From Points to Prints

Generation of building roofprints and footprints from ALS point clouds

Project Proposal

by

Alexandre Bry
6277535
abry@tudelft.nl

13 January 2026

IGN supervisor:	Bruno Vallet
1st TU Delft supervisor:	Hugo Ledoux
2nd TU Delft supervisor:	Ravi Peters

1. Introduction

In 2024, the Institut national de l'information géographique et forestière (IGN)[1] and two other French public entities have launched an initiative to bring together partners who can contribute to the development of a nation-wide digital twin (IGN, 2024a). In the same blog post, the IGN mentions ecological planning and sustainable land use as some of the priorities this project should accommodate. In a different post, it is explained how the LiDAR HD — the first project to collect high-density point clouds on almost the whole territory of France — is central to the future digital twin (IGN, 2024b). This comes from the unprecedented precision that it brings compared to previous data used and maintained by the IGN.

In this context, one of the many components of the future digital twin is buildings. Many algorithms have been developed to try to reconstruct simple but accurate 3D building models from various data sources, including point clouds. Some researchers from Delft University of Technology (TU Delft) especially developed an algorithm called roofer[2] which produced great results and was then applied to the whole of the Netherlands. This successfully created the 3DBAG, the first complete dataset of Dutch buildings in Level of Detail (LoD) 2.2 (Peters et al., 2022). This algorithm however requires two input data: a dense 3D point cloud and 2D building roofprints. In the Netherlands, the Actueel Hoogtebestand Nederland (AHN) was used for the point cloud and the Basisregistratie Adressen en Gebouwen (BAG) was used for the roofprints.

This is where things become more technical and where the precision provided by LiDAR HD becomes interesting. There are mainly two kinds of 2D building outlines, which are often used interchangeably, even though they can be significantly different once reaching the scale of centimetres or decimetres:

- **Footprints**: the 2D outer boundary defined by the vertical projection of the *outer walls/façades* of a building.
- **Roofprints**: the 2D outer boundary defined by the vertical projection of the *roof* of a building.

Usually, due to roof overhangs and gutters, the roof extends further than the walls, meaning that the footprint is included in the roofprint. In the rest of this document, I will use the terms roofprint and footprint when possible, and otherwise talk about outline when talking about any of them or when the differentiation was not made. As an example, the roofprint is what matters in estimating solar energy potential — in combination with other factors such as roof orientation and angle. But in many other applications — such as taxes or energy consumption — an accurate estimation of the area of buildings is necessary, which will be better with a footprint.

Adding this distinction to models is also what makes the difference between LoD 2.2 and LoD 2.3, as shown in Figure 1. Since roofer uses the points from the roof to reconstruct buildings, it requires a roofprint to work properly, but therefore reconstructs the buildings in LoD 2.2.

Moreover, different sources of data often make it easier to get either of the two:

- Experts on the field mostly use the walls and therefore measure the footprint.
- Experts working on aerial imagery can only use the roof as some walls will not be visible, meaning that they measure the roofprint.
- Airborne Laser Scanning (ALS) point clouds (such as the LiDAR HD and the AHN) give many points on the roofs and therefore make it easier to extract the roofprint.

[1] https://www.ign.fr/institut/identity-card
[2] https://github.com/3DBAG/roofer

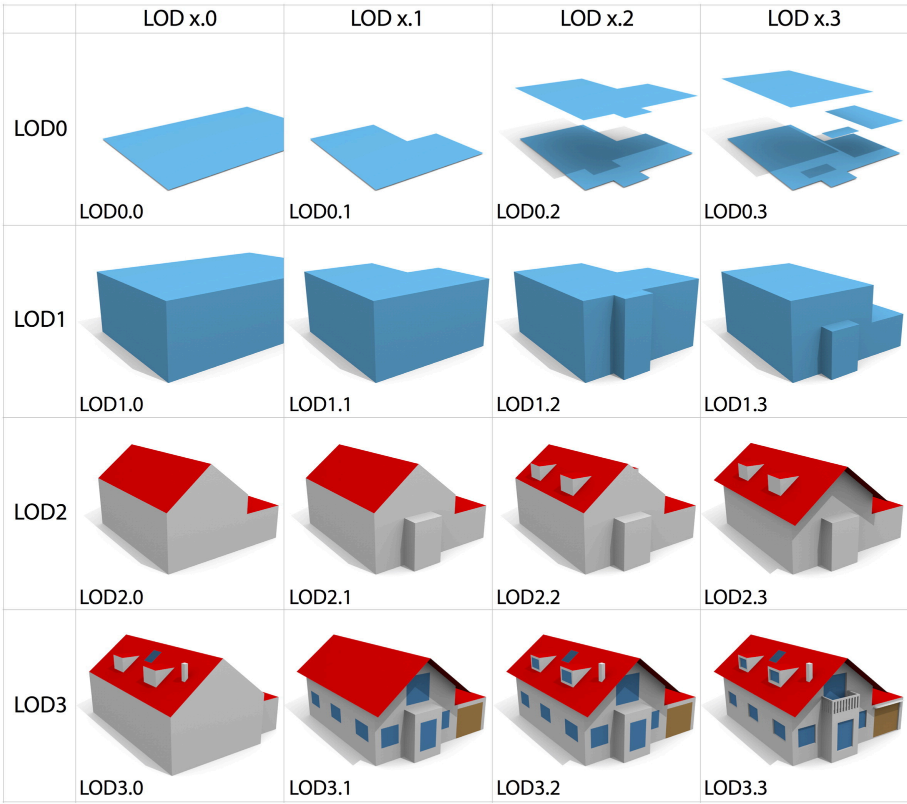

Figure 1: Visual example of the refined LoDs for a residential building (Biljecki et al., 2016).

- Terrestrial Laser Scanning (TLS) and Mobile Laser Scanning (MLS) point clouds give many points on the walls and therefore make it easier to extract the footprint.

The IGN already has a dataset containing building outlines, called BD TOPO. However this dataset has some issues, that can be explained by how it was historically built from different sources. First, some outlines come from terrain measurements and are therefore footprints, while others come from aerial image detection and are therefore roofprints. The dataset contains a column specifying for each outline which of the two it is, but it is missing for some buildings. Then, the georeferencing of these building outlines is often wrong by up to a few meters (see Figure 2 (*Chapter 5*)). This makes combining them with correctly georeferenced point clouds more complicated.

