



Vaasan yliopisto
UNIVERSITY OF VAASA

WAFA ZAINAB

Urban Digital Twin

An Open Data-Based Methodology for Automated Roof and Photovoltaic
Characterisation in 3D City Models

School of Technology and Innovations
Smart Grids
M.Sc. in Smart Cities and Communities

Vaasa 2026

UNIVERSITY OF VAASA**School of Technology and Innovations****Author:** Wafa Zainab**Title of the thesis:** Urban Digital Twin: An Open Data-Based Methodology for Automated Roof and Photovoltaic Characterisation in 3D City Models**Degree:** Master of Science**Degree Programme:** Smart Cities and Communities**Supervisor:** Professor Jon TERÉS ZUBIAGA, University of Basque Country, Spain**Year:** 2026 **Pages:** 73

ABSTRACT:

Urban Digital Twins (UDTs) have emerged as important tools for supporting sustainable urban planning, renewable energy integration, and smart city applications. However, generating geometrically accurate and semantically rich urban models remains challenging due to limitations in airborne LiDAR data and existing rooftop reconstruction workflows. This thesis presents a multi-stage geospatial methodology for developing an Urban Digital Twin of the Le Vigne district in Cesena, Italy, with a focus on rooftop reconstruction and photovoltaic suitability assessment. The proposed workflow integrates cadastral building footprints, institutional LiDAR datasets, DSM/DTM products, Google Earth Pro measurements, Google Photorealistic 3D Tiles, semantic segmentation, and automated reconstruction using Roofer. Geospatial harmonisation was performed in QGIS, while dense photogrammetric point clouds were generated and processed using Blender and CloudCompare. The resulting LoD2 building models were semantically enriched and exported in CityJSON format. The findings indicate that sparse airborne LiDAR data are effective for general urban modelling but insufficient for accurately reconstructing complex residential roofs. Dense photogrammetric reconstruction significantly improved roof geometry, ridge continuity, and roof-type representation. The final Urban Digital Twin incorporated semantic attributes such as roof slope, roof type, and photovoltaic suitability, enabling rooftop solar assessment and urban energy analysis. The study demonstrates the potential of combining LiDAR and photogrammetric data to create interoperable Urban Digital Twins that support data-driven urban planning and smart city decision-making.

KEYWORDS: Urban Digital Twin (UDT), LiDAR Analytics, 3D City Modelling, Semantic City Modelling, Rooftop Solar Assessment

Use of Artificial Intelligence

The author acknowledges the use of ChatGPT (OpenAI) as a supplementary tool during the preparation of this thesis. The tool was utilized for language editing, idea refinement, structuring content, and improving readability. The author takes full responsibility of the content of the thesis and confirms that all analysis, interpretations, results, and conclusions were developed independently by the author.

Contents

1	Introduction	5
1.1	Geospatial Databases, Data Sources, and Urban Informatics	6
1.2	Urban Digital Twins as a Framework for Smart-City Analysis	7
1.3	Role of Rooftop Photovoltaic Assessment within Urban Digital Twins	8
1.4	Challenges in Existing Urban Reconstruction Workflows	8
2	Objectives	11
2.1	Main Research Objective	11
2.2	Specific Objectives	11
2.3	Research Questions	12
3	State of the Art on Roof and Photovoltaic Detection	14
4	Methodology	21
4.1	Research Workflow	21
4.2	Proposed Methodology	23
4.2.1	Conceptual Framework	23
4.2.2	Methodological Inputs	24
4.2.3	Data Acquisition and Preparation	25
4.2.4	Data Integration and Interoperability	25
4.2.5	Methodological Outputs	26
5	Case Study	27
5.1	Data Sources	28
5.1.1	Cadastral Building Footprints	28
5.1.2	Institutional LiDAR Data	29
5.1.3	Google Photorealistic 3D Tiles	30
5.1.4	Google Earth Engine	30
5.2	Implementation of Proposed Methodology in Case study	31
5.2.1	Phase I – Manual Photogrammetric Roof Analysis	31
5.2.2	Phase II – Spatial Data Harmonisation in QGIS	32
5.2.3	Phase III – Automated Reconstruction Using Roofer	34

5.2.4	Phase IV – Slope Validation Using Google Earth Engine	35
5.2.5	Phase V – Dense Point Cloud Generation from Google 3D Tile	36
5.2.6	Mesh Acquisition in Blender using Blosm	37
5.2.7	Point Cloud Processing in CloudCompare, CANUPO	37
5.2.8	Phase VI – Dense Reconstruction Using Roofer	38
5.2.9	Phase VII – Semantic Enrichment and Classification	38
6	Key Considerations	40
6.1	Coordinate Reference System Management	40
6.2	Python Scripting for Geometry Attribute Computation	40
6.3	COMPOD Machine Learning Integration	41
6.4	Experimental Workflow	41
7	Results	43
8	Discussion	47
9	Conclusion and Future Work	50
	References	53
	Appendix	60

Figures

Figure 1: Google Earth Pro Geometry Extraction	32
Figure 2: QGIS Processing Pipeline	33
Figure 3: Roofer Processing and Ninja Viewer	35
Figure 4: Google Earth Engine Data Extraction	36
Figure 5: Mesh in Blender	37
Figure 6: Comparison of Results	41
Figure 7: Overall Workflow Architecture	42
Figure 8: Mesh Comparison	43
Figure 9: Google Earth Geometry Extraction	60
Figure 10: Google Earth Geometry Extraction	60
Figure 11: Length Calculation	61
Figure 12: Length Calculation	61
Figure 13: QGIS 2D Footprints	62
Figure 14: LIDAR Data	62
Figure 15: QGIS DTM	63
Figure 16: Importing National LIDAR in QGIS	63
Figure 17: Overlaying DTM on Regional Footprints	64
Figure 18: Downloading Pointcloud from PST Database	64
Figure 19: Overlaying of DSM on 2D Regional Footprints	65
Figure 20: Mesh Overlaying on Polygons	65
Figure 21: Classification of Pointcloud	65
Figure 22: Resulting 3D View	66
Figure 23: Slope Degree Estimation in QGIS	66
Figure 24: Roofer Commands to Generate 3D Model	67
Figure 25: 3D Model Rendering in Ninja	67
Figure 26: GEE Geometry Extraction	67
Figure 27: GEE Slope Calculation	68
Figure 28: JAVA Scrapping in GEE	68
Figure 29: Mesh Extraction via Blosm Plugin	69

Figure 30: Google 3D Tiles Mesh	69
Figure 31: Google 3D Tiles Mesh in Blender	69
Figure 32: Mesh Cleaning in CloudCompare	70
Figure 33: Clean Pointcloud Generation via CANUPO	70
Figure 34: 3D Model View in Autodesk Viewer	71
Figure 35: Classification of Sloped Roofs and Walls	71

Tables

Table 1: Classification and Semantic Enrichment	39
---	----

1 Introduction

Rapid urbanization, increasing energy demand, and the accelerating impacts of climate change have intensified the need for sustainable urban development strategies capable of supporting resilient and energy-efficient cities (United Nations, 2019). Contemporary urban environments are responsible for a substantial proportion of global energy consumption and greenhouse gas emissions, particularly through the operation of buildings and associated infrastructure systems (International Energy Agency [IEA], 2023). As cities continue to expand both spatially and demographically, the transition toward low-carbon urban energy systems has become a critical scientific, technological, and policy challenge (Intergovernmental Panel on Climate Change [IPCC], 2023). Within this context, rooftop photovoltaic (PV) deployment has emerged as one of the most promising distributed renewable energy solutions for reducing dependence on centralized fossil-fuel-based electricity generation while simultaneously improving urban energy resilience and decentralization. However, effective rooftop solar planning requires highly detailed geospatial intelligence capable of accurately representing building morphology, roof orientation, slope geometry, and surrounding urban context.

The growing complexity of urban systems has simultaneously accelerated the emergence of smart city paradigms centred around real-time urban analytics, interoperable geospatial infrastructures (Batty et al., 2012), and data-driven decision-making frameworks. Modern smart city initiatives increasingly rely on the integration of spatial data, sensor networks, remote sensing technologies, and computational modelling platforms to support sustainable urban governance. In this evolving landscape, geospatial technologies such as airborne LiDAR, photogrammetry, semantic 3D modelling, and Urban Digital Twins (UDTs) have become fundamental components of urban informatics research (Kitchin, 2014). These technologies enable the generation of high-resolution digital representations of cities capable of supporting applications including urban energy analysis, environmental monitoring, infrastructure management, disaster resilience assessment, and mobility optimization. Consequently, the quality, accuracy, and semantic richness of urban geospatial models directly influence the reliability of downstream analytical and simulation-based smart city applications.

Although smart-city initiatives and renewable energy technologies offer significant opportunities for sustainable urban development, several challenges continue to hinder their effective implementation. Reliable rooftop photovoltaic planning depends on accurate information regarding roof morphology, slope geometry, aspect orientation, usable rooftop area, and surrounding urban conditions (Deren et al., 2021). Even minor inaccuracies in these characteristics can significantly affect solar irradiance estimation, photovoltaic energy yield calculations, and installation feasibility assessments. Traditional two-dimensional geospatial representations are insufficient for capturing the geometric complexity of modern urban environments (White et al., 2021). Furthermore, urban areas are characterized by heterogeneous building morphologies, dense residential structures, vegetation interference, and complex rooftop configurations, making accurate rooftop characterization increasingly difficult.

Additional challenges arise from the need to integrate multiple datasets originating from different sources, resolutions, acquisition methods, and semantic standards. Misalignment between spatial datasets often introduces geometric inconsistencies and reduces the reliability of urban analyses. Consequently, urban planners and energy stakeholders require advanced methodologies capable of generating accurate, semantically meaningful, and interoperable representations of urban environments for sustainable decision-making.

1.1 Geospatial Databases, Data Sources, and Urban Informatics

As a result, the growing complexity of urban systems has accelerated the development of geospatial databases, urban informatics platforms, remote sensing technologies, and data-driven decision-support systems. Modern smart-city initiatives increasingly rely on the integration of spatial data, sensor networks, geospatial databases, and computational modelling frameworks to support sustainable urban governance. Advanced geospatial technologies such as airborne LiDAR, photogrammetry, semantic 3D modelling, Geographic Information Systems (GIS), and remote sensing have become fundamental components of urban informatics research (Jones et al., 2020). These technologies enable the generation of high-resolution digital representations of cities capable of supporting applications including urban energy analysis, environmental

monitoring, infrastructure management, disaster resilience assessment, and mobility optimization.

Airborne LiDAR point clouds provide direct elevation measurements and reliable terrain modelling capabilities, whereas photogrammetric reconstruction techniques generate dense surface representations through multi-view image processing. Recent developments, such as Google Photorealistic 3D Tiles, further enhance the availability of detailed urban surface information. Simultaneously, semantic modelling standards such as CityGML and CityJSON enable structured representation of buildings, rooftops, terrain, and infrastructure elements within machine-readable urban environments (Biljecki et al., 2016; Ledoux et al., 2019). As a result, the quality, accuracy, and semantic richness of urban geospatial datasets have become critical factors influencing the reliability of analytical and simulation-based urban applications.

1.2 Urban Digital Twins as a Framework for Smart-City Analysis

In this context, the increasing convergence of geospatial intelligence, urban informatics, and sustainable energy planning has accelerated the development of advanced three-dimensional urban modelling methodologies. Among these developments, Urban Digital Twins (UDTs) have emerged as critical computational frameworks for integrating heterogeneous geospatial datasets, semantic city models, remote sensing products, and analytical simulations within unified digital environments (Kolbe, 2009). Urban Digital Twins can be understood as semantically enriched digital replicas of urban environments that combine spatial geometry, semantic attributes, sensor data, and analytical models within interoperable computational ecosystems. Unlike traditional three-dimensional city models that primarily focus on visualization, UDTs are designed to support dynamic analysis, simulation, and decision-making processes.

The development of UDTs has been strongly supported by advancements in airborne LiDAR, photogrammetric reconstruction, Building Information Modelling (BIM), semantic city standards, and machine-learning-based geospatial analytics. These technologies enable urban models to evolve beyond static visualization and support applications such as energy simulation, environmental modelling, urban resilience assessment, infrastructure monitoring, and solar potential analysis (Ledoux et al., 2019;

Tao et al., 2019). However, the effectiveness of Urban Digital Twins depends heavily on the geometric fidelity and semantic completeness of reconstructed urban models. Existing workflows frequently struggle to achieve an optimal balance between geometric accuracy, semantic richness, interoperability, and computational efficiency. Consequently, many current UDT implementations still contain simplified roof geometries, incomplete semantic structures, and inconsistencies among heterogeneous datasets.

1.3 Role of Rooftop Photovoltaic Assessment within Urban Digital Twins

Furthermore, rooftop photovoltaic suitability assessment has gained increasing scientific and societal importance due to the urgent need for decentralized renewable energy integration within urban environments. Accurate rooftop solar assessment requires detailed information regarding roof slope, orientation, roof typology, shading behaviour, and available installation area (Haala & Brenner, 1999). Therefore, rooftop geometry reconstruction represents one of the most critical components of energy-oriented Urban Digital Twins. Airborne LiDAR datasets have historically served as primary sources for rooftop reconstruction because of their reliable elevation measurements. However, publicly available LiDAR datasets often exhibit sparse point distributions that limit the accurate reconstruction of complex residential rooftops. In contrast, photogrammetric reconstruction techniques can generate denser urban surface representations but may introduce reconstruction artefacts, image-matching instability, and uncertain geometric accuracy (Rottensteiner et al., 2005). Consequently, contemporary urban reconstruction research increasingly explores hybrid approaches that combine LiDAR, photogrammetry, semantic segmentation, and machine-learning techniques to improve both geometric quality and semantic completeness. Such approaches provide the foundation for more reliable photovoltaic suitability assessment and urban energy analysis within Urban Digital Twin environments.

1.4 Challenges in Existing Urban Reconstruction Workflows

However, despite substantial advances in remote sensing and urban reconstruction technologies, several limitations continue to affect existing workflows for semantically

enriched rooftop modelling and Urban Digital Twin generation. One of the most persistent challenges concerns the dependence on sparse airborne LiDAR datasets that are often incapable of accurately reconstructing complex residential roof structures. In many institutional datasets, point densities remain insufficient for resolving narrow ridges, dormers, intersecting roof planes, and architecturally heterogeneous rooftops, leading to simplified or flat-roof approximations during automated reconstruction.

Another major challenge concerns interoperability across heterogeneous geospatial datasets and software ecosystems. Urban reconstruction workflows frequently integrate data acquired from different spatial infrastructures, coordinate reference systems, resolutions, and semantic standards. Misalignment between cadastral footprints, LiDAR point clouds, DSM/DTM rasters, and photogrammetric datasets can introduce significant reconstruction errors and semantic inconsistencies (Brito et al., 2018). Consequently, spatial harmonisation and geospatial preprocessing constitute critical but often underestimated stages within urban modelling pipelines.

Existing workflows additionally suffer from limited semantic richness. Many reconstruction approaches focus primarily on geometric visualization without embedding semantically meaningful building attributes required for analytical urban applications. As a result, numerous three-dimensional city models remain unsuitable for energy-oriented analyses such as rooftop photovoltaic assessment, solar irradiance simulation, or infrastructure intelligence generation. Furthermore, fully automated reconstruction pipelines frequently struggle with vegetation occlusion, noisy point clouds, classification instability, and inconsistent roof-plane segmentation, particularly within dense residential neighbourhoods.

Manual reconstruction approaches, while capable of improving local geometric accuracy, are generally non-scalable and computationally inefficient for neighbourhood-scale or city-scale implementation. This creates an important methodological gap between high-fidelity but labour-intensive reconstruction methods and scalable but semantically simplified automated workflows. Addressing this gap requires reconstruction frameworks capable of integrating dense photogrammetric datasets, airborne LiDAR,

semantic segmentation techniques, and interoperable city modelling standards within unified computational pipelines (Mainzer et al., 2014).

2 Objectives

In response to these challenges, this research proposes a multi-stage Urban Digital Twin reconstruction workflow with the following:

2.1 Main Research Objective

The main objective of this research is to define an automatized Urban Digital Twin (UDT) workflow for enriching 3D city models with accurate roof slope geometries and photovoltaic (PV) panel representations using open-source geospatial data, with the aim of supporting energy-related urban analysis and semantic 3D city modelling within a CityGML 3.0 environment. The scope is designed to be applied to the Italian or EU context in general.

The research further seeks to evaluate the technical feasibility, geometric reliability, and practical limitations of integrating heterogeneous geospatial datasets, photogrammetric reconstruction techniques, and semantic urban modelling frameworks into a scalable neighbourhood-scale case study.

2.2 Specific Objectives

To achieve this overarching objective, a set of specific objectives has been defined, addressing the key stages of data acquisition, roof reconstruction, photovoltaic modelling, semantic integration, and workflow evaluation within the proposed Urban Digital Twin framework:

SO1: Roof Geometry Extraction

To develop a methodology for extracting accurate sloped roof geometries from Google Earth and complementary open geospatial datasets to enhance the geometric realism of Level of Detail (LOD) 2.0 3D building models.

SO2: Photovoltaic Panel Detection and Modelling

To define a procedure for detecting, modelling, and representing existing rooftop photovoltaic panels as geometric and semantic objects within a 3D urban modelling environment using Google Earth imagery and photogrammetric data sources.

SO3: Semantic Integration within 3DCityDB

To integrate roof slope attributes and photovoltaic panel information into 3DCityDB through a semantically structured database workflow aligned with CityGML 3.0 standards.

SO4: Generation of an Enhanced CityGML Model

To export an enriched CityGML 3.0 LOD 2.0 model demonstrating enhanced roof representation, semantic building attributes, and photovoltaic object integration suitable for Urban Digital Twin applications and energy-related urban analysis.

SO5: Case Study Implementation and Evaluation

To implement and evaluate the proposed workflow within the Le Vigne district in Cesena, Italy, to assess the robustness, interoperability, scalability, and limitations of the proposed methodology under real urban conditions.

SO6: Analysis of Available Geospatial Data Sources

To identify, analyse, and evaluate the suitability of available open-source and publicly accessible geospatial datasets, including Google Earth imagery, photogrammetric data, LiDAR data, cadastral information, and semantic city modelling resources, for supporting Urban Digital Twin development and rooftop photovoltaic assessment.

2.3 Research Questions

The research questions are directly linked to the specific objectives and are intended to guide the development, implementation, and evaluation of the proposed Urban Digital Twin workflow:

- Is it feasible to accurately extract sloped roof geometries from open-source geospatial data and Google Earth imagery for the enhancement of CityGML 3.0 LOD 2.0 building models?
- Can rooftop photovoltaic panels be reliably detected, geometrically modelled, and semantically represented within 3D city models using Google Earth imagery and photogrammetric reconstruction techniques?
- How effectively can enriched roof geometry and photovoltaic panel information be integrated into 3DCityDB and exported within an interoperable CityGML 3.0 framework?

