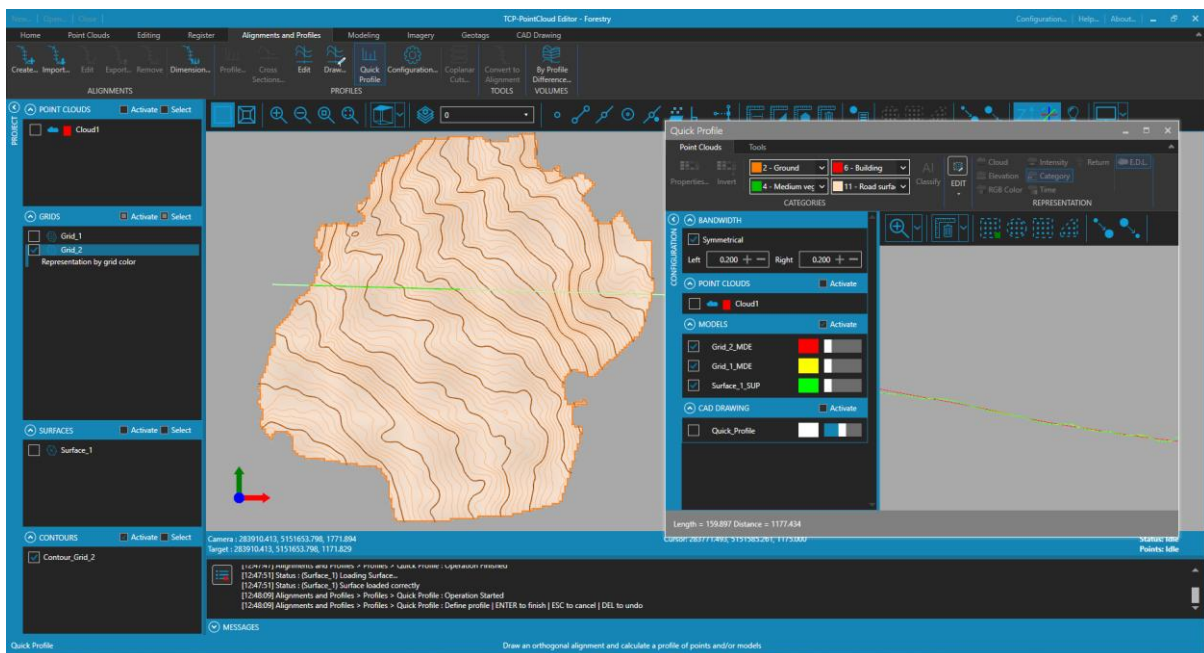

Figure 10. Contour Lines and Profile Comparison

If the result is not satisfactory, the manual editing of the classification would be repeated and the DTM would be regenerated with the corrected input data. This iterative refinement ensures that the final product meets the expected quality standards.

Conclusion

Transforming raw LiDAR data into an accurate digital model of the terrain is a multi-stage process, requiring careful attention to detail. Each stage (cleaning, classification, modeling, and validation) contributes to the geometric fidelity and usability of the final DTM.

Understanding this workflow is essential for professionals in geospatial sciences, civil engineering, and environmental analysis. Whether applying rule-based methods or leveraging AI-driven classification, the goal remains the same: to represent the terrain as accurately as the data allows.