All in all, the current context combines:
- newly available data with high precision and correct georeferencing (LiDAR HD),
- an objective to build a digital twin of France, including 3D buildings with algorithms which would benefit from or require correct building outlines (such as roofer),
- an existing dataset that provides nation-wide and potentially great data but is however missing harmonization and precise georeferencing (BD TOPO),
- the example of the Netherlands where a great dataset of 3D building models was built from similar point cloud data (3DBAG from AHN),
- an interesting and not yet fully explored question of the possibility of extracting both an accurate footprint and an accurate roofprint from point clouds (more in *Chapter 2*).

2. Related work

2.a. Roofprints and footprints

I could only find a few research papers about identifying or reconstructing roof overhangs. The difference between footprints and roofprints is often not acknowledged, and most papers talk about outlines. In the few papers that try to find and compute roof overhangs, the methods are often similar.

The most common method consists in sweeping vertical planes perpendicularly to the roofprint edges or to the footprint edges depending on what was computed first. Then, a best-fitting plane is determined among these planes with different criteria. In (Panday & Gerke, 2012), a correlation score is computed for each plane, and the best result is kept only if it represents a sharp enough peak compared to its neighbours. For each edge of the footprint, Dahlke et al. (2015) computes the median height on segments parallel to the edge using a precise 2.5D Digital Surface Model (DSM) with a resolution between 5 and 20 cm. Then, they use the inflection point of the height variation as the roofprint edge. In (Frommholz et al., 2017), the roofprint is projected onto the 5 cm resolution DSM and the zero-crossings of the second-order derivative of height variation are used to estimate the size of the roof overhang.

Other methods are proposed by Goebbels & Pohle-Fröhlich (2023) to extend LoD 2 models by identifying potential overhang edges from the footprints and computing the size of the overhangs from either oblique images or point clouds. Using obliques images and assuming angles of 45°, they identify the roof overhang in the texture using either edges detection or colour regions, and compute the size of the overhang from this. The method with point clouds extends the roof planes and identify the inliers with a threshold.

Some interesting machine-learning methods were also proposed to compute building outlines from point clouds and could potentially be used to compute both footprints and roofprints. Girard et al. (2020) uses a machine learning model to compute, for each pixel of a RGB aerial image, a classification of building and building edges, as well as a frame field defining tangents and normals of the buildings. Then a multi-step geometric process is used to construct roofprints as polygons. Dai et al. (2025) uses a deep learning model that inputs a point cloud and outputs footprints as a binary raster. The model uses sparse voxel representations for the point cloud and decoder/encoder architectures with a specific 3D attention module. I also want to mention (Saadaoui et al., 2025) which uses an almost full deep learning pipeline to produce roofprints as polygons. A first model identifies building pixels, followed by a residual autoencoder to regularize the segmentation, and finally a lightweight CNN that extracts building corners that can be used for polygonization.

As a final note, Panday & Gerke (2012) use an interesting property of buildings that it worth bringing up. They assume that building roofs are often symmetrical when they are slanted, meaning that the overhang is often the same on both side. This allows to use more data to compute the size of the roof overhang, and could also help making a reasonable guess if there are no points on one side of a building but many points on the opposite side.

2.b. Point cloud semantic segmentation

Point cloud semantic segmentation consists in assigning a class among a list of predefined classes to each individual point of a point cloud. Since the roofprint is computed with the points on the roof and the footprint comes from the points on the façades, being able to tell them apart would really help the process. This is why I looked into point cloud semantic segmentation methods.

Some models take range-images as input, which makes processing much more efficient and allows to re-uses all the knowledge developed for processing 2D images (Biasutti et al., 2019; Xu et al., 2020). Range-images represent the 3D points clouds in 2D by looking at it from the viewpoint of the sensor, which often scans point clouds in lines with almost uniform angular steps. This is especially beneficial when fast inference is necessary, such as real-time applications.

Other models extend 2D convolutions to 3D in different ways. Thomas et al. (2019) introduced KPConv, a pointwise convolution with weights carried by points that are correlated with the points in the neighbourhood. Li et al. (2020) uses their own geometry-aware convolution and combine it with attention based on points elevation to process ALS point clouds. Zhao et al. (2020) took inspiration from self-attention networks for image analysis to build a Point Transformer Layer that can work as a convolutional layer by using the neighbourhood of the points. With farthest point sampling and trilinear interpolation it allows to build a U-net architecture called PointTransformer that outperformed other models at that time in semantic segmentation. Its latest iteration PointTransformerv3 improved accuracy, speed and memory usage by getting rid of costly 3D operations like neighbourhood queries which were replaced by serialization of the point clouds using space-filling curves (Wu et al., 2023).

Since ALS point clouds can also be very imbalanced, especially with roof points being easier to acquire than façade points, I also looked into papers dealing with class-imbalanced datasets. Li et al. (2024) advocated to decouple the optimization of the backbone and the classifier, by alternating them. Optimization of the backbone happens on a rebalanced set of points, using ground-truth points and pseudo-labelled points, and a custom loss is added against imbalance. Pan et al. (2025) proposed a framework to deal with sparse and inhomogeneous annotations, with a label-aware downsampling strategy and a gradient calibration function to compensate the bias introduced by annotation inhomogeneity.

2.c. Datasets

To be able to train a model and/or to evaluate the results, it would be useful to have datasets that can be used as ground-truth. Therefore, I looked for ALS datasets with semantic segmentation classes differentiating façades and roofs; and for datasets containing both footprints and roofprints.

Regarding building outlines, Dai et al. (2025) should publish a dataset with more than 3000 building footprints based on the ALS dataset called DALES (Varney et al., 2020). This one would be interested as footprints are the most difficult thing to extract from ALS point clouds as explained before. However, it is not yet available. Besides this one, I did not find other interesting datasets of footprints or roofprints.

Regarding ALS datasets, a list of the datasets I found are shown in Table 1. They vary significantly in size and point density, and only some of them have separate classes for roof and façade points. Some of them even have more precise classes than roof and façades, such as CENAGIS-ALS benchmark, which classifies separately stairs, balconies and chimneys as well, or DublinCity with windows and doors (Zachar et al., 2023; Zolanvari et al., 2019). Some of them are generated from images using computer vision techniques. The one called FRACTAL is interesting because it was made by the IGN with the LiDAR HD, even if it does not separate

Table 1: List of point cloud datasets for semantic segmentation of point clouds.