- What are the principal technical limitations, geometric uncertainties, and interoperability challenges associated with using open-source and publicly accessible geospatial datasets for energy-related Urban Digital Twin applications?

Achieving these objectives requires a comprehensive understanding of the current state of knowledge, methodologies, technologies, and challenges related to Urban Digital Twins, semantic 3D city modelling, rooftop photovoltaic assessment, LiDAR and photogrammetric reconstruction, and geospatial interoperability standards. Therefore, the following chapter reviews the relevant scientific literature and existing approaches in these domains, identifying current achievements, limitations, and research gaps that motivate the proposed methodology of this study.

3 State of the Art on Roof and Photovoltaic Detection

Unlike conventional Geographic Information Systems (GIS), Urban Digital Twins are not merely spatial visualization platforms; rather, they function as dynamic, semantically enriched, multi-layered representations of urban environments capable of supporting simulation, prediction, and policy-oriented decision-making (Tao et al., 2019). Within this context, three-dimensional (3D) building reconstruction has become a central research problem due to its critical role in energy modelling, urban sustainability assessment, infrastructure management, and photovoltaic (PV) potential estimation.

Among the various applications of UDTs, rooftop solar suitability analysis has received considerable attention in recent years because accurate building-envelope reconstruction directly influences solar irradiance modelling and energy yield prediction (Bódis et al., 2021). However, existing literature demonstrates substantial inconsistencies regarding the optimal methodology for extracting reliable roof geometry information at neighbourhood and city scales. Current research remains fragmented between LiDAR-based reconstruction, photogrammetric mesh reconstruction, procedural modelling approaches, and increasingly, hybrid machine learning-driven pipelines. Consequently, the field continues to face unresolved challenges concerning geometric accuracy, scalability, semantic enrichment, and data interoperability.

Early research on 3D urban reconstruction relied predominantly on procedural extrusion techniques using 2D cadastral footprints combined with coarse elevation models. Although computationally efficient, these methods generated simplified Level of Detail (LoD1) representations incapable of accurately modelling roof typologies, pitch angles, or obstruction geometries (Kolbe et al., 2005). Biljecki et al. argued that such simplifications introduce substantial uncertainty into downstream urban simulations, particularly solar irradiation analysis, because generalized rooftops fail to preserve critical geometric attributes (Biljecki et al., 2016). Similarly, Kolbe emphasized that semantically poor building models significantly constrain the operational utility of digital twins, as geometry alone cannot support advanced urban analytics without semantic contextualization (Kolbe, 2009).

To overcome these limitations, airborne LiDAR rapidly emerged as the dominant technology for urban reconstruction due to its ability to capture direct elevation measurements with relatively high vertical precision. LiDAR-based workflows typically employ point cloud segmentation followed by planar fitting algorithms such as RANSAC, Hough Transform, or region-growing methods for roof plane extraction (Awrangjeb et al., 2011). Vosselman and Dijkman demonstrated that LiDAR-based segmentation performs effectively for regular roof structures characterized by large planar surfaces (Vosselman & Dijkman, 2001). However, subsequent studies identified severe limitations when these approaches are applied to dense urban environments containing occlusions, vegetation interference, and architecturally complex rooftops.

One of the most persistent challenges concerns point cloud density. Oude Elberink and Vosselman observed that sparse airborne LiDAR acquisitions frequently fail to capture fine-grained roof discontinuities, resulting in incomplete segmentation and inaccurate roof topology reconstruction (Oude Elberink & Vosselman, 2010). This issue becomes particularly problematic for photovoltaic suitability analysis because even small slope estimation errors may propagate into substantial deviations in solar irradiance calculations. In comparative studies, Yan et al. reported that sparse LiDAR datasets systematically underestimate roof slope values by overgeneralizing multi-faceted roofs into flat or simplified planar structures (Yan et al., 2021). Similar conclusions were reached by Nex and Remondino, who argued that conventional airborne LiDAR often lacks sufficient surface continuity for detailed urban reconstruction despite its strong metric reliability (Nex & Remondino, 2014).

These shortcomings have stimulated increasing interest in photogrammetric reconstruction techniques based on Structure-from-Motion (SfM) and Multi-View Stereo (MVS). Unlike LiDAR, photogrammetric pipelines reconstruct dense surface geometry from overlapping imagery, often generating point clouds with substantially higher spatial density. Research by Murtiyoso et al. demonstrated that dense image matching can produce centimetric surface detail suitable for heritage and urban reconstruction applications (Murtiyoso et al., 2017). More recent studies indicate that dense photogrammetric

meshes frequently outperform conventional LiDAR in rooftop-detail preservation, especially for ridge delineation and small architectural features (Zhang et al., 2013).

Nevertheless, photogrammetric reconstruction introduces a different class of limitations. Whereas LiDAR directly measures elevation, photogrammetric geometry is indirectly inferred through image correspondence estimation, making reconstruction quality highly sensitive to illumination conditions, surface texture, viewing geometry, and occlusions (Haala & Kada, 2010). Fan et al. demonstrated that urban photogrammetric meshes frequently contain surface noise, stitching artefacts, and geometric distortions near building edges (Fan et al., 2020). Moreover, the metric reliability of commercial photogrammetric products remains insufficiently documented because reconstruction algorithms and calibration parameters are often proprietary. Consequently, while photogrammetry offers superior surface density, concerns persist regarding geometric fidelity and scientific reproducibility.

The emergence of Google Photorealistic 3D Tiles has intensified scholarly interest in large-scale photogrammetric urban reconstruction. Generated through extensive multi-view stereo pipelines using aerial and satellite imagery, these datasets provide visually detailed city-scale meshes at resolutions unattainable through most publicly available LiDAR acquisitions (McClure et al., 2022). However, despite growing adoption within visualization platforms, relatively few studies have critically evaluated their scientific applicability for quantitative urban analysis. Existing research primarily focuses on rendering quality and visualization performance rather than metric reconstruction accuracy (Degan et al., 2021). Furthermore, most commercial photogrammetric meshes lack semantic segmentation, requiring substantial preprocessing before analytical use.

This limitation has contributed to growing research attention toward semantic point cloud classification. Traditional geometric segmentation approaches rely heavily on handcrafted heuristics, which often fail in heterogeneous urban scenes. Consequently, machine learning-based segmentation has become increasingly dominant. Weinmann et al. demonstrated that semantic classification using geometric descriptors significantly improves urban point cloud interpretation compared to rule-based segmentation (Weinmann et al., 2015). CANUPO, introduced by Brodu and Lague, represented a major

advancement using multiscale dimensionality analysis combined with Support Vector Machine (SVM) classification (Brodu & Lague, 2012). By computing local covariance-based descriptors such as linearity, planarity, and sphericity across multiple neighbourhood scales, CANUPO achieved strong discrimination between vegetation, ground surfaces, and buildings.

Despite these improvements, semantic segmentation research continues to face important unresolved challenges. Classification performance remains highly dependent on training sample selection and neighbourhood scale configuration (Griffiths & Boehm, 2019). In urban environments characterized by mixed vegetation and dense architectural complexity, overhanging trees and façade-shadow interactions frequently produce misclassifications. Moreover, most existing segmentation studies evaluate algorithms using terrestrial laser scanning datasets rather than photogrammetric urban meshes, leaving the robustness of these methods under photogrammetric noise conditions insufficiently investigated.

In parallel with advances in geometric reconstruction, the field has increasingly shifted toward semantic urban modelling frameworks such as CityGML and CityJSON. Unlike traditional mesh representations, semantic city models encode explicit relationships between urban objects and their attributes, thereby enabling interoperability across analytical platforms (Ledoux, 2018). Agugiaro argued that semantic enrichment constitutes a prerequisite for operational Urban Digital Twins because energy modelling, mobility simulation, and infrastructure analysis all require machine-readable building semantics rather than purely visual geometries (Agugiaro, 2016). However, maintaining semantic consistency during automated reconstruction remains a major challenge. Several studies report that automated pipelines frequently generate geometrically plausible models with inconsistent semantic labelling, particularly when roof planes are fragmented or partially reconstructed (Peters et al., 2022).

Recent literature has therefore explored the integration of machine learning into semantically enriched urban reconstruction pipelines. Deep learning frameworks based on PointNet++, KPConv, and graph neural networks have demonstrated promising performance in point cloud segmentation and urban object classification (Qi et al., 2017).

Nevertheless, these methods generally require extensive labelled datasets and computational resources, limiting their applicability in municipal-scale digital twin generation. Additionally, most deep learning approaches prioritize semantic classification accuracy without adequately addressing geometric reconstruction fidelity. Consequently, the integration of accurate roof geometry extraction with semantic intelligence remains insufficiently developed.

Within the context of solar energy modelling, the literature reveals similarly persistent methodological fragmentation. Rooftop PV estimation studies typically rely on one of three approaches: DSM-based irradiance modelling, LiDAR-driven roof segmentation, or image-based deep learning classification (Catita et al., 2017). DSM-based approaches are computationally scalable but suffer from coarse spatial resolution and poor roof-plane delineation. LiDAR-based approaches achieve stronger geometric precision but remain constrained by data accessibility and point density limitations (Hofierka & Kaňuk, 2009). Meanwhile, image-based deep learning methods can identify rooftop features effectively but often lack accurate elevation modelling necessary for reliable irradiance estimation.

Jakica observed that roof pitch and orientation constitute the most influential variables affecting rooftop PV estimation accuracy, yet these parameters are frequently oversimplified in large-scale studies (Jakica, 2016). Similarly, Mainzer et al. demonstrated that uncertainties in roof-plane extraction propagate directly into irradiance estimation errors, particularly in dense European urban environments with complex roof morphologies (Mainzer et al., 2017). Although several urban solar cadastre systems have been proposed in recent years, most continue to rely on simplified geometric assumptions that inadequately represent real-world roof complexity (Alobaidi et al., 2021).

A notable trend in recent research is the movement toward hybrid reconstruction architectures integrating multiple heterogeneous data sources. Rather than relying exclusively on LiDAR or photogrammetry, contemporary studies increasingly combine airborne point clouds, dense image matching, procedural modelling, and machine learning classification within unified workflows (Biljecki et al., 2015). This evolution reflects a broader recognition that no single reconstruction modality adequately satisfies all

requirements concerning accuracy, scalability, accessibility, and semantic richness. However, despite this transition, interoperability between heterogeneous geospatial tools remains poorly standardized. Coordinate reference inconsistencies, incompatible file formats, and varying semantic schemas continue to complicate end-to-end pipeline reproducibility (Ledoux et al., 2021).

Another unresolved issue concerns validation methodology. Many urban reconstruction studies evaluate performance primarily through visual inspection rather than rigorous geometric benchmarking. Quantitative validation using terrestrial laser scanning or high-precision reference datasets remains comparatively rare due to cost and data availability constraints (Kada & Wichmann, 2015). Consequently, reported reconstruction accuracies often lack sufficient external verification, particularly in studies utilizing proprietary photogrammetric datasets. This limitation becomes especially significant in photovoltaic applications where small geometric inaccuracies may substantially affect energy simulations.

Collectively, the existing literature reveals several critical research gaps directly relevant to this thesis. First, there remains no widely accepted methodology capable of simultaneously achieving high roof-geometry fidelity, semantic enrichment, automation scalability, and accessibility using open or reproducible data sources. Second, although dense photogrammetric urban meshes such as Google 3D Tiles provide unprecedented geometric detail, their integration into semantically structured reconstruction workflows remains largely underexplored. Third, most semantic segmentation studies focus either on LiDAR datasets or deep learning frameworks, with comparatively limited investigation into multiscale machine learning approaches for photogrammetric urban point clouds. Finally, existing rooftop PV assessment frameworks frequently depend on simplified geometric assumptions that inadequately capture roof complexity in dense urban environments.

These unresolved challenges collectively motivate the need for integrated reconstruction pipelines capable of combining dense urban geometry, semantic segmentation, machine learning classification, and interoperable digital twin modelling within a unified architecture. The broader technical evolution of the field suggests that future Urban

Digital Twin systems will increasingly depend on hybrid reconstruction methodologies that bridge the divide between geometric precision and semantic intelligence while remaining computationally scalable and operationally reproducible.

4 Methodology

This chapter presents the methodological framework developed to achieve the research objectives outlined in Chapter 3. The proposed methodology was designed to support the generation of a semantically enriched Urban Digital Twin (UDT) for rooftop photovoltaic assessment through the integration of heterogeneous geospatial datasets, including airborne LiDAR, photogrammetric data, cadastral information, and semantic city-modelling standards. The chapter first introduces the overall research workflow, describing the logical progression from data-source analysis to methodology design and case-study implementation. Subsequently, the proposed methodological framework, data acquisition procedures, interoperability strategies, and processing stages are presented in detail, followed by the implementation of the workflow and the generation of the final Urban Digital Twin outputs.

4.1 Research Workflow

The research workflow was designed in response to the limitations identified in the literature review. Existing rooftop reconstruction methodologies frequently rely on a single source of geospatial information, most commonly airborne LiDAR data, and therefore inherit the limitations associated with point density, acquisition geometry, vegetation occlusion, and incomplete roof representation. Similarly, many existing Urban Digital Twin implementations emphasize geometric visualization while providing only limited semantic information suitable for urban energy analysis. These limitations motivated the development of a multi-source reconstruction framework capable of combining airborne LiDAR data, photogrammetric datasets, semantic classification techniques, and interoperable city-modelling standards within a unified methodological architecture.

The workflow consists of three major research stages. The first stage focuses on the analysis and evaluation of available geospatial databases and data sources. The second stage concerns the design and implementation of the proposed reconstruction methodology. The third stage involves the application and validation of the methodology

within a real-world urban environment. These stages are interconnected and collectively contribute to the achievement of the research objectives defined in Chapter 3.

The first stage involved a systematic assessment of available geospatial resources capable of supporting rooftop reconstruction and Urban Digital Twin generation. This analysis was necessary because the quality of any reconstruction workflow is fundamentally constrained by the quality of the input data. Rather than selecting datasets solely based on availability, each candidate data source was evaluated according to its spatial resolution, geometric accuracy, semantic richness, coverage, accessibility, interoperability, and suitability for rooftop photovoltaic assessment. Attention was devoted to understanding the complementary strengths and weaknesses of airborne LiDAR datasets and photogrammetric urban meshes. While LiDAR datasets provide highly reliable elevation measurements and terrain characterization capabilities, they frequently exhibit insufficient sampling density over roof surfaces. Conversely, photogrammetric datasets provide dense surface continuity and visually complete roof structures but may contain geometric artefacts and uncertain elevation accuracy. This analysis established the rationale for integrating both data types within a hybrid reconstruction workflow.

The second stage focused on the design of the methodological framework itself. Rather than implementing a conventional linear reconstruction approach, the methodology was developed as a modular architecture in which intermediate outputs generated by one processing stage become inputs for subsequent stages. This design philosophy was adopted to improve workflow transparency, facilitate independent validation of individual components, and enhance reproducibility. The workflow integrates multiple software environments, including QGIS, Roofer, Blender, CloudCompare, Google Earth Engine, Google Earth Pro, and custom Python scripts. Each environment performs a specialized analytical function within the broader reconstruction process. QGIS serves as the primary platform for spatial preprocessing and harmonization; Roofer performs automated roof reconstruction; Blender and Blism facilitate access to Google Photorealistic 3D Tiles; CloudCompare supports dense point-cloud processing and

semantic filtering; Google Earth Engine provides independent slope validation; and Python scripts automate attribute calculations and data transformations.

The third stage consists of the practical implementation of the methodology within the selected case study area, namely the Le Vigne district of Cesena, Italy. The purpose of the implementation phase extends beyond simple demonstration of the workflow. It provides an opportunity to evaluate the feasibility, robustness, scalability, and interoperability of the proposed framework under realistic urban conditions. Furthermore, the case study enables comparative assessment between reconstruction results obtained from sparse airborne LiDAR data and those generated from dense photogrammetric point clouds, thereby allowing investigation of the extent to which dense photogrammetric reconstruction can improve rooftop representation for photovoltaic applications.

The overall workflow therefore represents a complete research pipeline beginning with data-source evaluation, progressing through methodological design and implementation, and culminating in the generation of a semantically enriched Urban Digital Twin. By structuring the research process in this manner, the study ensures methodological traceability while providing a reproducible framework that can potentially be adapted to other urban environments where high-quality rooftop information is required for energy-related analyses.

4.2 Proposed Methodology

The proposed methodology is based on the premise that accurate rooftop photovoltaic assessment requires both geometrically reliable roof reconstruction and semantically meaningful building representation.

4.2.1 Conceptual Framework

This research proposes a multi-stage Urban Digital Twin reconstruction framework that combines geospatial data acquisition, spatial harmonization, automated roof reconstruction, photogrammetric enhancement, semantic segmentation, and CityJSON-based semantic enrichment within a unified processing architecture. The framework was

designed according to four fundamental principles: geometric fidelity, semantic richness, interoperability, and reproducibility.

Geometric fidelity refers to the ability of the reconstructed model to accurately represent roof morphology, including slope, ridge configuration, roof typology, and aspect orientation. Semantic richness refers to the incorporation of building-level attributes capable of supporting analytical applications beyond visualization. Interoperability ensures compatibility between heterogeneous datasets and software ecosystems. Reproducibility guarantees that the workflow can be replicated using publicly accessible datasets and open-source software environments.

4.2.2 Methodological Inputs

The proposed methodology relies on multiple categories of geospatial information, each contributing distinct forms of spatial knowledge required for Urban Digital Twin generation.

Cadastral building footprints provide the geometric boundaries used for building identification, spatial clipping, and roof reconstruction. These footprints establish the fundamental spatial units upon which all subsequent analyses are performed.

Airborne LiDAR point clouds provide three-dimensional elevation measurements that support terrain modelling, roof-plane extraction, and building height estimation. Despite their importance, preliminary analysis revealed that the available LiDAR datasets exhibit limited roof-surface density in certain portions of the study area, motivating the exploration of complementary reconstruction approaches.

Digital Surface Models and Digital Terrain Models provide raster-based representations of above-ground and bare-earth elevation respectively. These datasets support slope analysis, terrain correction, and validation of reconstructed geometries.

Google Earth imagery and Google Photorealistic 3D Tiles provide dense photogrammetric representations of urban surfaces. These datasets play a crucial role in compensating for the geometric limitations of sparse LiDAR observations by providing continuous roof-surface coverage across the study area.

Finally, semantic city-modelling standards, particularly CityJSON, provide the data structure required for transforming reconstructed geometries into semantically enriched

Urban Digital Twin entities capable of supporting downstream urban-energy applications.