Name	Number of points	Point density (pts/m^2)	Classes for roof/façade	Source	Paper
H3D	74M	800	Yes	ALS	(Kölle et al., 2021)
V3D	0.78M	8	Yes	ALS	(Cramer, 2010; Niemeyer et al., 2014)
DublinCity	260M	348	Yes	ALS	(Zolanvari et al., 2019)
Campus3D	937.1M		Yes	Aerial images	(Li et al., 2020)
CENAGIS-ALS	550M	275	Yes	ALS	(Zachar et al., 2023)
DALES	505M	50	No	ALS	(Varney et al., 2020)
SensatUrban	2847M		No	Aerial images	(Hu et al., 2020)
LASDU	3.12M	4	No	ALS	(Ye et al., 2020)
TALD	121M	12	No	ALS	(Vijaywargiya & Ramiya, 2025)
FRACTAL	9261M	37	No	ALS	(Gaydon et al., 2024)

roofs and façades (Gaydon et al., 2024). The paper also explains how other French datasets such as BD TOPO were used to select the areas to include in order to reduce class imbalance.

Finally, to be able to create a dataset more quickly, Mérizette et al. (2025) explained how they created a new dataset for semantic segmentation of indoor TLS with two different processes: manual labelling and automatic generation of pseudo-labels from a BIM model of the objects. For the automatic generation, they had experts make a 3D BIM model of the rooms, and then used this model to classify points in the class of the closest object if close enough.

3. Research questions

The research question of this thesis will be:

> ***What characteristics of ALS point clouds influence the ability to generate accurate building roofprints and footprints?***

I will also try to look into the following sub-questions:

- How do the characteristics of the point cloud acquisition (geometry of the scanning device, trajectory of the flying vehicle) impact the whole process?
- Are there common characteristics of buildings that can help compensate the lack of points on some façades?
- How meaningful are the differences obtained between roofprints and footprints in France in terms of distances and areas?
- Can existing but inaccurate building outlines be helpful in guiding the creation of new accurate footprints and roofprints?

4. Methodology

Based on the literature review and on discussions with my supervisors, I plan to experiment with a process combining machine learning for point cloud semantic segmentation and geometric processing to compute the roofprints and footprints. The possibility of a full machine learning pipeline was rejected because time constraints and the lack of training dataset would likely make it difficult to get robust results.

The first step will be to identify in the point clouds the points corresponding to the façades and the points corresponding to the roof. To do so, I will train a deep learning model on a custom dataset. To build the custom dataset, I will use Dutch data mostly, by overlaying the AHN with the 3DBAG. By doing so, I should be able to classify the AHN between façade points, roof points and others. Then, since Dutch buildings have similarities with French buildings, and the AHN is also similar in density to the LiDAR HD, I hope to be able to simply fine-tune the model on a limited amount of French buildings. In the end, I should end up with the ability to extract points corresponding to buildings, separated between roof and façades.

Then, the geometric process will start by computing the roofprint, as ALS point clouds give many more points on the roofs than on the façades. How this will be done exactly will depend on the quality of the point cloud semantic segmentation, but many algorithms have been studied to extract roof planes, and I am only really interested in getting the boundaries of the roof. Once the roofprint will be computed, I will use it to identify the best vertical plane to match the façade points, parallel to each edge of the roofprint. The details still need to be decided and tested, but I hope that this method will work on façades even if they contain few points, and other principles such as symmetry, or the orientation of the roof planes could help to make the best guesses possible with façades containing too few points or no point at all. This combination of logical decisions and geometric computations should ensure that I can get a realistic result in every situation, compared to using machine learning. Moreover, I will try to integrate the existing outlines from BAG in the Netherlands and from BD TOPO in France to guide the reconstruction process. This could help because even though they are shifted in BD TOPO, they still provide interesting information in terms of shape, length of edges and angles between the edges.

Finally, assessing the results may be a difficult task. Any dataset with correctly modelled roof overhangs would be useful both in training the semantic segmentation model and in evaluating the method. However, without such high-quality ground-truths, I may have to build my own dataset of footprints and roofprints, in areas of varying difficulty, to assess the accuracy of the algorithm in different scenarios. One solution could be to generate buildings with overhangs by extending the roof planes of LoD 2.2 buildings, and then use a tool such as Heidelberg LiDAR Operations Simulator ++ (HELIOS++)[3] to simulate an ALS point cloud acquired on these buildings in different scenarios. The robustness of not using machine learning for the last part could also mean that the focus should be on assessing how the algorithm performs when the classification given by the model on the point cloud is inaccurate or even misleading. Therefore, there will likely be several iterations to get a robust algorithm in every possible configuration of point cloud and semantic segmentation.

[3] https://github.com/3dgeo-heidelberg/helios

(a) *First example with aerial images.*

(b) *Second example with aerial images.*

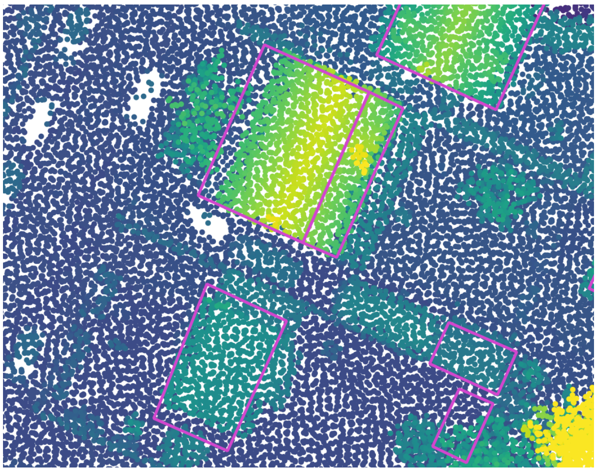

(c) *First example with LiDAR HD.*

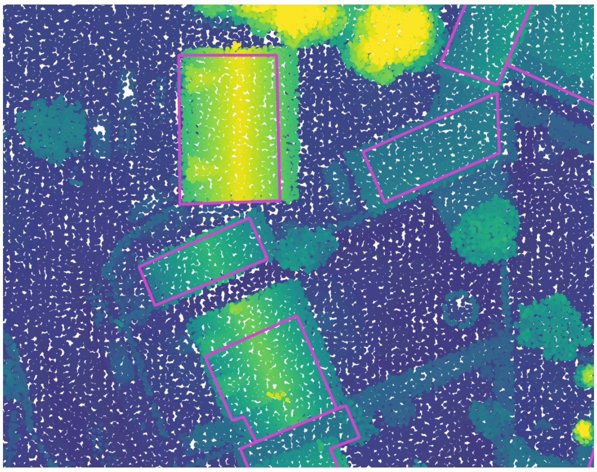

(d) *Second example with LiDAR HD.*

Figure 2: The misalignment of the outlines (in purple) with aerial images and LiDAR HD.

5. Preliminary results

I do not have yet many results to show, as most of the time until now was dedicated to studying literature to decide in which direction to go. I did however look into the different datasets in France and in the Netherlands that will be interesting for this project, and can illustrate why the methodology was designed as explained in *Chapter 4*.

First, the current outlines in BD TOPO are misaligned with both aerial images and with the LiDAR HD dataset. This is visible in Figure 2, especially compared to the point clouds.

Then, before starting to implement the classification of the AHN with the 3DBAG, I tried it with CloudCompare on a subset to check if it seemed reasonable and feasible. First, I looked at the distance between the points and the buildings, keeping only the points classified a “Not classified” and “Building”. The results are displayed in Figure 3, and show that the results are indeed promising, with most points from the buildings being at less than 1 m from the 3DBAG buildings.

Finally, I tried to classify the points based on distances to the 3DBAG with 3 classes (roof, facade and other) with roof being assigned to the point p if

$$\text{distance}\,(p, \text{roof}) < \min(\text{distance}(p, \text{facade}), 1) \tag{5.1}$$

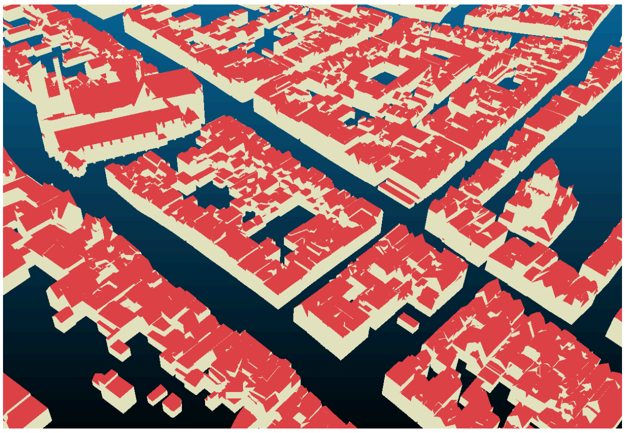

(a) First example: the buildings in LoD 2.2 from 3DBAG.

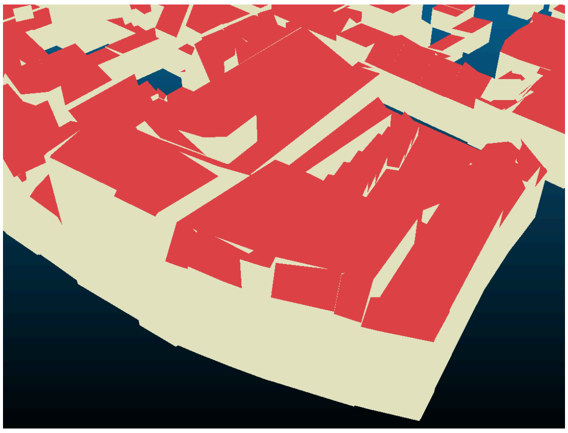

(b) Second example: the buildings in LoD 2.2 from 3DBAG.

(c) First example: the point cloud from AHN 4 filtered to keep only “Not classified” (green) and “Building” (white).

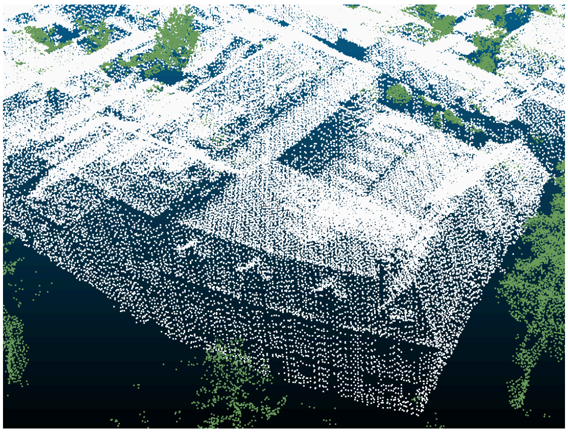

(d) Second example: the point cloud from AHN 4 filtered to keep only “Not classified” (green) and “Building” (white).

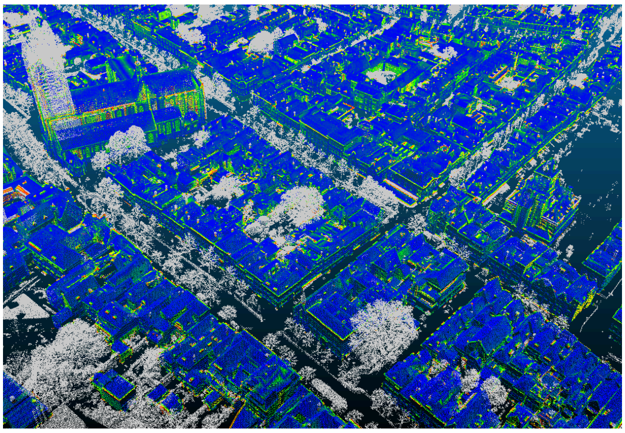

(e) First example: distance from AHN 4 points to 3DBAG buildings with 0 m for blue, 0.333 m for green, 0.666 m for yellow and 1 m for red. Values higher than 1 m are grey.

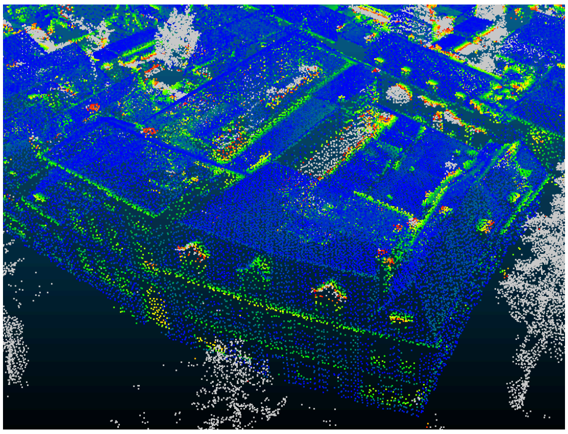

(f) Second example: distance from AHN 4 points to 3DBAG buildings with 0 m for blue, 0.333 m for green, 0.666 m for yellow and 1 m for red. Values higher than 1 m are grey.