4.2.3 Data Acquisition and Preparation

The acquisition strategy adopted in this research was intentionally designed to combine authoritative institutional datasets with globally accessible photogrammetric resources. This approach reflects a central objective of the study: developing a workflow that remains transferable and reproducible even in locations where dense airborne LiDAR datasets may not be available.

Authoritative datasets were obtained from the Emilia-Romagna Regional Geoportal and include LiDAR point clouds, DSM products, DTM products, and cadastral building footprints. These datasets provide survey-grade spatial information characterized by well-documented acquisition procedures and known positional accuracy.

Photogrammetric datasets were acquired through Google Earth Pro and Google Photorealistic 3D Tiles. Unlike LiDAR datasets, which measure elevation directly, these resources are generated through large-scale multi-view stereo reconstruction processes and provide dense surface representations of urban environments. The integration of these datasets introduces an important methodological trade-off: increased geometric completeness at the cost of potentially greater geometric uncertainty. Rather than treating this as a limitation, the methodology explicitly investigates the potential benefits and constraints associated with this trade-off.

4.2.4 Data Integration and Interoperability

A central challenge addressed by this research concerns interoperability between heterogeneous datasets and software environments. The acquired datasets differ substantially in coordinate reference systems, geometric representations, semantic structures, storage formats, and acquisition epochs.

To overcome these challenges, all datasets underwent harmonization procedures involving coordinate transformation, spatial alignment, topology verification, geometric validation, and format standardization. Attention was devoted to CRS consistency

because inaccurate transformations can introduce geometric distortions that propagate throughout subsequent reconstruction stages.

Interoperability was further achieved using standardized exchange formats including KML, GeoJSON, LAZ, OBJ, and CityJSON. These formats enabled the transfer of information between GIS platforms, point-cloud processing environments, photogrammetric modelling software, and semantic city-modelling frameworks without significant loss of information.

4.2.5 Methodological Outputs

The final output of the proposed methodology is a semantically enriched Urban Digital Twin representing the study area at neighbourhood scale. Unlike conventional three-dimensional city models focused primarily on visualization, the generated Urban Digital Twin contains both geometric and semantic information required for photovoltaic-oriented urban analysis.

The resulting model incorporates reconstructed roof geometries together with attributes describing roof slope, roof typology, aspect orientation, building height, photovoltaic suitability indicators, and estimated rooftop energy potential. Intermediate outputs generated throughout the workflow include harmonized spatial datasets, filtered point clouds, segmented building surfaces, reconstructed roof models, validation datasets, and semantic enrichment layers.

Collectively, these outputs establish the foundation for rooftop photovoltaic assessment, Urban Digital Twin development, and future smart-city applications while simultaneously providing a reproducible methodological framework capable of being transferred to other urban environments.

5 Case Study

The study area selected for this research is the Le Vigne district located within the municipality of Cesena in the Emilia-Romagna region of northern Italy. Cesena represents a particularly suitable urban environment for investigating neighbourhood-scale Urban Digital Twin (UDT) reconstruction and rooftop photovoltaic (PV) suitability analysis due to its heterogeneous urban morphology, diverse building typologies, and availability of high-quality regional geospatial datasets.

Le Vigne is characterized by a mixed residential urban fabric composed primarily of detached and semi-detached residential buildings, low-rise apartment blocks, commercial structures, and service-related urban infrastructure. The district exhibits considerable variation in roof typologies, including flat, gabled, hipped, shed, and multi-faceted pitched roofs, thereby providing a representative environment for evaluating the robustness of automated roof reconstruction methodologies. This architectural diversity is particularly significant in the context of rooftop solar analysis because roof geometry directly influences solar irradiance interception, shading behaviour, panel orientation, and energy yield estimation.

From a geospatial modelling perspective, the study area presents several challenges commonly encountered in dense urban environments. These include partial roof occlusions caused by vegetation, heterogeneous building heights, narrow inter-building spacing, and complex roof discontinuities. Such conditions create substantial difficulties for automated plane-fitting algorithms and semantic segmentation techniques operating on sparse airborne LiDAR datasets. Consequently, the Le Vigne district constitutes an appropriate experimental environment for assessing the comparative performance of sparse LiDAR reconstruction and dense photogrammetric reconstruction workflows.

The district is also particularly relevant for urban-scale solar suitability analysis due to the prevalence of sloped residential rooftops with varying aspect orientations and pitch angles. In photovoltaic energy modelling, roof slope and orientation are among the most influential geometric variables affecting annual solar irradiance and potential energy production. The presence of diverse roof morphologies within the study area therefore

enables comprehensive evaluation of the proposed reconstruction pipeline under realistic urban energy-planning conditions.

In addition, the Emilia-Romagna region maintains one of the most advanced regional geospatial infrastructures in Italy through its institutional geoportal, providing open access to airborne LiDAR datasets, Digital Surface Models (DSM), Digital Terrain Models (DTM), cadastral data, and Web Map Services (WMS). The availability of these authoritative datasets significantly enhances the reproducibility and scientific validity of the proposed Urban Digital Twin workflow.

5.1 Data Sources

The reconstruction pipeline developed in this research integrates multiple heterogeneous geospatial datasets originating from both institutional and commercial sources. The selection of these datasets was guided by four principal criteria: geometric accuracy, spatial resolution, semantic relevance, and interoperability within Urban Digital Twin environments.

The adopted data sources collectively support the extraction, reconstruction, validation, and semantic enrichment of rooftop geometries required for photovoltaic suitability analysis. Because no single dataset independently satisfies all requirements related to roof reconstruction accuracy, semantic completeness, and neighbourhood-scale scalability, a multi-source integration strategy was implemented.

5.1.1 Cadastral Building Footprints

The foundational vector dataset used in this research consists of cadastral building footprint polygons acquired in KML and GeoJSON formats. These datasets provide two-dimensional building boundaries and serve as the geometric basis for three-dimensional roof reconstruction. The footprints were used to define building extents, constrain roof reconstruction, enable point cloud clipping, and establish building segmentation boundaries within the Roofer framework. Acting as geometric priors, they guide roof-plane extraction by limiting point cloud analysis to relevant building regions, making geometric accuracy essential to avoid reconstruction errors and semantic inconsistencies. Originally encoded in EPSG:4326, the datasets were reprojected to EPSG:25832 and

EPSG:32632 to ensure spatial consistency with LiDAR point clouds, DSMs, DTMs, and photogrammetric meshes. GeoJSON was selected as the primary workflow format due to its compatibility with Roofer, QGIS, and Python-based geospatial tools, while KML was additionally used for manual roof geometry extraction in Google Earth Pro. Although cadastral footprints provide accurate planimetric building delineation, they lack vertical geometric information and semantic roof attributes, requiring integration with elevation datasets such as LiDAR and photogrammetric meshes.

5.1.2 Institutional LiDAR Data

Institutional airborne LiDAR datasets were acquired from the Emilia-Romagna regional geoportal through the QU_LIDAR_RER2023_2024 acquisition campaign and included compressed LAZ point clouds, Digital Surface Models (DSM), and Digital Terrain Models (DTM), which served as the authoritative elevation source within the reconstruction pipeline. The airborne LiDAR point clouds, provided in EPSG:25832 using a 1×1 km tiling structure, contain three-dimensional laser return measurements and were selected for their high vertical accuracy, direct elevation measurement capability, relative independence from illumination conditions, and suitability for roof-plane extraction. As an active sensing technology, LiDAR enables reliable measurement of roof heights and terrain morphology and was therefore used as the primary source for automated roof modelling in Roofer.

However, the moderate point density of approximately $5\text{--}10$ points/m² limited its ability to capture small roof discontinuities, narrow ridges, and complex roof structures, resulting in incomplete segmentation and flat-roof approximations for some buildings. DSM datasets were used for elevation verification, roof-height assessment, and terrain analysis, while DTMs provided bare-earth elevation for terrain normalization, building-height estimation, and elevation referencing. The difference between DSM and DTM surfaces enabled the calculation of normalized building heights, which is important for shading analysis and solar potential assessment. Overall, although the institutional LiDAR datasets provide highly reliable elevation information, their limitations in representing complex rooftop geometries motivated the investigation of denser reconstruction alternatives.

5.1.3 Google Photorealistic 3D Tiles

To overcome the sparse sampling limitations of institutional LiDAR, this research employed Google Photorealistic 3D Tiles as an alternative dense urban reconstruction source accessed through the Blosm plugin in Blender. Generated through large-scale photogrammetric reconstruction using multi-view stereo processing of high-resolution aerial and satellite imagery, these datasets provide highly detailed textured meshes capable of representing fine-scale rooftop geometry. Their primary advantage lies in their high surface density and geometric continuity, allowing the visualization of complex roof morphologies, ridge lines, dormers, and roof discontinuities that are often unresolved in sparse LiDAR datasets. Consequently, the meshes were used for dense point cloud generation, high-resolution roof reconstruction, and improved roof-plane segmentation. The imported meshes were converted into dense point clouds using CloudCompare and subsequently processed through semantic segmentation and automated reconstruction with Roofer. However, because the reconstruction is image-based rather than measurement-based, the datasets may contain stitching artefacts, surface noise, geometric distortions, and local irregularities. Furthermore, the proprietary nature of Google's reconstruction pipeline limits transparency regarding acquisition and accuracy metrics, requiring careful validation. Another limitation is semantic ambiguity, as the meshes include buildings, vegetation, roads, terrain, and street furniture, necessitating CANUPO-based semantic segmentation to isolate building surfaces before reconstruction.

5.1.4 Google Earth Engine

Google Earth Engine (GEE) was employed as a cloud-based geospatial analysis platform for terrain slope validation and comparative roof geometry assessment. The platform provides access to large-scale remote sensing and elevation datasets such as Shuttle Radar Topography Mission (SRTM), regional DEM products, and terrain derivatives. Within this research, GEE was used to compute terrain slope values using the `ee.Terrain.slope()` function over manually defined building polygons. The resulting slope estimates were compared qualitatively against Roofer-derived roof slopes to assess geometric consistency.

The principal advantage of GEE lies in its scalable cloud-computing architecture, which enables rapid geospatial processing without local computational constraints. The platform also supports reproducible terrain analysis workflows and integration of multi-source elevation datasets.

Nevertheless, GEE-based slope estimation remains constrained by the spatial resolution of available DEM datasets. Satellite-derived terrain products generally lack the resolution required for precise rooftop geometry analysis, particularly for small residential structures with complex roof morphologies. Consequently, GEE was not employed as a primary reconstruction methodology but rather as an auxiliary validation framework for cross-checking slope estimation trends derived from higher-resolution reconstruction workflows.

5.2 Implementation of Proposed Methodology in Case study

Following are the phases which were implemented to test the proposed methodology:

5.2.1 Phase I – Manual Photogrammetric Roof Analysis

The initial phase of data collection employed Google Earth Pro as a photogrammetric measurement environment for extracting building geometry attributes within the Cesena study neighbourhood. This approach was motivated by the immediate availability of high-resolution oblique and nadir 3D imagery within Google Earth Pro's geospatial browser, which embeds photogrammetrically reconstructed 3D building meshes derived from multi-view stereo imagery. Using the software's integrated Ruler toolset, which supports 3D path and polygon measurement modes, the following geometric attributes were manually digitised for each building of interest: geographic coordinates (latitude, longitude) of key roof vertices; Z-values (absolute elevation in metres above mean sea level); eave height and ridge height, representing the lower and upper horizontal termini of a sloped roof face, respectively.

From the eave-to-ridge vertical differential (ΔZ) and the horizontal run distance derived from the 2D cadastral footprint polygon (sourced as a .KML file from the regional geportal), the roof slope angle θ was computed using the inverse tangent relationship: $\theta = \arctan (\Delta Z / \text{horizontal run})$. This computation was implemented in a purpose-written

Python script that performed a spatial join between the manually collected CSV of geometric attributes and the footprint centroids, enabling slope degree assignment on a per-building basis.

While this approach yielded geometrically coherent results for isolated buildings where clear ridge-to-eave measurement paths could be established, its fundamental limitation was its non-scalable, entirely manual character. The methodology required operator intervention for each individual building, precluding automation across the full study area. Furthermore, measurement accuracy was subject to the native resolution and positional accuracy of Google Earth Pro's 3D imagery, which in dense urban environments may exhibit mesh artefacts, occlusion-induced gaps, and vertical datum inconsistencies. This phase therefore served as a proof-of-concept and calibration baseline rather than as a production-grade data source.




Figure 1: Google Earth Pro Geometry Extraction

5.2.2 Phase II – Spatial Data Harmonisation in QGIS

To transition from manual measurement to a reproducible automated pipeline, the research subsequently focused on integrating institutional-grade geospatial data sourced from the official Emilia-Romagna regional geoportal (<https://geoportale.regione.emilia-romagna.it/>). The following data products were acquired: DSM tiles (first-return surface

model capturing both built structures and vegetation); DTM tiles (bare-earth terrain model produced through ground-point classification of the LiDAR data); and LiDAR point cloud tiles in LAZ format, representing the 2023-2024 airborne LiDAR acquisition campaign (QU_LIDAR_RER2023_2024) delivered on a 1x1 km grid in EPSG:25832 projection.

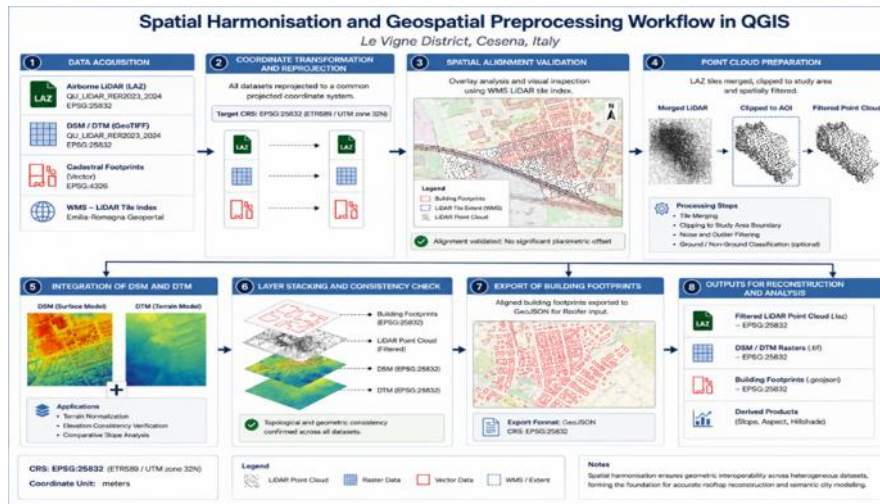


Figure 2: QGIS Processing Pipeline

Within QGIS 3.0, the 2D cadastral building footprints, originally encoded in an EPSG:4326 KML file were reprojected and co-registered with the LAZ point cloud tiles using spatial overlay operations. A critical pre-processing step involved validating the geometric alignment between the vector footprints and the point cloud extents through visual inspection using the WMS layer interface connecting to the regional LiDAR tile index (QU_LIDAR_RER2023_2024). This alignment was necessary because cadastral footprints may exhibit planimetric offsets relative to surveyed ground positions due to differing acquisition epochs, positional accuracy tolerances, and coordinate transformation residuals.

Following alignment verification, the spatially consistent footprint polygons were exported as a GeoJSON file, and the corresponding LiDAR tiles were merged and clipped to the study area extent to produce a study-area-specific LAZ file. Both outputs, the .laz point cloud and the .geojson footprint layer, were encoded in a consistent reference system (WGS 84, EPSG:4326) to satisfy the input requirements of the Roofer reconstruction engine.

5.2.3 Phase III – Automated Reconstruction Using Roofer

Roofer (<https://github.com/3DBAG/roofer>) is an open-source, command-line tool developed by the 3DBAG project at Delft University of Technology, designed to automatically reconstruct 3D building models from a combination of an airborne LiDAR point cloud and 2D building footprint polygons. Its reconstruction methodology is grounded in plane-fitting algorithms that identify dominant planar facets within the subset of point cloud returns overlying each footprint polygon. From these fitted planes, a geometric roof model is assembled that captures roof type (flat, shed, gabled, hipped, pyramidal, complex), slope angle, ridge orientation, and eave height.

Roofer was invoked via a Windows command-line environment, with the LAZ and GeoJSON files as positional arguments. The tool operated on 1,153 source footprints within the region of interest, identifying candidate reconstruction targets and producing a CityJSON output file containing the 3D building geometries enriched with per-building semantic attributes including: `building_id`, `height_m` (mean roof height above ground), `area_m2`, `roof_type` classification, `slope_deg` (estimated roof pitch), `pv_suitable` (boolean flag for photovoltaic suitability), and `ert_kwp` (estimated peak kilowatt capacity).

The primary limitation encountered in this phase was insufficient point cloud density over individual roof surfaces. The regional LiDAR dataset, acquired at an approximate pulse density of 5–10 pts/m², proved inadequate for Roofer's plane-fitting routines in cases of small-footprint buildings, complex multi-faceted roofs, or buildings whose roof surfaces were occluded by adjacent structures or vegetation at the time of acquisition. As a result, a significant proportion of reconstructed roofs exhibited partial geometries, flat roof approximations were returned for buildings known from visual inspection to possess sloped roofs, representing a systematic underestimation of slope angle and a consequent bias in PV suitability classification. The Roofer outputs were visually validated using the Ninja CityJSON viewer, which confirmed these incomplete reconstructions, necessitating an alternative dense point cloud generation strategy.