Figure 3: Experiment with the distances between AHN 4 points and 3DBAG buildings in LoD 2.2.

This is possible because each polygon in the 3DBAG buildings is classified as either roof, wall, or ground surface. This gave the results shown in Figure 4. This shows that this method is promising to automatically produce a training dataset, and it also raises a few caveats:

(a) First example: the buildings in LoD 2.2 from 3DBAG.

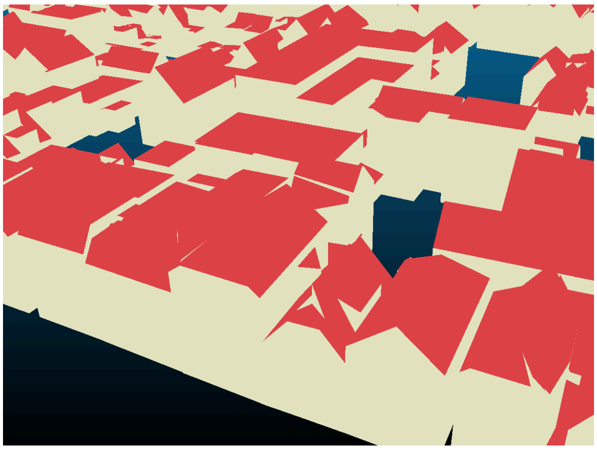

(b) Second example: the buildings in LoD 2.2 from 3DBAG.

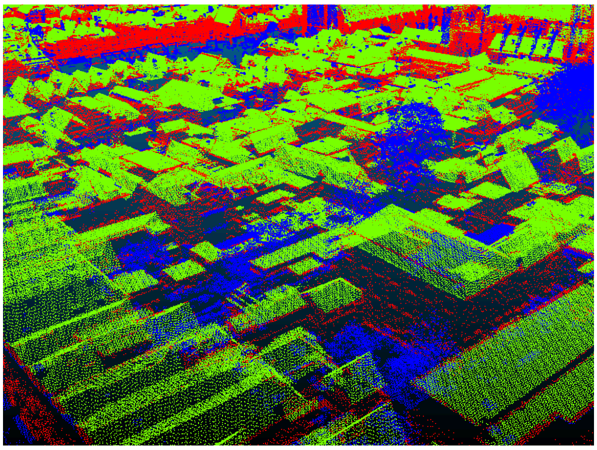

(c) First example: the custom classification of the points (green for roof, red for façade and blue for other).

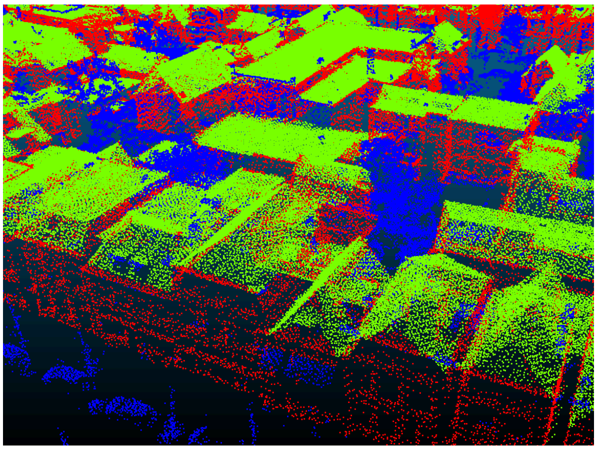

(d) Second example: the custom classification of the points (green for roof, red for façade and blue for other).

Figure 4: Experiment with classifying AHN 4 points into façade, roof and other from 3DBAG buildings in LoD 2.2.

- Most building parts that are not included in the LoD 2.2 models are not classified in the buildings: balconies, chimneys, dormers and in general all the objects on the roofs or on the walls. But since our focus is on the façades and roof boundaries, this is not a problem.
- Roof overhangs and/or the boundaries of the roofs are sometimes/often classified as façades because the LoD 2.2 models don't expand far enough on the sides. I plan to improve this by expanding the roof polygons to simulate the overhangs.
- As expected, monuments and modified buildings perform poorly, but are rare enough and not the focus of this project.

Currently, I am working on automating this process with C++ and Computational Geometry Algorithms Library (CGAL), in order to actually build a large enough dataset. The process will be:

1. Download the 3DBAG and AHN for a region
2. Separate the roof, façade and floor from the 3DBAG
3. Extract the points classified as "Building" or "No classification" in the AHN
4. Extend the polygons of the roofs downwards by a fixed amount to be determined
5. Compute for each point its distance to the roofs and to the façades