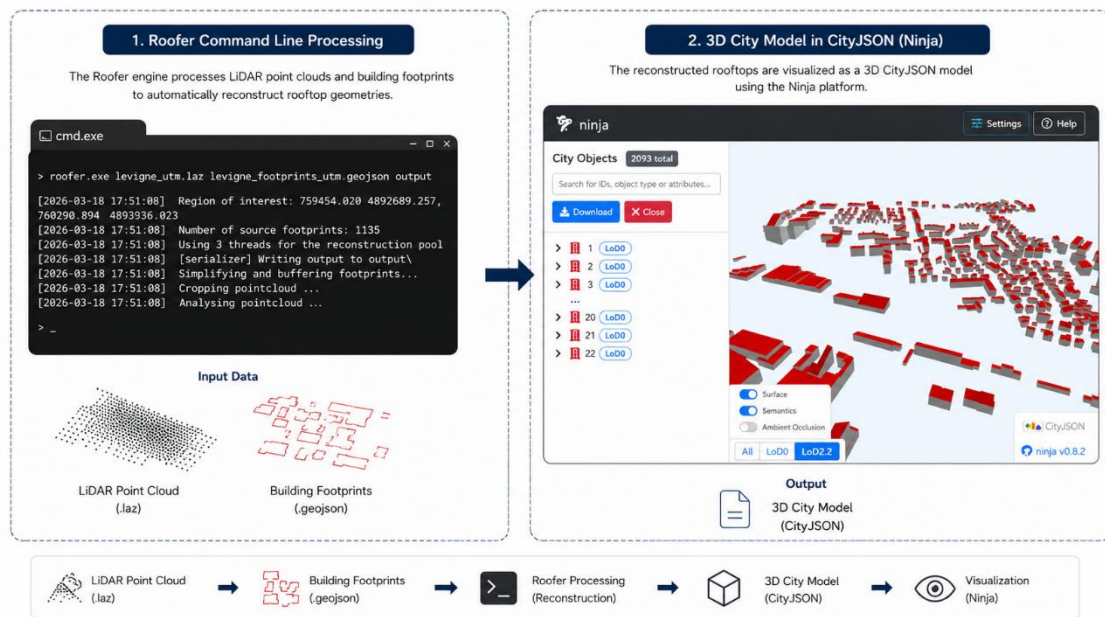


Figure 3: Roofer Processing and Ninja Viewer

5.2.4 Phase IV – Slope Validation Using Google Earth Engine

To provide an independent, pixel-based validation of roof slope angles, a secondary methodology was developed using the Google Earth Engine (GEE) JavaScript API. GEE provides cloud-based access to the NASA Shuttle Radar Topography Mission (SRTM) and regional elevation datasets, as well as the capability to compute terrain derivatives, including slope (in degrees) and aspect, from DSM imagery at sub-metre resolution via the `ee.Terrain.slope()` function.

For each building polygon of interest, a hand-digitised geometry was drawn interactively on the GEE code editor map interface. The `buildingAnalysis.reduceRegions()` function was then applied to compute statistical aggregations (mean, maximum) of the slope raster within each polygon boundary. The scale parameter was set to 5 metres to leverage higher-resolution regional elevation data where available, with fallback to 10-metre SRTM derivatives for buildings without corresponding high-resolution coverage.

While GEE provided a computationally tractable means of slope estimation for individual buildings, its application to the full study area was impractical due to the manual polygon digitisation requirement. Each building polygon had to be independently drawn, and slope statistics were computed per digitisation session rather than as a batch operation

over the entire footprint layer. This absence of batch automation over the full cadastral inventory precluded its adoption as a primary reconstruction methodology, relegating GEE to an auxiliary validation role for qualitatively cross-checking Roofer-derived slope values against independently derived terrain analysis estimates.

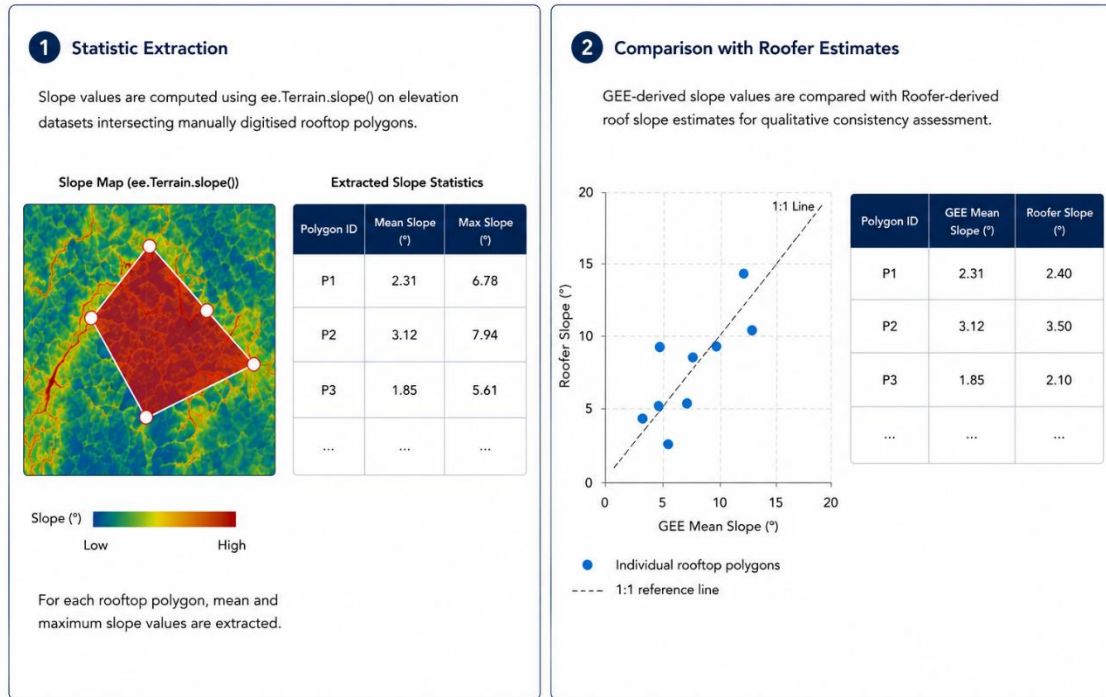


Figure 4: Google Earth Engine Data Extraction

5.2.5 Phase V – Dense Point Cloud Generation from Google 3D Tile

The fifth methodological phase constituted the principal innovation component of the workflow and focused on generating dense urban point clouds from Google Photorealistic 3D Tiles.

This phase was developed specifically to overcome the sparse sampling limitations observed during airborne LiDAR reconstruction. Google Photorealistic 3D Tiles are generated through large-scale multi-view stereo (MVS) photogrammetric reconstruction using aerial and satellite imagery. Unlike airborne LiDAR, these datasets provide highly detailed urban surface meshes with dense rooftop continuity. The workflow involved: mesh acquisition, mesh-to-point-cloud conversion, semantic segmentation, and point cloud refinement.

5.2.6 Mesh Acquisition in Blender using Blosm

The Blosm add-in for Blender 5.0 provides a graphical interface for streaming and importing Google 3D Tiles directly into a Blender scene as triangulated mesh objects. The study area extent was defined as the bounding box corresponding to the Cesena neighbourhood (Le Vigne district), and the Google 3D Tiles data was imported at the highest available level of detail (LOD). The resulting Blender scene contained a textured 3D mesh encompassing not only building structures but also terrain, vegetation canopies, street furniture, and other non-building urban features, collectively referred to as scene noise.



Figure 5: Mesh in Blender

5.2.7 Point Cloud Processing in CloudCompare, CANUPO

The raw Blender mesh was exported as an OBJ file and imported into CloudCompare 2.0, where it was first sampled into a dense point cloud representation. Point cloud subsampling was performed using an octree-based spatial subsampling strategy to achieve uniform spatial density, yielding a point cloud of several million points over the study area extent. To isolate building surfaces from non-building elements, specifically to remove ground, vegetation, and noise returns, the CANUPO (Context-Aware Classification from Normal and Orientation Point cloud Operations) machine-learning classifier plugin was employed.

CANUPO is a supervised classification framework that computes multi-scale geometric descriptors at each point (including dimensionality features derived from eigenvalue decomposition of local neighbourhood covariance matrices: linearity, planarity, sphericity, and verticality) and trains a Support Vector Machine (SVM) classifier using labelled training samples. Three binary classifiers were configured: (1) ground vs. non-ground, (2) vegetation vs. non-vegetation, and (3) building vs. non-building. By applying the classifiers sequentially and retaining only points classified as belonging to building surfaces, a semantically filtered point cloud containing exclusively roof and façade geometries was produced.

This cleaned point cloud was subsequently exported from CloudCompare as a LAZ file using custom LAS scale and offset parameters to preserve spatial precision, and served as the dense, high-fidelity point cloud input to the Roofer reconstruction engine in a second reconstruction pass.

5.2.8 Phase VI – Dense Reconstruction Using Roofer

With the Google 3D Tiles-derived dense point cloud substituting the sparse LiDAR input, the Roofer reconstruction pipeline was re-executed over the study area footprints. The substantially higher point density over roof surfaces enabled Roofer's plane-fitting algorithms to resolve fine-grained roof geometry features that had been unresolvable with the sparse regional LiDAR. The resulting CityJSON model, validated in both the Ninja CityJSON viewer and the Autodesk Viewer (AEC format), demonstrated significantly improved roof completeness, with the proportion of flat-roof approximations substantially reduced and the accuracy of slope angle estimates considerably improved.

5.2.9 Phase VII – Semantic Enrichment and Classification

The final methodological phase focused on semantic enrichment and Urban Digital Twin generation. The reconstructed CityJSON models were enriched with building-level semantic attributes including `slope_deg`, `roof_type`, `aspect`, and `ert_kwp`.

Table 1: Classification and Semantic Enrichment

Semantic Enrichment Component	Description / Function within Urban Digital Twin Framework
Purpose of Semantic Enrichment	Transformed reconstructed building geometry from purely visual 3D models into analytically interpretable Urban Digital Twin entities suitable for urban energy analysis.
Integrated Methodology	COMPOD-based machine learning classification was incorporated to automate semantic categorisation of reconstructed buildings.
Classification Parameters	Building geometries were semantically analysed using roof morphology, building height, aspect distribution, roof slope characteristics
Semantic Output	Generated semantically enriched building entities with rooftop energy intelligence attributes.
Integrated Urban Digital Twin Components	The final Urban Digital Twin framework combined: geometric reconstruction, semantic classification, rooftop energy intelligence
Data Framework / Interoperability	All enriched entities were organised within a unified interoperable CityJSON framework.
Application Domain	Suitable for neighbourhood-scale photovoltaic suitability assessment and smart city energy analysis.

6 Key Considerations

The proposed methodology was developed while keeping in view the following considerations including coordinate reference systems, python scripting and a machine learning algorithm.

6.1 Coordinate Reference System Management

A persistent challenge throughout the pipeline was the management of heterogeneous coordinate reference systems across data sources. The cadastral footprints were delivered in EPSG:4326 (geographic, WGS 84), the LiDAR tiles in EPSG:25832 (UTM Zone 32N, ETRS89), the DSM/DTM rasters in EPSG:25832, and the Google 3D Tiles mesh in a local Blender coordinate system anchored to a user-defined geographic origin. All spatial data operations requiring metric distance computation (slope angle, area, height) were performed in a projected UTM coordinate system (EPSG:25832 or EPSG:32632 depending on the tool), while file interchange formats (GeoJSON, KML) were encoded in EPSG:4326 to ensure cross-tool compatibility. Projection transformations were handled using QGIS's built-in reprojection algorithms (GDAL OGR) and custom Python scripts leveraging the pyproj library.

6.2 Python Scripting for Geometry Attribute Computation

A suite of Python scripts was developed to bridge tool-specific output formats and to compute derived geometric attributes not natively produced by any single tool. Key scripts included: a slope angle computation module that ingested the manually collected eave/ridge coordinate CSV alongside the cadastral KML footprints to calculate per-building slope degree via trigonometric analysis; a coordinate transformation script managing EPSG:4326 to EPSG:25832 reprojection for LiDAR alignment; and a mesh-to-point-cloud conversion script that processed Blender OBJ exports into CloudCompare-compatible formats using Open3D and laspy libraries.

6.3 COMPOD Machine Learning Integration

In the terminal phase of the pipeline, the COMPOD (COMplex Polygonal Object Detection) machine learning algorithm was applied to the enriched 3D building model outputs for automated classification of building attributes. COMPOD operates on geometric and contextual features extracted from CityJSON building objects, including roof slope distribution, building height, footprint compactness, and aspect orientation, to produce building-level classification outputs. The integration of COMPOD within the pipeline represents a transition from purely geometric reconstruction to semantically enriched building intelligence.

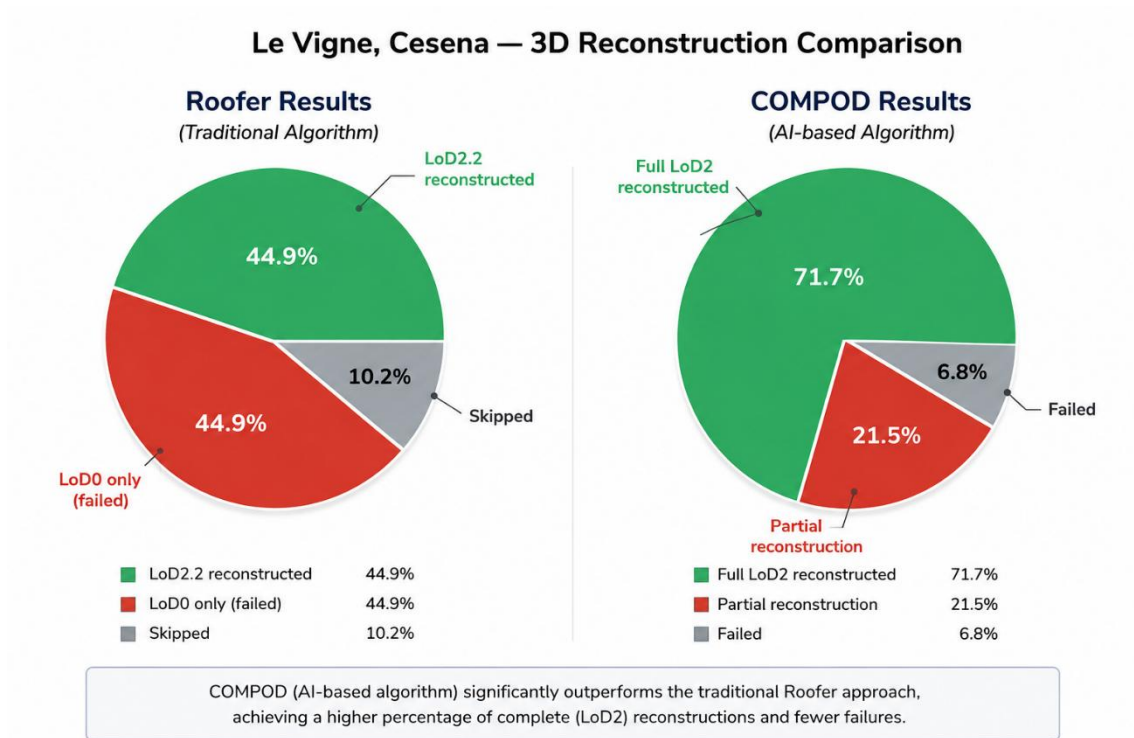


Figure 6: Comparison of Results

6.4 Experimental Workflow

Figure 7 presents the overall methodological framework developed for the generation of a semantically enriched Urban Digital Twin (UDT) of the Le Vigne district in Cesena, Italy. The workflow follows a multi-stage geospatial reconstruction pipeline integrating heterogeneous spatial datasets, airborne LiDAR point clouds, dense photogrammetric

meshes, semantic segmentation techniques, and automated roof reconstruction algorithms within a unified and interoperable architecture. The methodology was specifically designed to address limitations associated with sparse LiDAR-based roof reconstruction and the lack of semantically enriched 3D building models. Through progressive stages of data acquisition, point cloud preprocessing, semantic classification, rooftop reconstruction, validation, and CityJSON-based model generation, the workflow incrementally improves both geometric fidelity and semantic completeness, enabling the extraction of roof slope, roof typology, and aspect orientation for neighbourhood-scale urban energy analysis.

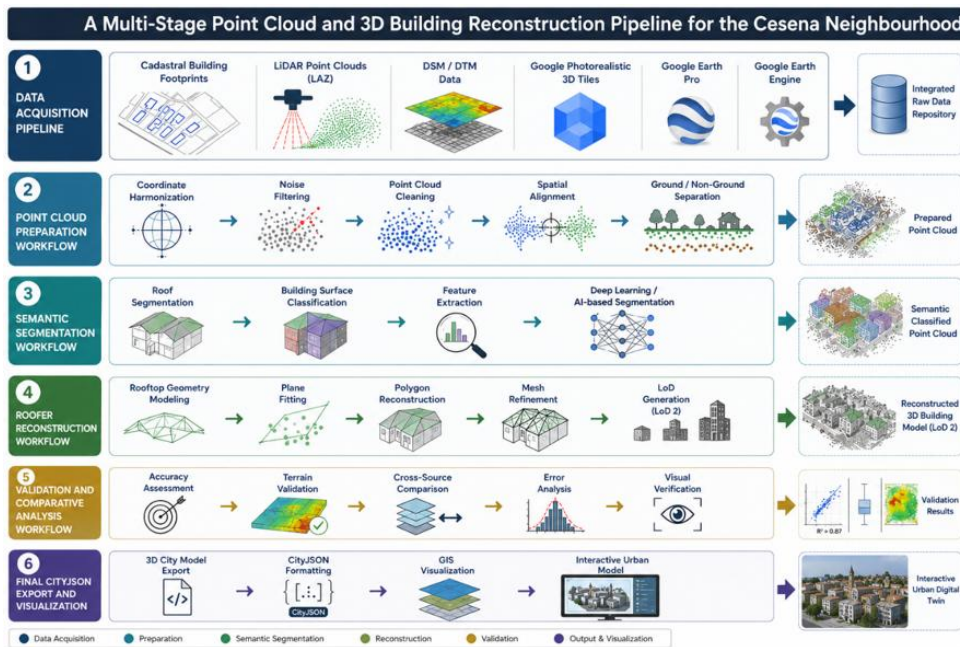


Figure 7: Overall Workflow Architecture

As illustrated in Figure 7, the proposed workflow follows a progressive refinement strategy in which geometric reconstruction accuracy and semantic completeness are incrementally enhanced through successive processing stages. The integration of sparse institutional LiDAR data with dense photogrammetric point clouds enables comparative evaluation of reconstruction methodologies while simultaneously improving roof-plane extraction reliability for complex urban morphologies. The final output of the pipeline consists of a semantically enriched CityJSON-based Urban Digital Twin containing interoperable building-level attributes suitable for neighbourhood-scale rooftop solar analysis, urban energy assessment, and future smart-city applications.

7 Results

The first reconstruction experiment employed the institutional airborne LiDAR dataset acquired from the Emilia-Romagna regional geoportal as the primary elevation source for automated roof reconstruction using Roofer.

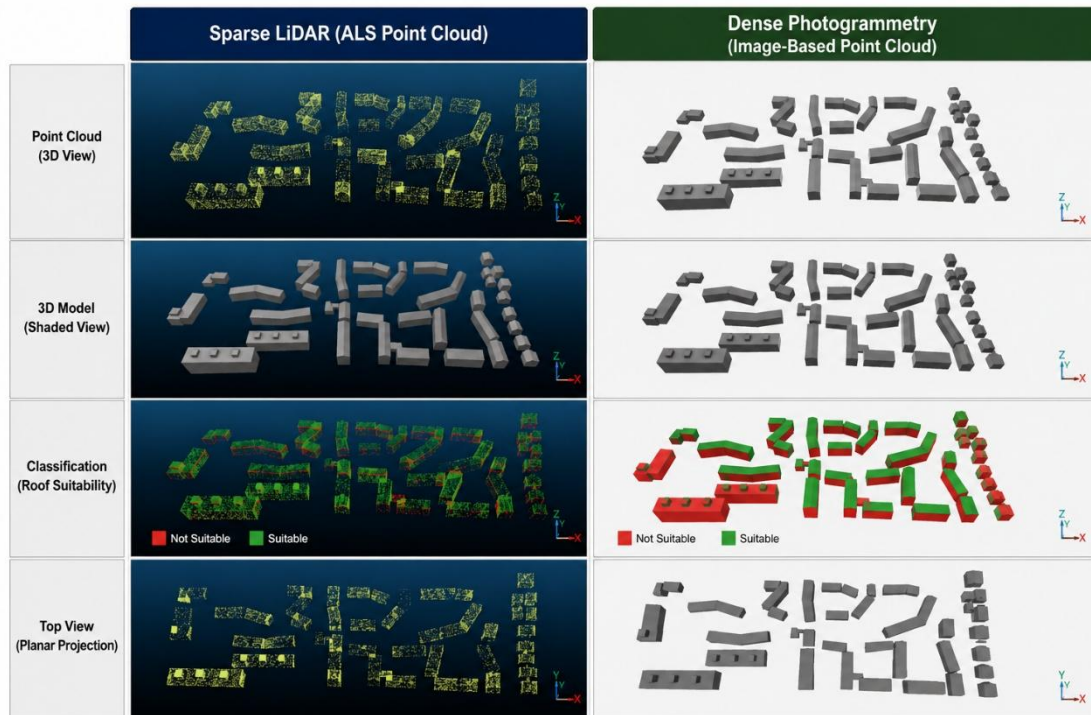


Figure 8: Mesh Comparison

The results demonstrated that sparse airborne LiDAR could support generalized LoD2 urban reconstruction; however, important geometric limitations emerged when modelling complex residential rooftops. The LiDAR dataset exhibited an approximate density of 5–10 pts/m², which proved sufficient for reconstructing large planar surfaces and overall building volumes but inadequate for resolving narrow ridges, intersecting roof planes, dormers, and multi-faceted roof geometries. Consequently, several buildings were reconstructed with simplified or flat roof approximations due to incomplete roof-plane segmentation.

These findings are consistent with previous LiDAR reconstruction studies reporting reduced plane-fitting reliability under sparse sampling conditions. Geometric incompleteness was particularly evident in areas affected by vegetation occlusion and

dense urban clustering, where reduced rooftop sampling negatively influenced ridge continuity and roof-edge delineation. Although the resulting CityJSON outputs preserved overall volumetric consistency, the geometric simplifications introduced uncertainty into roof slope estimation and photovoltaic suitability classification. Nevertheless, the sparse LiDAR workflow demonstrated strong vertical consistency and reliable terrain representation, indicating that airborne laser scanning remains suitable for generalized LoD2 urban modelling where high-resolution rooftop analysis is not required.

The dense reconstruction experiment employed point clouds generated from Google Photorealistic 3D Tiles following mesh extraction in Blender and semantic segmentation in CloudCompare. Compared with the sparse LiDAR workflow, the dense photogrammetric reconstruction produced substantially improved rooftop continuity, geometric completeness, and architectural detail. The higher point density enabled Roofer's plane-fitting algorithms to reconstruct complex roof geometries with greater stability. Multi-faceted residential roofs, gabled structures, and sloped roof planes that previously collapsed into flat approximations were reconstructed with improved geometric coherence. Ridge continuity and roof-edge definition were also significantly enhanced.

A major improvement concerned the reduction of flat-roof fallback classifications observed during the sparse LiDAR reconstruction. The dense photogrammetric point cloud enabled more accurate extraction of roof slope, roof orientation, roof typology, and usable roof area, all of which are critical variables for rooftop photovoltaic analysis. However, several limitations remained. Since Google 3D Tiles are derived from image-based reconstruction rather than direct elevation measurement, the generated meshes occasionally contained stitching artefacts, surface irregularities, and reconstruction noise. These distortions were most visible near vegetation boundaries and shadowed regions where photogrammetric surface matching becomes unstable. Despite these limitations, the dense reconstruction workflow produced substantially higher urban representation fidelity than the sparse LiDAR approach.

Comparative evaluation between sparse airborne LiDAR and dense photogrammetric reconstruction revealed significant differences in geometric quality, semantic

consistency, and rooftop analytical suitability. The most influential factor was point-cloud density. Sparse LiDAR provided reliable elevation measurements but insufficient rooftop sampling for stable roof-plane extraction. In contrast, the dense photogrammetric workflow significantly improved roof continuity, ridge preservation, roof-plane segmentation, and roof typology reconstruction. These findings support previous research indicating that dense photogrammetric point clouds often outperform sparse airborne LiDAR in residential rooftop reconstruction. However, important trade-offs were identified. Airborne LiDAR exhibited stronger metric reliability and reduced susceptibility to surface artefacts because laser scanning directly measures elevation independently of illumination conditions. Conversely, photogrammetric reconstruction provided greater geometric richness but introduced localized distortions caused by image matching and mesh stitching. The comparative analysis therefore highlights a balance between elevation reliability and geometric completeness. The integration of both datasets proved advantageous because LiDAR contributed vertical consistency while photogrammetric reconstruction enhanced rooftop detail and semantic reconstruction quality.

Semantic enrichment transformed the reconstructed geometries from visualization-oriented models into analytically usable Urban Digital Twin entities. The final CityJSON outputs contained building-level semantic attributes including roof slope degree (`slope_deg`), roof typology (`roof_type`), roof orientation, photovoltaic suitability (`pv_suitable`), and estimated peak photovoltaic capacity (`ert_kwp`). The dense photogrammetric workflow generated substantially more coherent semantic outputs because improved roof-plane segmentation reduced ambiguities in roof classification and slope estimation. The extracted slope values reflected realistic residential roof characteristics within the Le Vigne district, while aspect analysis revealed varying solar exposure conditions across the study area. Buildings exhibiting coherent roof planes, moderate roof slopes, and sufficient usable rooftop area were classified as suitable for photovoltaic installation. The semantic enrichment process therefore extended the contribution of the research beyond geometric reconstruction by producing a

semantically interpretable Urban Digital Twin suitable for urban energy analysis, rooftop solar assessment, and future smart-city applications.

Visualization quality also played an important role in assessing the operational usability of the reconstructed Urban Digital Twin. The final CityJSON outputs were visualized using Ninja CityJSON Viewer, QGIS, and Autodesk Viewer. The dense reconstruction workflow produced visibly improved architectural realism compared with the sparse LiDAR outputs, with roof structures exhibiting stronger geometric continuity, more coherent roof typologies, and improved semantic interpretability. CityJSON proved particularly effective for representing semantically enriched urban models because it preserved interoperability while maintaining lightweight data structures suitable for visualization and analysis. The resulting Urban Digital Twin successfully integrated heterogeneous geospatial datasets, semantic roof reconstruction, machine learning-based segmentation, and rooftop solar suitability analysis within a unified interoperable framework. Consequently, the generated model represents not only a visualization environment but also an analytical geospatial infrastructure capable of supporting future applications in urban energy planning, rooftop solar assessment, and sustainable smart-city development.

However, several residual sources of geometric error and processing uncertainty remain within the pipeline. First, the Google 3D Tiles mesh, while exhibiting high surface density, is itself a photogrammetric reconstruction artefact, small-scale surface irregularities, mesh stitching seams, and texturing artefacts may introduce localised noise into the derived point cloud that biases plane-fitting. Second, the CANUPO classification, while effective at separating ground and vegetation from building surfaces at neighbourhood scale, may misclassify overhanging trees adjacent to building rooftops as building points, introducing spurious points into the roof point cloud subset. Third, the absence of a dense, independent validation reference dataset, such as a high-accuracy terrestrial laser scanning survey of representative buildings, precludes rigorous quantitative accuracy assessment of slope angle estimates. The GEE-derived slope values were used as a qualitative cross-check only, given the coarser resolution of available satellite-derived elevation products in the study area.

8 Discussion

The research makes several important technical contributions to the fields of Urban Digital Twins (UDTs), geospatial reconstruction, and rooftop photovoltaic analysis through the development of a modular multi-stage reconstruction pipeline integrating heterogeneous geospatial datasets within a unified semantic framework. A primary contribution lies in the demonstrated integration of sparse institutional airborne LiDAR datasets with dense photogrammetric urban meshes derived from Google Photorealistic 3D Tiles for neighbourhood-scale roof reconstruction. Existing urban reconstruction workflows frequently rely exclusively on either airborne laser scanning or photogrammetric reconstruction, each of which presents inherent limitations related to point density, geometric reliability, or semantic completeness.

The proposed methodology demonstrates that combining these data modalities within a sequential reconstruction architecture substantially improves roof-plane continuity, ridge preservation, and rooftop semantic extraction. In particular, the integration of dense photogrammetric point clouds with Roofer reconstruction workflows enabled significantly improved extraction of roof slope, roof typology, and roof orientation attributes, which are critical variables for photovoltaic suitability assessment. The incorporation of CANUPO-based semantic segmentation further extends the contribution of the research by introducing machine learning-driven filtering of urban point clouds prior to reconstruction, thereby improving building-surface isolation within geometrically noisy photogrammetric environments. The resulting CityJSON-based Urban Digital Twin therefore represents not only a geometric reconstruction framework but also a semantically enriched analytical infrastructure capable of supporting energy-oriented urban applications.

Beyond the technical reconstruction workflow itself, the findings carry broader scientific implications for geospatial analytics, semantic city modelling, and smart-city research. The comparative evaluation between sparse airborne LiDAR and dense photogrammetric reconstruction highlights the increasing importance of hybrid geospatial reconstruction paradigms within contemporary urban informatics. The results demonstrate that point cloud density exerts a greater influence on roof-plane reconstruction

completeness than previously emphasized in many generalized LoD2 modelling studies. This has direct implications for urban energy modelling because rooftop photovoltaic suitability assessment depends heavily on accurate representation of roof morphology, aspect orientation, and slope geometry.

The research additionally contributes to the growing transition from visualization-oriented 3D city models toward semantically interpretable Urban Digital Twins capable of supporting analytical and simulation-based applications. By embedding semantic attributes such as `slope_deg`, `roof_type` within the CityJSON structure, the workflow demonstrates how semantic enrichment can transform reconstructed geometry into machine-readable urban intelligence suitable for integration into smart-city decision-support systems. Furthermore, the use of open geospatial standards and interoperable formats strengthens the methodological relevance of the research within broader GIS and digital-twin ecosystems where cross-platform compatibility and data exchange remain persistent challenges.

Despite the methodological advances achieved, several limitations remain that must be critically acknowledged. The first limitation concerns the sparse institutional LiDAR dataset itself, whose moderate point density constrained reliable reconstruction of complex residential rooftops. Although the integration of dense photogrammetric reconstruction partially addressed this issue, the sparse LiDAR workflow still produced incomplete roof-plane segmentation and flat-roof approximations for geometrically complex structures. A second limitation arises from the use of Google Photorealistic 3D Tiles as a dense reconstruction source.

While the photogrammetric meshes exhibited substantially improved rooftop continuity, they also introduced reconstruction artefacts associated with image-based multi-view stereo processing, including stitching irregularities, local surface distortions, and vegetation-induced noise. Because Google's reconstruction pipeline is proprietary, rigorous quantitative assessment of geometric accuracy was not feasible. Similarly, the semantic segmentation stage remained sensitive to classifier training quality and neighbourhood-scale parameterization, particularly in regions where vegetation and rooftop boundaries intersected.

Another important limitation concerns the absence of high-precision terrestrial laser scanning or ground-truth survey data for quantitative benchmarking of roof slope accuracy. Consequently, validation relied primarily on comparative reconstruction consistency, visual inspection, and supplementary terrain slope verification rather than absolute geometric error metrics. Nevertheless, these limitations do not diminish the methodological value of the research; rather, they reflect broader unresolved challenges currently affecting neighbourhood-scale Urban Digital Twin reconstruction within open and heterogeneous geospatial ecosystems.

9 Conclusion and Future Work

The research presented in this thesis developed and validated a multi-stage Urban Digital Twin (UDT) reconstruction pipeline for the Le Vigne district in Cesena, Italy, integrating heterogeneous geospatial datasets, airborne LiDAR point clouds, dense photogrammetric meshes, semantic segmentation techniques, and automated roof reconstruction algorithms within a unified semantic modelling framework. The findings demonstrated that sparse institutional airborne LiDAR datasets, while geometrically reliable for generalized LoD2 urban modelling, exhibit important limitations when reconstructing complex residential roof morphologies required for rooftop photovoltaic suitability assessment. In contrast, the integration of dense photogrammetric point clouds derived from Google Photorealistic 3D Tiles significantly improved roof-plane continuity, ridge preservation, and rooftop semantic extraction.

The comparative reconstruction experiments confirmed that point cloud density plays a decisive role in the accuracy and completeness of automated roof reconstruction workflows. Furthermore, the integration of semantic enrichment procedures enabled the generation of CityJSON-based Urban Digital Twin models containing analytically meaningful building attributes including roof slope, roof typology, aspect orientation, and photovoltaic suitability indicators. Collectively, the results demonstrate the feasibility of generating semantically enriched and geometrically coherent neighbourhood-scale Urban Digital Twins using interoperable geospatial processing pipelines built upon both institutional and publicly accessible datasets.

The research objectives defined at the outset of the study were successfully achieved through the implementation and evaluation of the proposed methodology. The first objective concerning the extraction of accurate sloped roof geometries from open geospatial datasets was addressed through both manual photogrammetric analysis and automated reconstruction using Roofer. The second objective related to rooftop semantic enrichment and photovoltaic suitability representation was achieved through the integration of roof slope, roof orientation, and energy-related semantic attributes within the final CityJSON model. The methodological objective of integrating heterogeneous datasets into an interoperable geospatial reconstruction framework was accomplished

through systematic spatial harmonisation, coordinate transformation, semantic segmentation, and automated reconstruction procedures implemented across QGIS, CloudCompare, Blender, Roofer, and Google Earth Engine environments.

Additionally, the implementation of the workflow within the Le Vigne district provided a realistic neighbourhood-scale validation environment through which the robustness, scalability, and practical applicability of the proposed Urban Digital Twin architecture could be critically evaluated. The research therefore addressed the identified problem of insufficient rooftop geometric fidelity within existing urban reconstruction workflows while simultaneously demonstrating the analytical value of semantically enriched 3D city models for rooftop solar assessment.

From a broader scientific perspective, this thesis contributes to the advancement of Urban Digital Twin research, LiDAR-based geospatial analytics, semantic city modelling, and smart-city geoinformatics. One of the principal contributions lies in demonstrating the operational integration of sparse airborne LiDAR datasets with dense photogrammetric urban meshes within a unified reconstruction framework. The research further contributes to the growing transition from visualization-oriented 3D city models toward semantically interpretable Urban Digital Twins capable of supporting analytical urban applications. By integrating machine learning-based semantic segmentation and CityJSON semantic enrichment into the reconstruction workflow, the study establishes a scalable methodological foundation for generating machine-readable urban intelligence suitable for rooftop solar analysis and future smart-city systems.

The findings additionally highlight the increasing relevance of hybrid reconstruction paradigms in contemporary geospatial research, where no single data modality independently satisfies the requirements for geometric accuracy, semantic completeness, and computational scalability. In this regard, the proposed workflow contributes to ongoing efforts within urban geospatial informatics to develop interoperable, semantically enriched, and analytically operational digital representations of urban environments.

Although the proposed methodology demonstrated strong performance within the neighbourhood-scale case study, several opportunities remain for future research and methodological refinement.

One important extension concerns the integration of deep learning-based rooftop object detection frameworks capable of automatically identifying existing photovoltaic installations directly from dense photogrammetric meshes or high-resolution aerial imagery. Such integration would significantly enhance the semantic richness of Urban Digital Twins by enabling direct modelling of rooftop solar infrastructure rather than relying solely on geometric suitability estimation.

Future research could additionally incorporate physically based solar irradiance simulation models to transform the current geometry-oriented workflow into a fully operational urban energy simulation environment. Another promising direction involves replacing conventional plane-fitting reconstruction algorithms with graph neural networks or transformer-based point cloud segmentation frameworks capable of handling highly irregular roof morphologies with greater robustness.

From a scalability perspective, future studies should investigate cloud-native geospatial processing architectures and distributed point cloud analytics capable of extending the proposed workflow from neighbourhood-scale implementation to city-scale or regional Urban Digital Twin generation. Such developments would support large-scale deployment of semantically enriched digital urban environments while maintaining computational efficiency and interoperability.

Finally, rigorous quantitative validation using terrestrial laser scanning or UAV-based photogrammetric surveys would provide more robust geometric benchmarking and support systematic evaluation of reconstruction uncertainty. Such developments would further strengthen the role of semantically enriched Urban Digital Twins as operational infrastructures for sustainable urban planning, rooftop solar deployment, and data-driven smart-city governance.

References

United Nations. (2019). World urbanization prospects 2018: Highlights. United Nations Department of Economic and Social Affairs.

International Energy Agency. (2023). Energy efficiency 2023. IEA Publications.

Intergovernmental Panel on Climate Change. (2023). Climate change 2023: Synthesis report. Contribution of Working Groups I, II and III to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change. IPCC.

Batty, M., Axhausen, K. W., Giannotti, F., Pozdnoukhov, A., Bazzani, A., Wachowicz, M., Ouzounis, G., & Portugali, Y. (2012). Smart cities of the future. *European Physical Journal Special Topics*, 214(1), 481–518. <https://doi.org/10.1140/epjst/e2012-01703-3>.

Kitchin, R. (2014). The real-time city? Big data and smart urbanism. *GeoJournal*, 79(1), 1–14. <https://doi.org/10.1007/s10708-013-9516-8>.

Deren, L., Wenbo, Y., & Zhenfeng, S. (2021). Smart city based on digital twins. *Computers, Environment and Urban Systems*, 85, 101531. <https://doi.org/10.1016/j.compenvurb-sys.2020.101531>.

White, G., Zink, A., Codecá, L., & Clarke, S. (2021). A digital twin smart city for citizen feedback. *Cities*, 110, 103064. <https://doi.org/10.1016/j.cities.2020.103064>.

Jones, D., Snider, C., Nassehi, A., Yon, J., & Hicks, B. (2020). Characterising the digital twin: A systematic literature review. *CIRP Journal of Manufacturing Science and Technology*, 29, 36–52. <https://doi.org/10.1016/j.cirpj.2020.02.002>.

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>.

Kolbe, T. H. (2009). Representing and exchanging 3D city models with CityGML. In J. Lee & S. Zlatanova (Eds.), *3D geo-information sciences* (pp. 15–31). Springer. https://doi.org/10.1007/978-3-540-87395-2_2.