6. Classify the points based these distances

6. Time planning

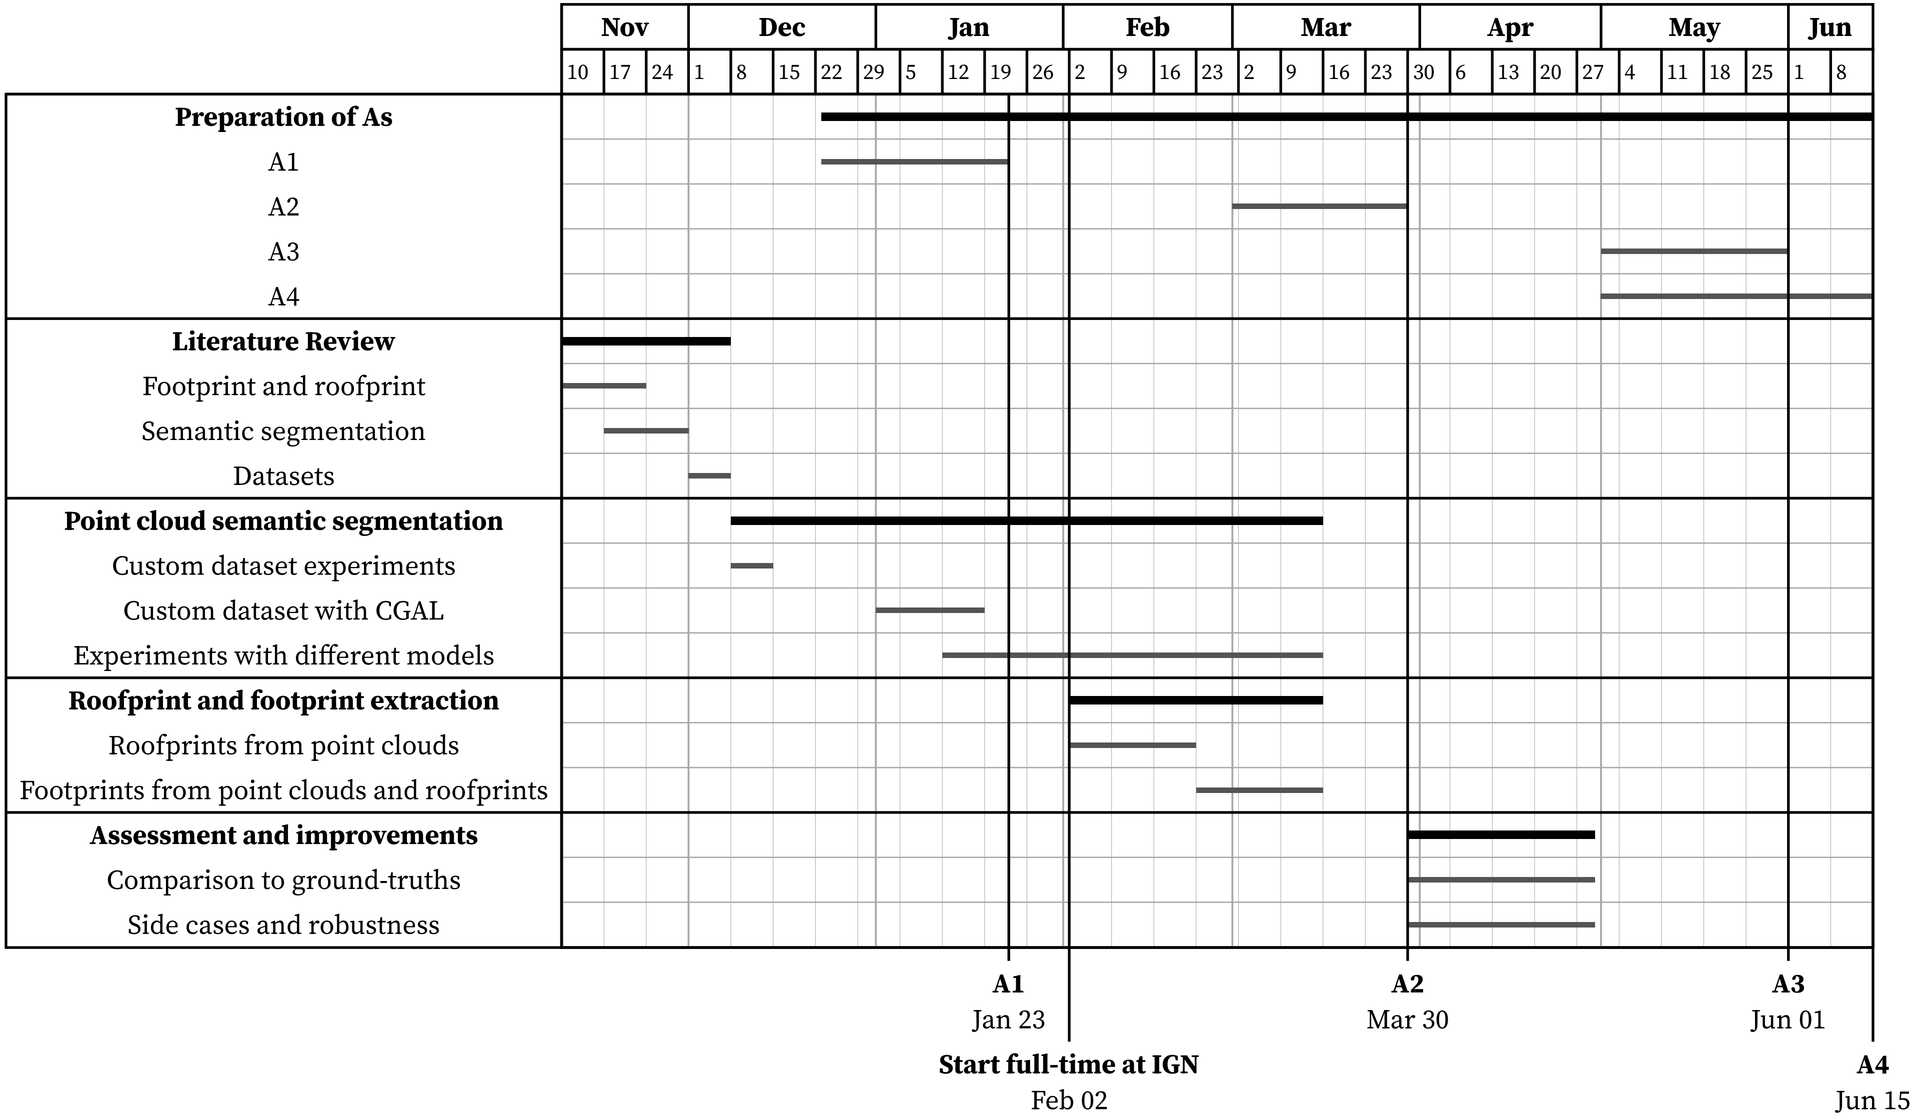

7. Tools and datasets used

The main datasets that will be used are:
- Point clouds: AHN for the Netherlands and LiDAR HD for France. The real target is LiDAR HD, but AHN will be used for training in combination with other Dutch datasets and potentially to assess the final method. I may also use other point cloud datasets if I find relevant ones with the necessary separation of façade and roof points, such as the Vaihingen dataset.
- 3D models of buildings: 3DBAG for the Netherlands and internal examples for France. These will be used mostly for the custom dataset and potentially as ground-truths for the roofprints and footprints if they can be extracted from them.

The main tools that will be used are:
- CGAL to implement geometric operations both to create the custom dataset and to perform the computation of the roofprints and footprints.
- PyTorch to train the semantic segmentation model. Libraries like Pointcept, openpoints or torch-points3d could be useful to use existing architectures and speed up the process.
- CloudCompare to visualise results and perform quick experiments.
- High-performance clusters to train the models: DAIC for TU Delft or Jean Zay in France.