Ledoux, H., Arroyo Ogori, K., Kumar, K., Dukai, B., Labetski, S., & Vitalis, G. (2019). CityJSON: A compact and easy-to-use encoding of the CityGML data model. *Open Geospatial Data, Software and Standards*, 4(4). <https://doi.org/10.1186/s40965-019-0064-0>.

Haala, N., & Brenner, C. (1999). Extraction of buildings and trees in urban environments. *ISPRS Journal of Photogrammetry and Remote Sensing*, 54(2–3), 130–137. [https://doi.org/10.1016/S0924-2716\(99\)00010-6](https://doi.org/10.1016/S0924-2716(99)00010-6).

Rottensteiner, F., Trinder, J., Clode, S., & Kubik, K. (2005). Building detection using LiDAR data and aerial imagery. *ISPRS Journal of Photogrammetry and Remote Sensing*, 59(1–2), 130–151. <https://doi.org/10.1016/j.isprsjprs.2005.02.006>.

Brito, M. C., Redweik, P., Catita, C., & Freitas, S. (2018). Solar potential in a Lisbon suburb using LiDAR data. *Solar Energy*, 159, 773–784. <https://doi.org/10.1016/j.solener.2017.11.007>.

Mainzer, K., Fath, K., McKenna, R., Stengel, J., Fichtner, W., & Schultmann, F. (2014). A high-resolution determination of the technical potential for residential-roof-mounted photovoltaic systems in Germany. *Solar Energy*, 105, 715–731. <https://doi.org/10.1016/j.solener.2014.04.015>.

Tao, F., Zhang, H., Liu, A., & Nee, A. Y. C. (2019). Digital twin in industry: State-of-the-art. *IEEE Transactions on Industrial Informatics*, 15(4), 2405–2415.

Bódis, C., Kougias, P., & Jäger-Waldau, A. (2021). High-resolution geospatial assessment of rooftop solar photovoltaics potential. *Applied Energy*, 284, 116434.

Kolbe, T. H., Gröger, G., & Plümer, L. (2005). CityGML: Interoperable access to 3D city models. In *Geo-information for disaster management* (pp. 883–899). Springer.

Biljecki, F., Ledoux, H., & Stoter, J. (2016). An improved LOD specification for 3D building models. *Computers, Environment and Urban Systems*, 59, 25–37.

Kolbe, T. H. (2009). Representing and exchanging 3D city models with CityGML. *Lecture Notes in Geoinformation and Cartography*, 15–31.

Awrangjeb, M., Ravanbakhsh, M., & Fraser, C. (2011). Automatic detection of residential buildings using LiDAR data and multispectral imagery. *ISPRS Journal of Photogrammetry and Remote Sensing*, 66(4), 457–467.

Awrangjeb, M., Ravanbakhsh, M., & Fraser, C. S. (2011). Automatic detection of residential buildings using LiDAR data and multispectral imagery. *ISPRS Journal of Photogrammetry and Remote Sensing*, 66(4), 457–467. <https://doi.org/10.1016/j.isprsjprs.2011.01.006>.

Vosselman, G., & Dijkman, S. (2001). 3D building model reconstruction from point clouds and ground plans. *ISPRS Journal of Photogrammetry and Remote Sensing*, 53(1), 37–43.

Oude Elberink, S., & Vosselman, G. (2010). Building reconstruction by target based graph matching on incomplete laser data. *ISPRS Journal of Photogrammetry and Remote Sensing*, 65(2), 149–160. <https://doi.org/10.1016/j.isprsjprs.2009.10.002>.

Yan, J., Wang, M., & Zhang, Y. (2021). Robust roof plane extraction from airborne LiDAR point clouds using adaptive segmentation. *Remote Sensing*, 13(9), 1775.

Nex, F., & Remondino, F. (2014). UAV for 3D mapping applications: A review. *Applied Geomatics*, 6(1), 1–15. <https://doi.org/10.1007/s12518-013-0120-x>.

Murtiyoso, A., Grussenmeyer, P., & Landes, T. (2017). Dense point clouds from image matching for heritage reconstruction. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XLII-2/W5, 535–542. <https://doi.org/10.5194/isprs-archives-XLII-2-W5-535-2017>.

Zhang, J., Lin, X., & Ning, X. (2013). SVM-based classification of segmented airborne LiDAR point clouds in urban areas. *Remote Sensing*, 5(8), 3749–3775. <https://doi.org/10.3390/rs5083749>.

Haala, N., & Kada, M. (2010). An update on automatic 3D building reconstruction. *ISPRS Journal of Photogrammetry and Remote Sensing*, 65(6), 570–580. <https://doi.org/10.1016/j.isprsjprs.2010.09.006>.

Fan, H., Zipf, A., Fu, Q., & Neis, P. (2020). Quality assessment of photogrammetric mesh models for urban applications. *Remote Sensing*, 12(14), 2215. <https://doi.org/10.3390/rs12142215>.

McClure, R., Hill, J., & Bogen, P. (2022). Photorealistic 3D tiles for geospatial visualization and analysis. In *Proceedings of the ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems* (pp. 45–54).

Degan, S., Remondino, F., & Nex, F. (2021). Large-scale urban mesh reconstruction from aerial imagery. *Remote Sensing*, 13(18), 3675. <https://doi.org/10.3390/rs13183675>.

Weinmann, M., Jutzi, B., & Mallet, C. (2015). Semantic point cloud interpretation based on optimal neighborhoods, relevant features and efficient classifiers. *ISPRS Journal of Photogrammetry and Remote Sensing*, 105, 286–304. <https://doi.org/10.1016/j.isprsjprs.2015.01.016>.

Brodu, N., & Lague, D. (2012). 3D terrestrial LiDAR data classification of complex natural scenes using a multi-scale dimensionality criterion. *ISPRS Journal of Photogrammetry and Remote Sensing*, 68, 121–134. <https://doi.org/10.1016/j.isprsjprs.2012.01.006>.

Griffiths, D., & Boehm, J. (2019). A review on deep learning techniques for 3D sensed data classification. *Remote Sensing*, 11(12), 1499. <https://doi.org/10.3390/rs11121499>.

Ledoux, H. (2018). CityJSON: A compact and easy-to-use encoding of the CityGML data model. *Open Geospatial Data, Software and Standards*, 3(1), 1–12. <https://doi.org/10.1186/s40965-018-0048-x>.

Agugiaro, G. (2016). Energy planning tools and CityGML-based 3D virtual city models. *Applied Geomatics*, 8(1), 41–56. <https://doi.org/10.1007/s12518-015-0152-5>.

Peters, R., Ledoux, H., & Stoter, J. (2022). Automatic reconstruction of 3D building models using CityJSON and airborne LiDAR. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, V-4-2022, 101–108. <https://doi.org/10.5194/isprs-annals-V-4-2022-101-2022>.

Qi, C. R., Yi, L., Su, H., & Guibas, L. J. (2017). PointNet++: Deep hierarchical feature learning on point sets in a metric space. In *Advances in Neural Information Processing Systems* (Vol. 30, pp. 5099–5108).

Catita, C., Redweik, P., Pereira, M., & Brito, M. C. (2017). Extending solar potential analysis in buildings to vertical facades. *Computers, Environment and Urban Systems*, 66, 1–11. <https://doi.org/10.1016/j.compenvurbsys.2017.08.002>.

Hofierka, J., & Kaňuk, J. (2009). Assessment of photovoltaic potential in urban areas using open-source solar radiation tools. *Renewable Energy*, 34(10), 2206–2214. <https://doi.org/10.1016/j.renene.2009.02.021>.

Jakica, N. (2016). State-of-the-art review of solar design tools and methods for assessing rooftop solar potential. *Renewable and Sustainable Energy Reviews*, 63, 810–820. <https://doi.org/10.1016/j.rser.2016.05.080>.

Mainzer, K., Killinger, S., McKenna, R., & Fichtner, W. (2017). Assessment of rooftop photovoltaic potentials at the urban level using publicly available geodata and image recognition techniques. *Solar Energy*, 155, 561–573. <https://doi.org/10.1016/j.solener.2017.06.065>.

Alobaidi, M., Al-Ansari, N., & Knutsson, S. (2021). Solar potential mapping using Google Earth Engine and high-resolution DEMs. *Energy Reports*, 7, 5890–5902. <https://doi.org/10.1016/j.egy.2021.08.090>.

Biljecki, F., Stoter, J., Ledoux, H., Zlatanova, S., & Çöltekin, A. (2015). Applications of 3D city models: State of the art review. *ISPRS International Journal of Geo-Information*, 4(4), 2842–2889. <https://doi.org/10.3390/ijgi4042842>.

Ledoux, H., Biljecki, F., & Stoter, J. (2021). 3DBAG: Generation of a national level of detail 2 dataset for the Netherlands. *Photogrammetric Engineering & Remote Sensing*, 87(4), 267–277. <https://doi.org/10.14358/PERS.87.4.267>.

Kada, M., & Wichmann, A. (2015). Feature-driven generation of LoD3 building models from cadastral data and LiDAR. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, II-3/W5, 197–202. <https://doi.org/10.5194/isprsannals-II-3-W5-197-2015>.

Appendix

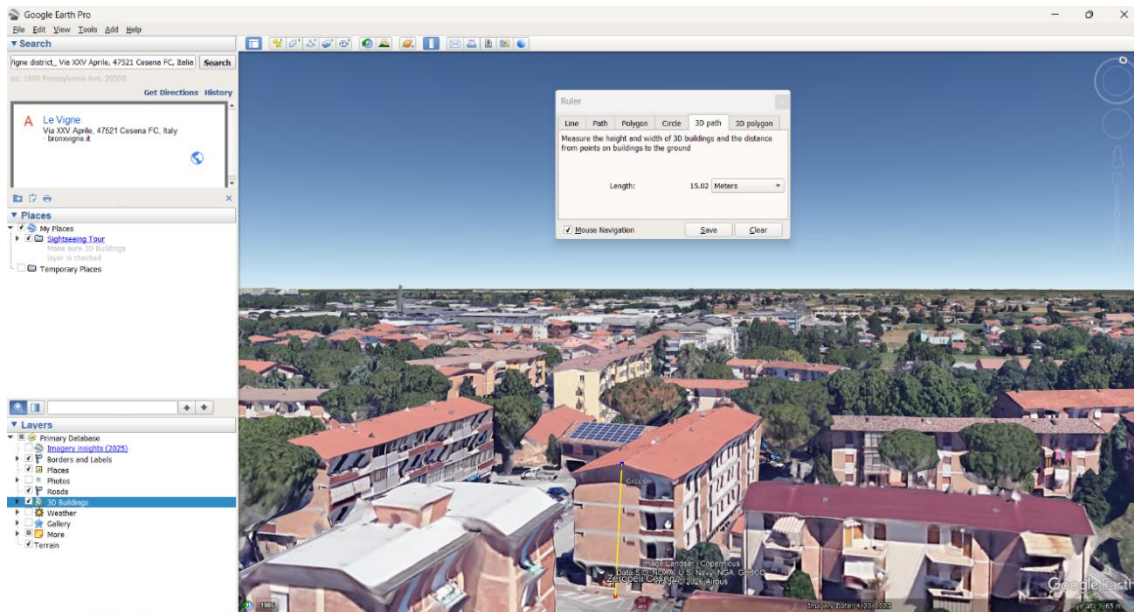


Figure 9: Google Earth Geometry Extraction

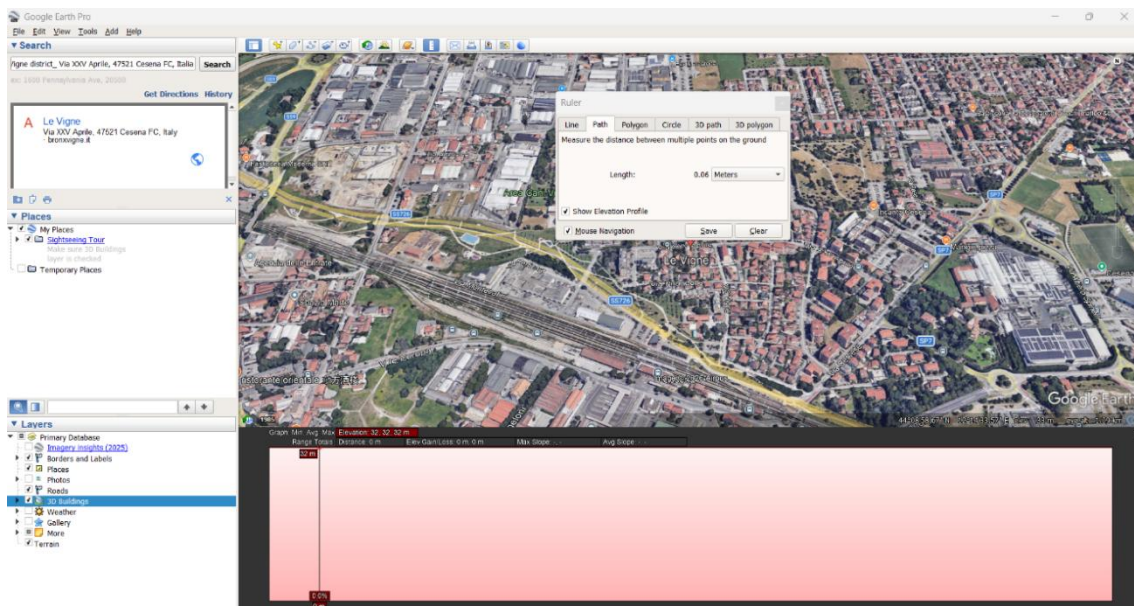


Figure 10: Google Earth Geometry Extraction

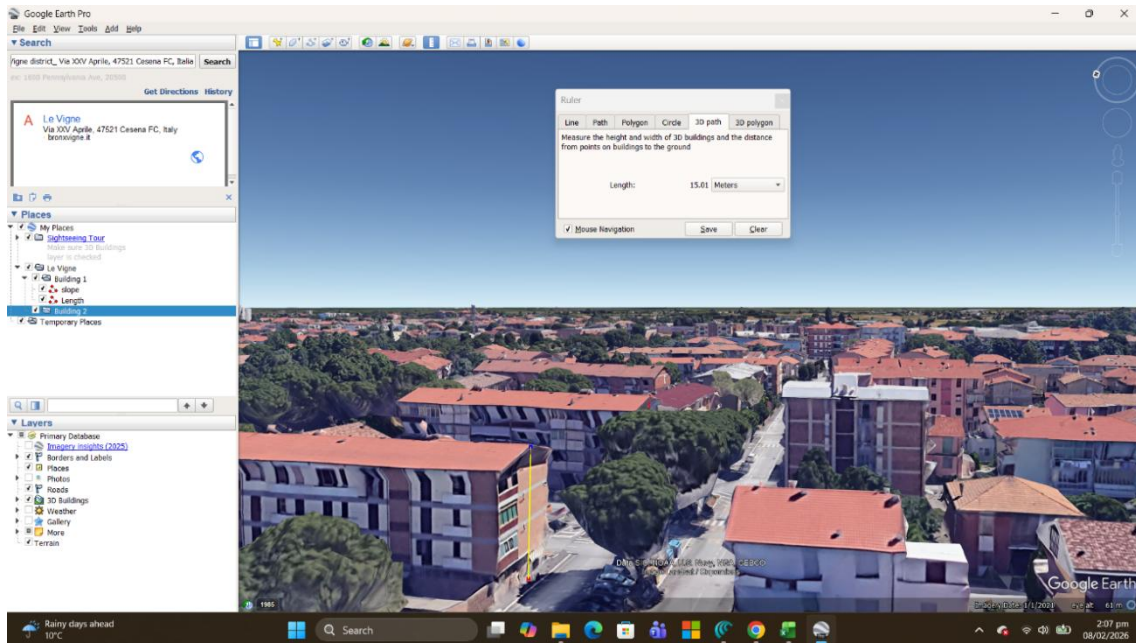


Figure 11: Length Calculation

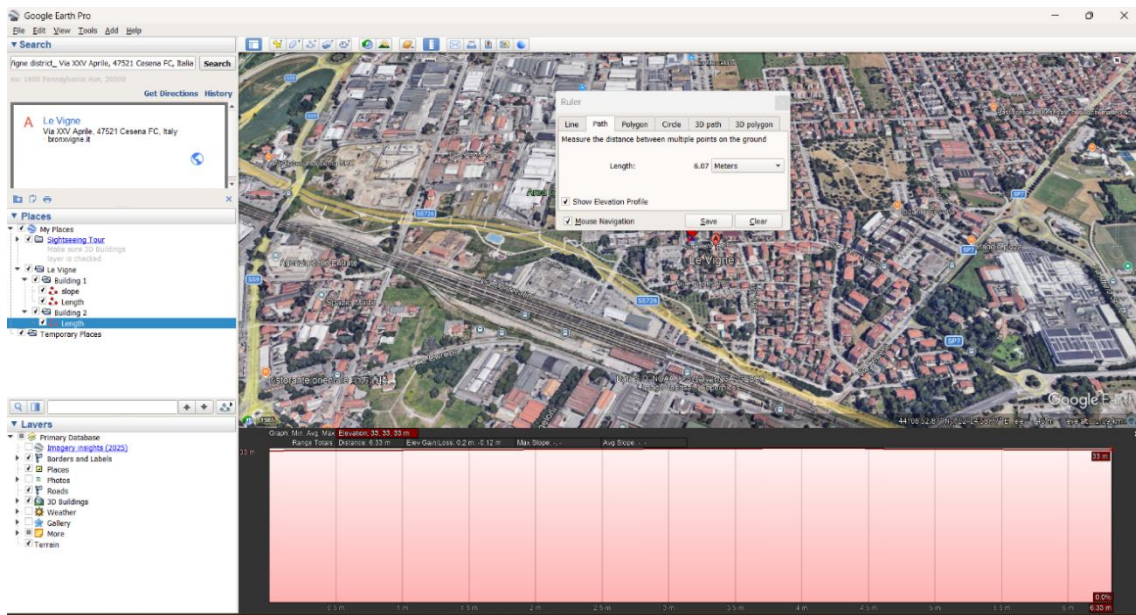


Figure 12: Length Calculation

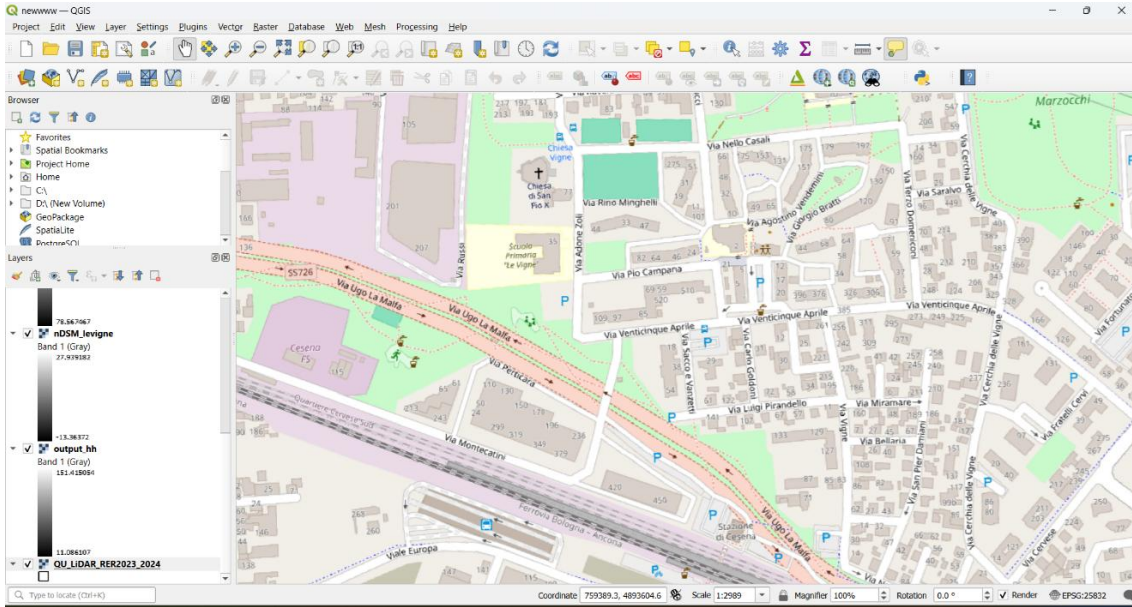


Figure 13: QGIS 2D Footprints

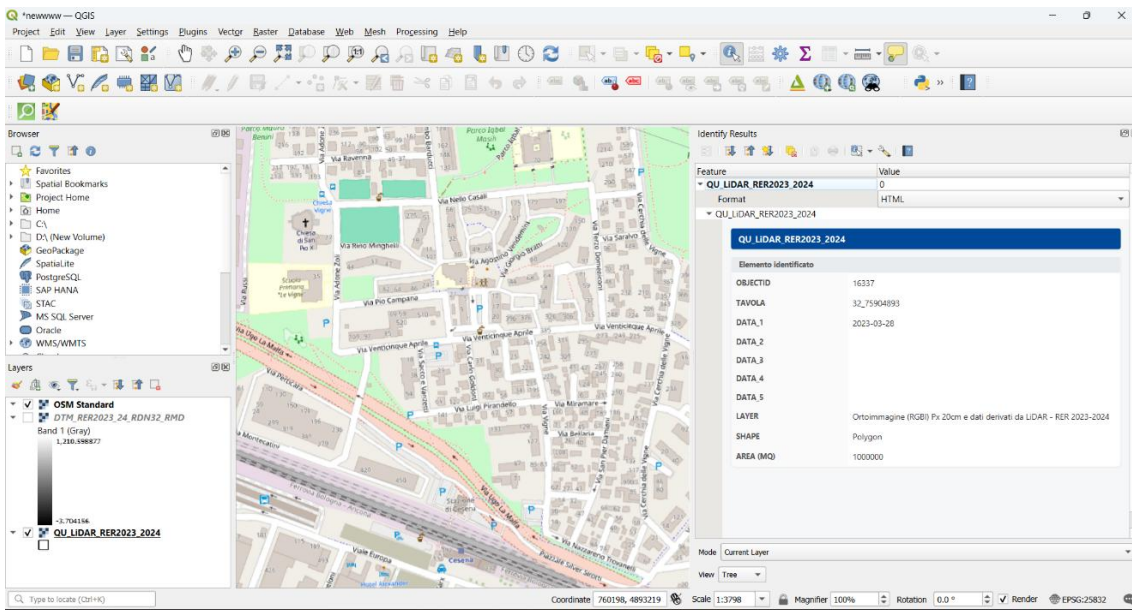


Figure 14: LIDAR Data

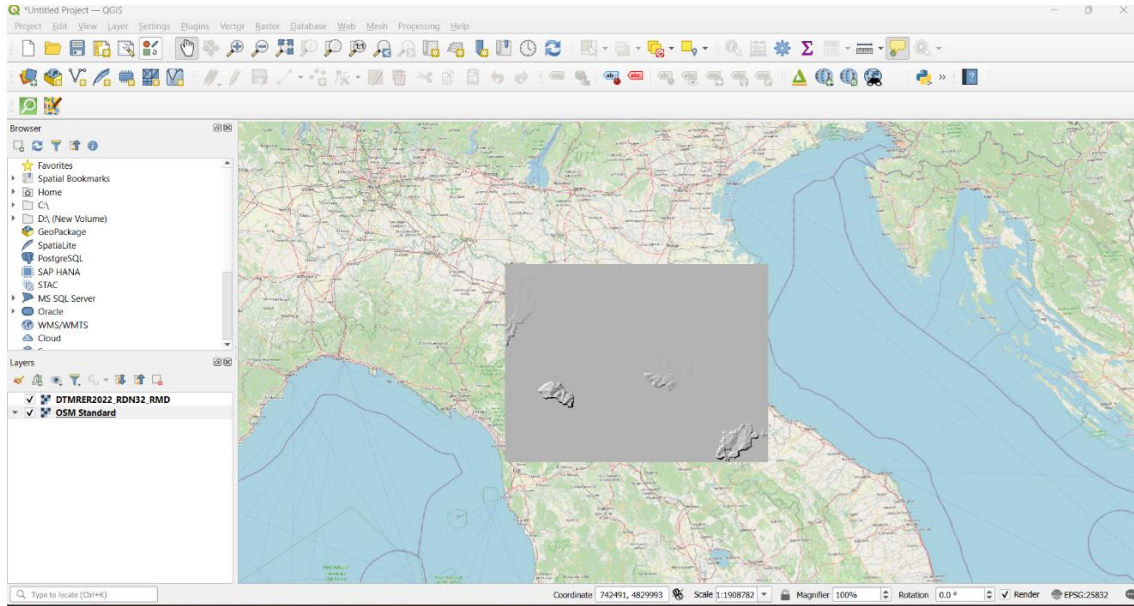


Figure 15: QGIS DTM

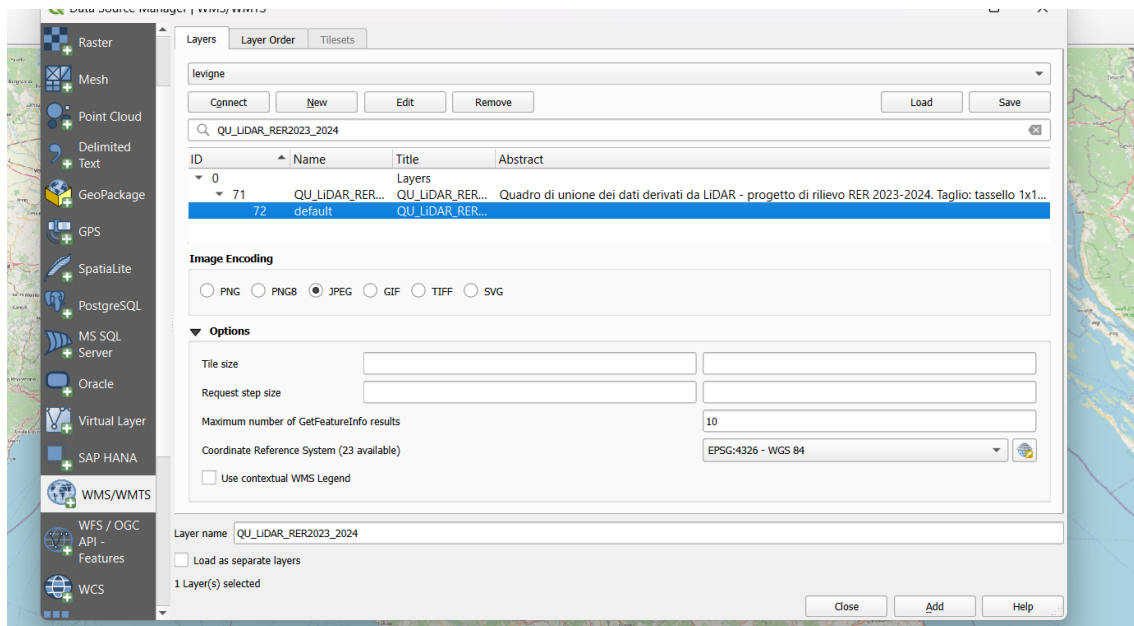


Figure 16: Importing National LIDAR in QGIS

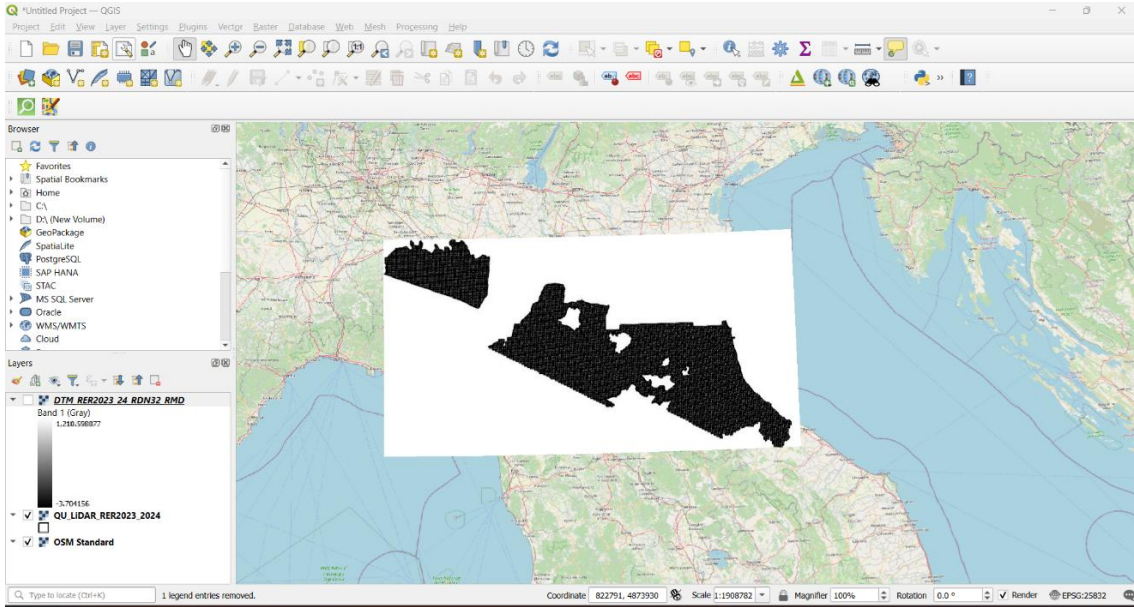


Figure 17: Overlaying DTM on Regional Footprints

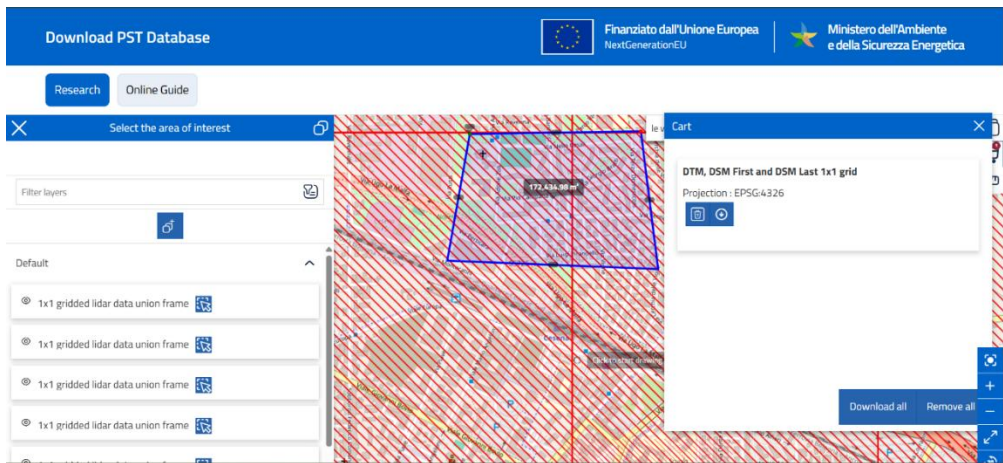


Figure 18: Downloading Pointcloud from PST Database

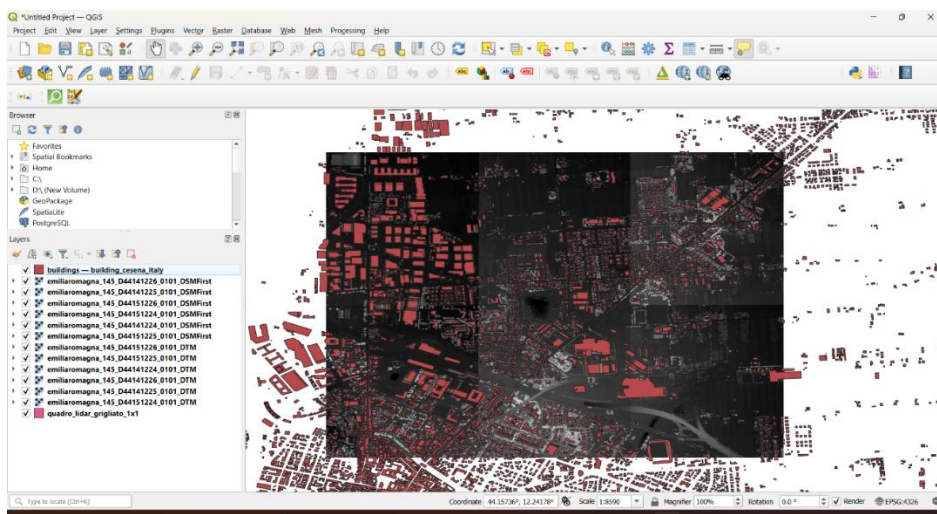


Figure 19: Overlaying of DSM on 2D Regional Footprints

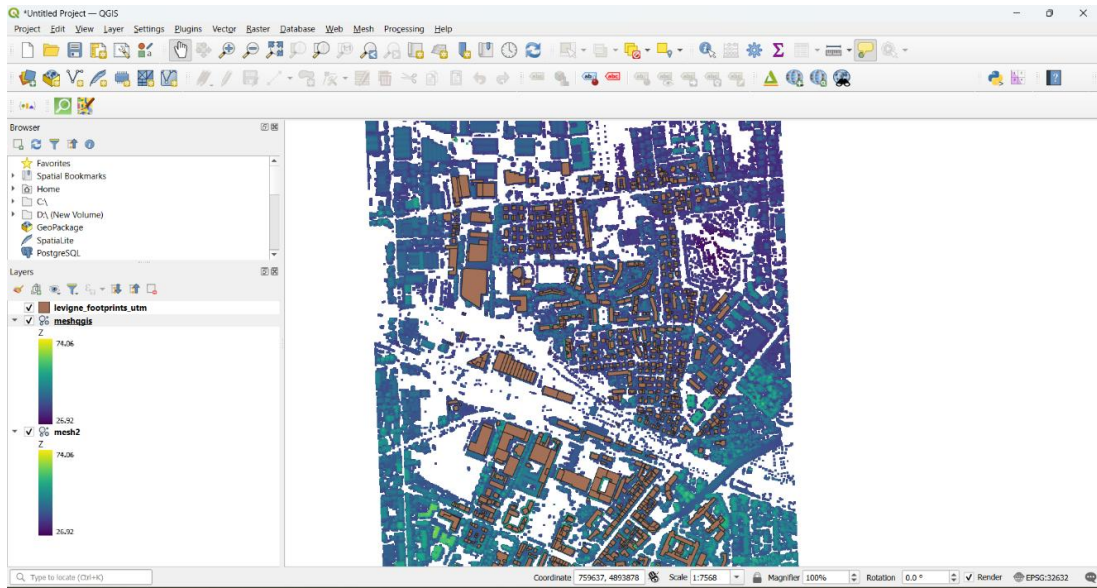


Figure 20: Mesh Overlaying on Polygons

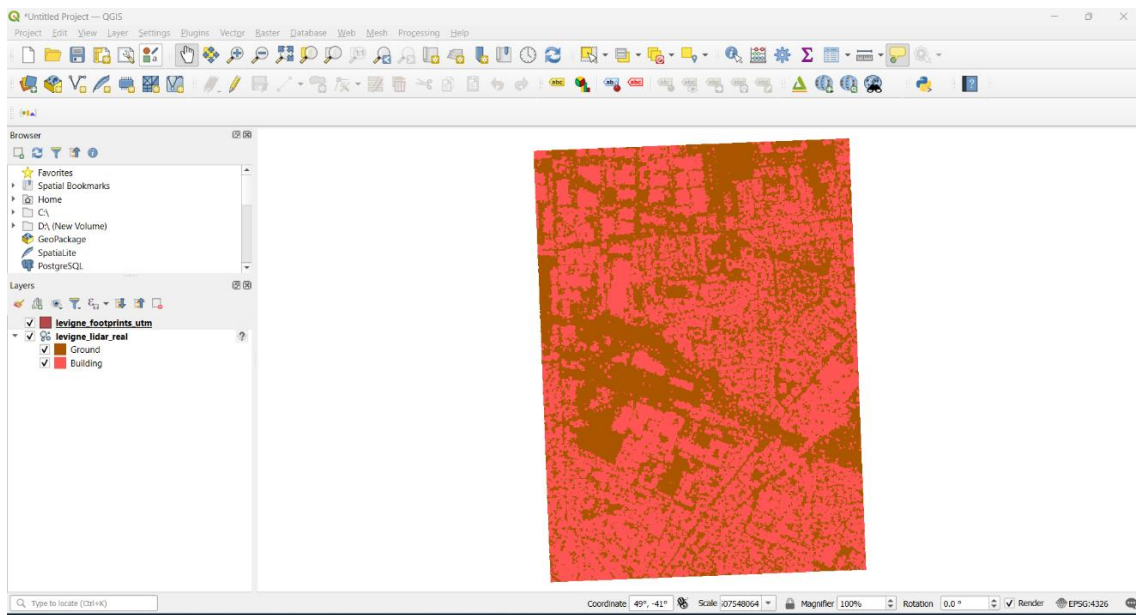


Figure 21: Classification of Pointcloud




Figure 22: Resulting 3D View

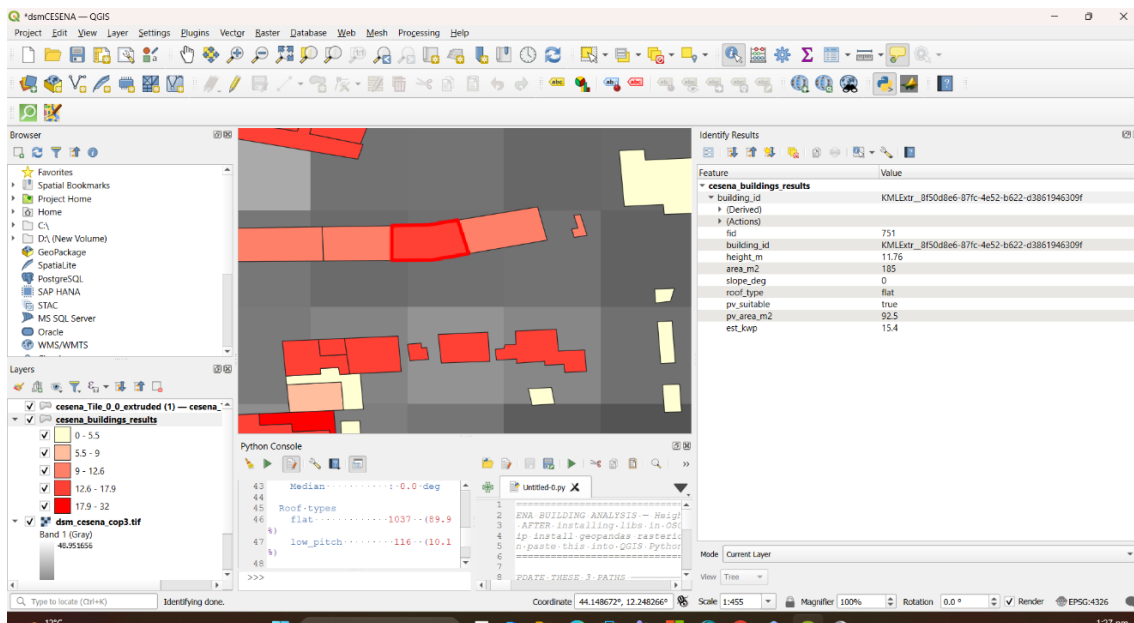


Figure 23: Slope Degree Estimation in QGIS

```

C:\Windows\System32\cmd.exe
Microsoft Windows [Version 10.0.26200.8037]
(c) Microsoft Corporation. All rights reserved.

C:\Users\DELL\Downloads\roofer-windows-x64-v1.0.0-beta.5>bin\roofer.exe levigne_utm.laz levigne_footprints_utm.geojson o