References

Biasutti, P., Bugeau, A., Aujol, J.-F., & Brédif, M. (2019). *RIU-Net: Embarrassingly simple semantic segmentation of 3D LiDAR point cloud*. https://doi.org/10.48550/ARXIV.1905.08748

Biljecki, F., Ledoux, H., & Stoter, J. (2016). An improved LOD specification for 3D building models. *Computers, Environment and Urban Systems*, *59*, 25–37. https://doi.org/10.1016/j.compenvurbsys.2016.04.005

Cramer, M. (2010). The DGPF-Test on Digital Airborne Camera Evaluation Overview and Test Design. *Photogrammetrie - Fernerkundung - Geoinformation*, *2010*(2), 73–82. https://doi.org/10.1127/1432-8364/2010/0041

Dahlke, D., Linkiewicz, M., & Meissner, H. (2015). True 3D building reconstruction: façade, roof and overhang modelling from oblique and vertical aerial imagery. *International Journal of Image and Data Fusion*, *6*(4), 314–329. https://doi.org/10.1080/19479832.2015.1071287

Dai, H., Xu, J., Hu, X., Shu, Z., Ma, W., & Zhao, Z. (2025). Deep projective prediction of building facade footprints from ALS point cloud. *International Journal of Applied Earth Observation and Geoinformation*, *139*, 104448. https://doi.org/10.1016/j.jag.2025.104448

Frommholz, D., Linkiewicz, M., Meissner, H., & Dahlke, D. (2017). Reconstructing Buildings with Discontinuities and Roof Overhangs from Oblique Aerial Imagery. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 465–471. https://doi.org/10.5194/isprs-archives-xlii-1-w1-465-2017

Gaydon, C., Daab, M., & Roche, F. (2024). *FRACTAL: An Ultra-Large-Scale Aerial Lidar Dataset for 3D Semantic Segmentation of Diverse Landscapes*. https://doi.org/10.48550/ARXIV.2405.04634

Girard, N., Smirnov, D., Solomon, J., & Tarabalka, Y. (2020). *Polygonal Building Segmentation by Frame Field Learning*. https://doi.org/10.48550/ARXIV.2004.14875

Goebbels, S., & Pohle-Fröhlich, R. (2023). Automatic Reconstruction of Roof Overhangs for 3D City Models. *Proceedings of the 18th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications*, 145–152. https://doi.org/10.5220/0011604200003417

Hu, Q., Yang, B., Khalid, S., Xiao, W., Trigoni, N., & Markham, A. (2020). *Towards Semantic Segmentation of Urban-Scale 3D Point Clouds: A Dataset, Benchmarks and Challenges*. https://doi.org/10.48550/ARXIV.2009.03137

IGN. (2024b, March). *La recherche au défi du jumeau numérique de la France — ign.fr*.

IGN. (2024a, October). *Un appel à communs pour la conception du Jumeau numérique de la France et de ses territoires — ign.fr*.

Kölle, M., Laupheimer, D., Schmohl, S., Haala, N., Rottensteiner, F., Wegner, J. D., & Ledoux, H. (2021). The Hessigheim 3D (H3D) benchmark on semantic segmentation of high-resolution 3D point clouds and textured meshes from UAV LiDAR and Multi-View-Stereo. *ISPRS Open Journal of Photogrammetry and Remote Sensing*, *1*, 100001. https://doi.org/10.1016/j.ophoto.2021.100001

Li, M., Lin, S., Wang, Z., Shen, Y., Zhang, B., & Ma, L. (2024). Class-imbalanced semi-supervised learning for large-scale point cloud semantic segmentation via decoupling optimization. *Pattern Recognition*, *156*, 110701. https://doi.org/10.1016/j.patcog.2024.110701

Li, W., Wang, F.-D., & Xia, G.-S. (2020). A geometry-attentional network for ALS point cloud classification. *ISPRS Journal of Photogrammetry and Remote Sensing*, *164*, 26–40. https://doi.org/10.1016/j.isprsjprs.2020.03.016

Li, X., Li, C., Tong, Z., Lim, A., Yuan, J., Wu, Y., Tang, J., & Huang, R. (2020). Campus3D: A Photogrammetry Point Cloud Benchmark for Hierarchical Understanding of Outdoor Scene. *Proceedings of the 28th ACM International Conference on Multimedia*, 238–246. https://doi.org/10.1145/3394171.3413661

Mérizette, M., Audebert, N., Kervella, P., & Verdun, J. (2025). *3DSES: an indoor Lidar point cloud segmentation dataset with real and pseudo-labels from a 3D model*. https://doi.org/10.48550/ARXIV.2501.17534

Niemeyer, J., Rottensteiner, F., & Soergel, U. (2014). Contextual classification of lidar data and building object detection in urban areas. *ISPRS Journal of Photogrammetry and Remote Sensing*, *87*, 152–165. https://doi.org/10.1016/j.isprsjprs.2013.11.001

Pan, Z., Zhang, N., Gao, W., Liu, S., & Li, G. (2025). Point Cloud Semantic Segmentation with Sparse and Inhomogeneous Annotations. *Proceedings of the AAAI Conference on Artificial Intelligence*, *39*(6), 6354–6362. https://doi.org/10.1609/aaai.v39i6.32680

Panday, U. S., & Gerke, M. (2012). Fitting of Parametric Building Models to Oblique Aerial Images. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 233–238. https://doi.org/10.5194/isprsarchives-xxxviii-4-w19-233-2011

Peters, R., Dukai, B., Vitalis, S., van Liempt, J., & Stoter, J. (2022). Automated 3D reconstruction of LoD2 and LoD1 models for all 10 million buildings of the Netherlands. *Photogrammetric Engineering and Remote Sensing*, *88*(3), 165–170. https://doi.org/10.14358/PERS.21-00032R2

Saadaoui, H., Farah, S., Lechgar, H., Ghennioui, A., & Rhinane, H. (2025). Advancing Urban Roof Segmentation: Transformative Deep Learning Models from CNNs to Transformers for Scalable and Accurate Urban Imaging Solutions—A Case Study in Ben Guerir City, Morocco. *Technologies*, *13*(10), 452. https://doi.org/10.3390/technologies13100452

Thomas, H., Qi, C. R., Deschaud, J.-E., Marcotegui, B., Goulette, F., & Guibas, L. J. (2019). *KPConv: Flexible and Deformable Convolution for Point Clouds*. https://doi.org/10.48550/ARXIV.1904.08889

Varney, N., Asari, V. K., & Graehling, Q. (2020). *DALES: A Large-scale Aerial LiDAR Data Set for Semantic Segmentation*. https://doi.org/10.48550/ARXIV.2004.11985

Vijaywargiya, J., & Ramiya, A. M. (2025). Thiruvananthapuram Aerial LiDAR Dataset (TALD): A Benchmark for Complex Urban Point Cloud. *IEEE Access*, *13*, 42350–42363. https://doi.org/10.1109/access.2025.3546628

Wu, X., Jiang, L., Wang, P.-S., Liu, Z., Liu, X., Qiao, Y., Ouyang, W., He, T., & Zhao, H. (2023). *Point Transformer V3: Simpler, Faster, Stronger*. https://doi.org/10.48550/ARXIV.2312.10035

Xu, C., Wu, B., Wang, Z., Zhan, W., Vajda, P., Keutzer, K., & Tomizuka, M. (2020). *SqueezeSegV3: Spatially-Adaptive Convolution for Efficient Point-Cloud Segmentation*. https://doi.org/10.48550/ARXIV.2004.01803

Ye, Z., Xu, Y., Huang, R., Tong, X., Li, X., Liu, X., Luan, K., Hoegner, L., & Stilla, U. (2020). LASDU: A Large-Scale Aerial LiDAR Dataset for Semantic Labeling in Dense Urban Areas. *ISPRS International Journal of Geo-Information*, *9*(7), 450. https://doi.org/10.3390/ijgi9070450

Zachar, P., Bakuła, K., & Ostrowski, W. (2023). CENAGIS-ALS Benchmark - New Proposal for Dense ALS Benchmark Based on the Review of Datasets and Benchmarks for 3D Point Cloud Segmentation. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 227–234. https://doi.org/10.5194/isprs-archives-xlviii-1-w3-2023-227-2023

Zhao, H., Jiang, L., Jia, J., Torr, P., & Koltun, V. (2020). *Point Transformer*. https://doi.org/10.48550/ARXIV.2012.09164

Zolanvari, I., Ruano, S., Rana, A., Cummins, A., Smolic, A., Da Silva, R., & Rahbar, M. (2019, September). *DublinCity: Annotated LiDAR Point Cloud and its Applications*.

Glossary

Datasets

BAG (Basisregistratie Adressen en Gebouwen) — The Dutch dataset of buildings outlines with addresses and other attributes. More information at https://www.kadaster.nl/zakelijk/registraties/basisregistraties/bag. 1, 6

3DBAG — An open 3D building data set that is generated fully automatically based on LiDAR data and covers the whole of the Netherlands. More information at https://docs.3dbag.nl/en/. 1, 2, 6, 7, 8, 9, 12

AHN (Actueel Hoogtebestand Nederland) — A programme to produce high-density point clouds on the whole Dutch territory, with a new version every few years. More information at https://www.ahn.nl/. 1, 2, 6, 7, 8, 9, 12

BD TOPO — The French dataset that contains (among many other types of objects) the buildings outlines, combining footprints and roofprints. 2, 5, 6, 7

LiDAR HD — The first project to collect high-density point clouds on the whole territory of France (except Guyane). More information at https://www.ign.fr/institut/programme-lidar-hd-vers-une-nouvelle-cartographie-3d-du-territoire. 1, 2, 4, 6, 7, 12

Entities

TU Delft (Delft University of Technology) — The Dutch university in which this project was carried on. More information at https://www.tudelft.nl/en/. 1

IGN (Institut national de l'information géographique et forestière) — The reference public operator for geographic and forest information in France. More information at https://www.ign.fr/institut/identity-card. 1, 2, 4

Software

C++ — A high-level, general-purpose programming language, focussed on performance, efficiency, and flexibility. More information at https://en.wikipedia.org/wiki/C%2B%2B. 9, 16

CGAL (Computational Geometry Algorithms Library) — An open source software project that provides easy access to efficient and reliable geometric algorithms in the form of a C++ library. More information at https://www.cgal.org/. 9, 11, 12

PyTorch — An optimized tensor library for deep learning using GPUs and CPUs More information at https://pytorch.org/. 12

roofer — An algorithm to reconstruct a 3D building model from a point cloud and a 2D roofprint polygon in different LoDs up to 2.2. More information at https://github.com/3DBAG/roofer. 1, 2

HELIOS++ (Heidelberg LiDAR Operations Simulator ++) — A general-purpose Python package for simulation of terrestrial, mobile and airborne laser scanning surveys. More information at https://github.com/3dgeo-heidelberg/helios. 6

Vocabulary

footprint — The 2D outer boundary defined by the vertical projection of the outer walls/façades of a building. 1, 2, 3, 4, 5, 6, 11, 12, 16

roofprint — The 2D outer boundary defined by the vertical projection of the roof of a building. 1, 2, 3, 4, 5, 6, 11, 12, 16

outline — Used to talk about either a roofprint or a footprint or when they are mixed and the differentiation cannot be made. 1, 2, 3, 5, 6, 7, 16

LoD (Level of Detail) — A concept to differentiate multi-scale representations of semantic 3D city models. It is in practice principally used to indicate the geometric detail of a model, primarily of buildings. More information at https://3d.bk.tudelft.nl/lod/. 1, 2, 3, 6, 8, 9, 16

ALS (Airborne Laser Scanning) — Aircraft-mounted lidar that emits laser pulses toward the ground, capturing millions of georeferenced points for large-scale terrain and vegetation models. More information at https://en.wikipedia.org/wiki/Lidar. 1, 4, 5, 6

MLS (Mobile Laser Scanning) — Vehicle-mounted lidar that records dense point clouds while moving, used for detailed street-level 3-D mapping. More information at https://en.wikipedia.org/wiki/Lidar. 2

TLS (Terrestrial Laser Scanning) — Tripod-mounted static lidar that scans surroundings from fixed positions, ideal for high-resolution models of structures and heritage sites. More information at https://en.wikipedia.org/wiki/Lidar. 2

DSM (Digital Surface Model) — A gridded digital representation of the highest point at every cell location, including all natural and man-made features such as vegetation, buildings, and other above-ground objects. 3