utput4
[2026-03-18 17:51:08.1357746] INFO Region of interest: 759454.020 4892689.257, 760290.894 4893936.023
[2026-03-18 17:51:08.1358067] INFO Number of source footprints: 1153
[2026-03-18 17:51:08.1358973] INFO Using 3 threads for the reconstructor pool, 8 threads in total (system offers 8)
[2026-03-18 17:51:08.1376838] INFO [Serializer] Writing output to output4
[2026-03-18 17:51:08.1761651] INFO Simplifying and buffering footprints...
[2026-03-18 17:51:08.2403236] INFO Cropping pointcloud ...
[2026-03-18 17:51:08.3160863] INFO Analysing pointcloud ...

C:\Users\DELL\Downloads\roofer-windows-x64-v1.0.0-beta.5>

```

Figure 24: Roofer Commands to Generate 3D Model

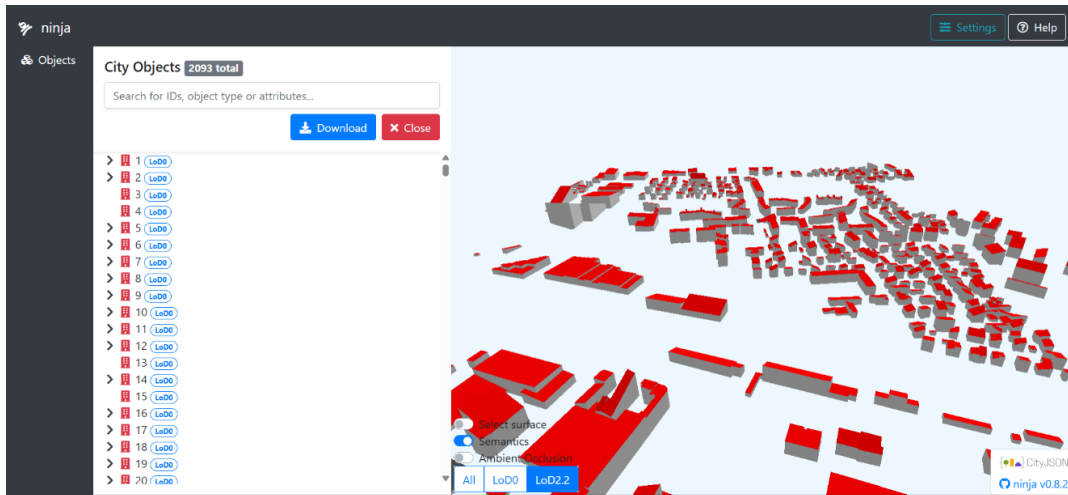


Figure 25: 3D Model Rendering in Ninja

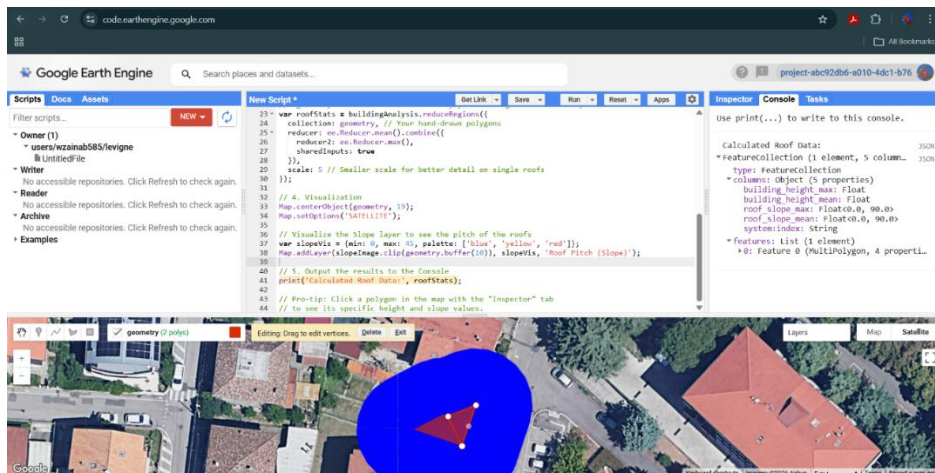


Figure 26: GEE Geometry Extraction

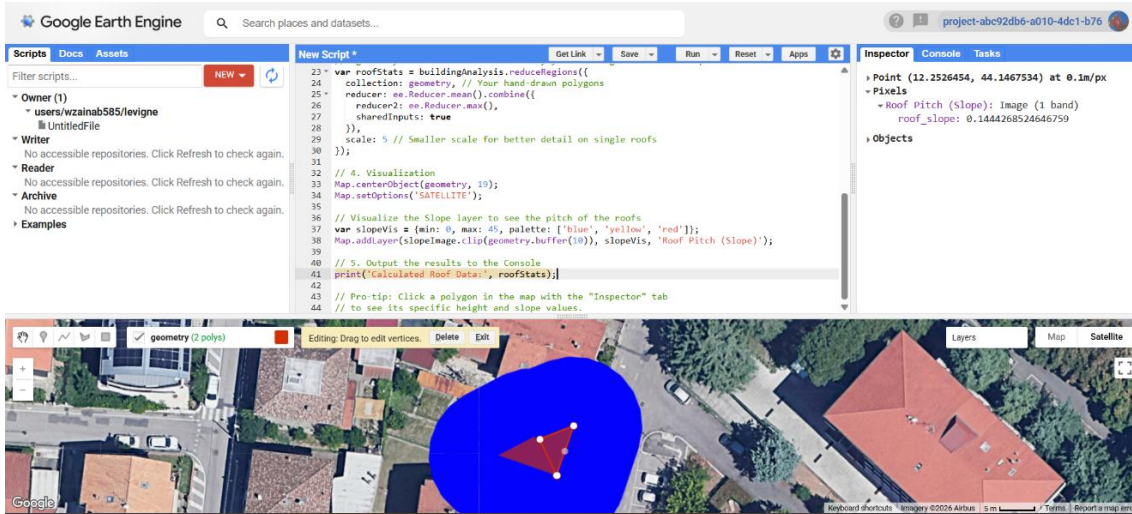


Figure 27: GEE Slope Calculation

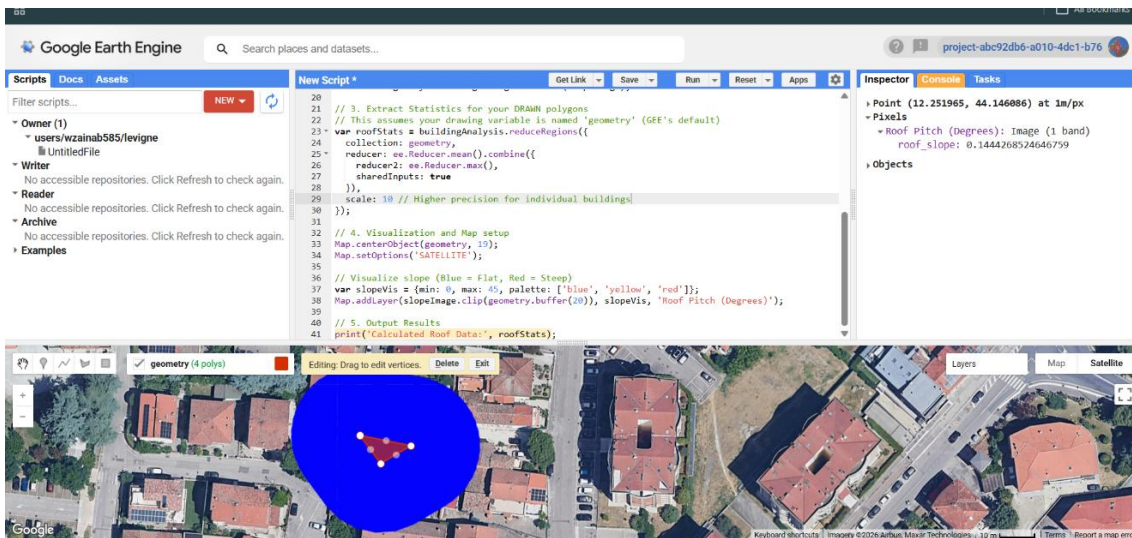


Figure 28: JAVA Scapping in GEE

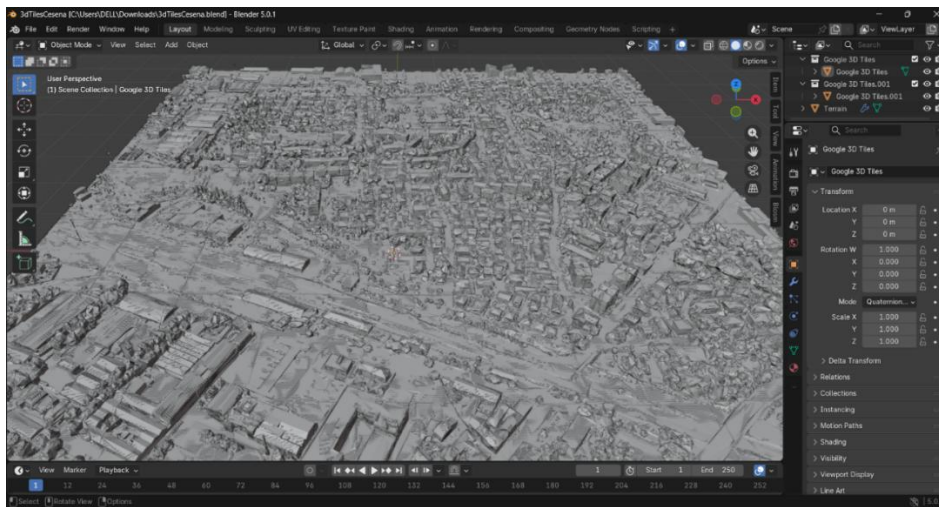


Figure 29: Mesh Extraction via Blossm Plugin

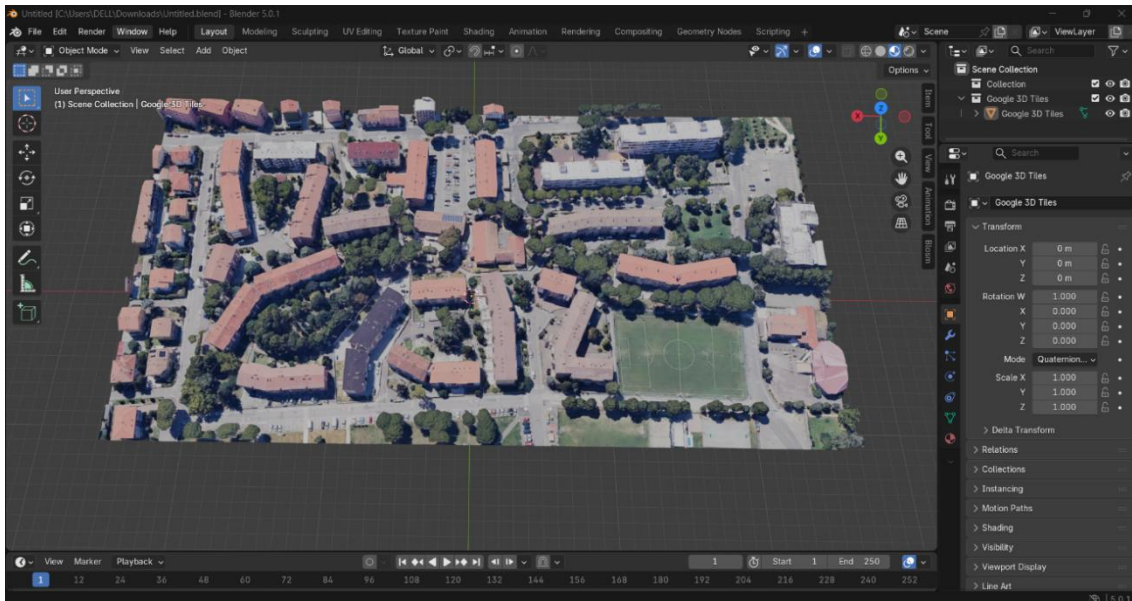


Figure 30: Google 3D Tiles Mesh

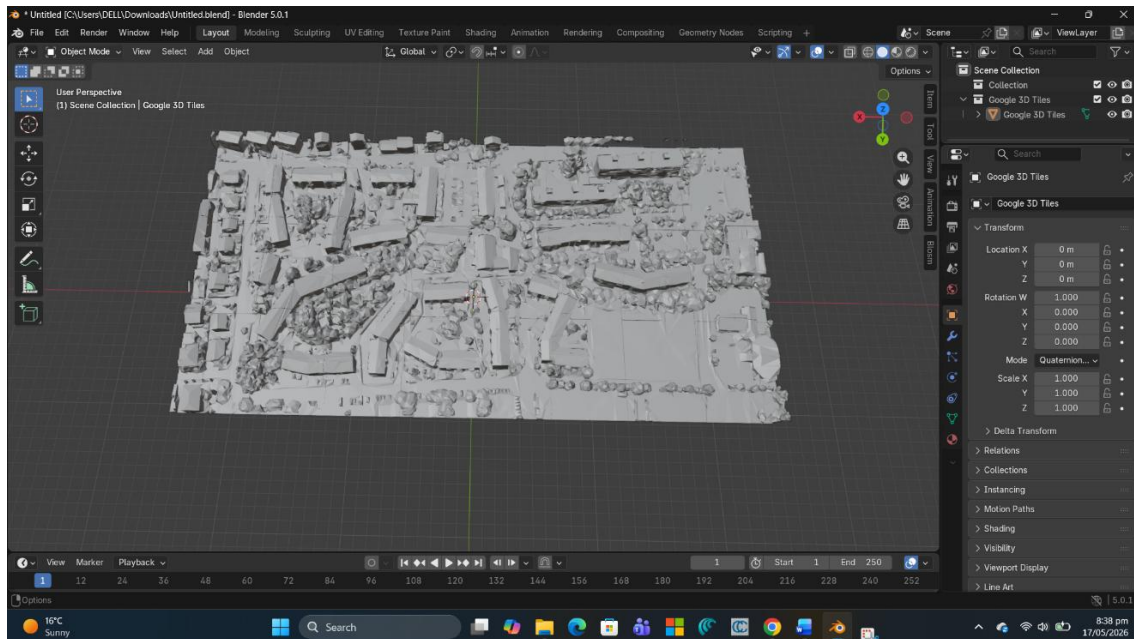


Figure 31: Google 3D Tiles Mesh in Blender

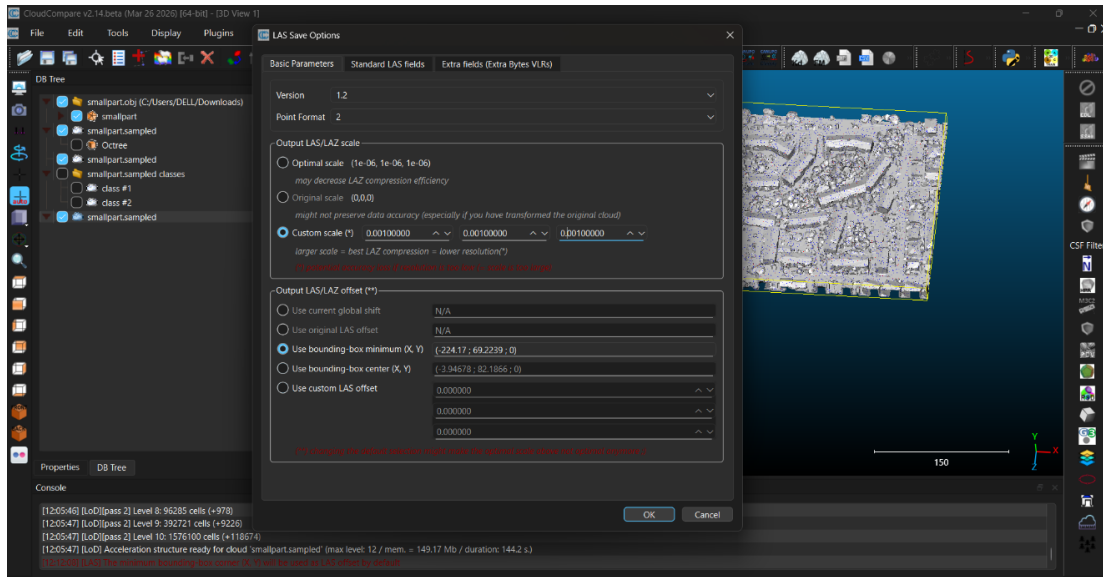


Figure 32: Mesh Cleaning in CloudCompare

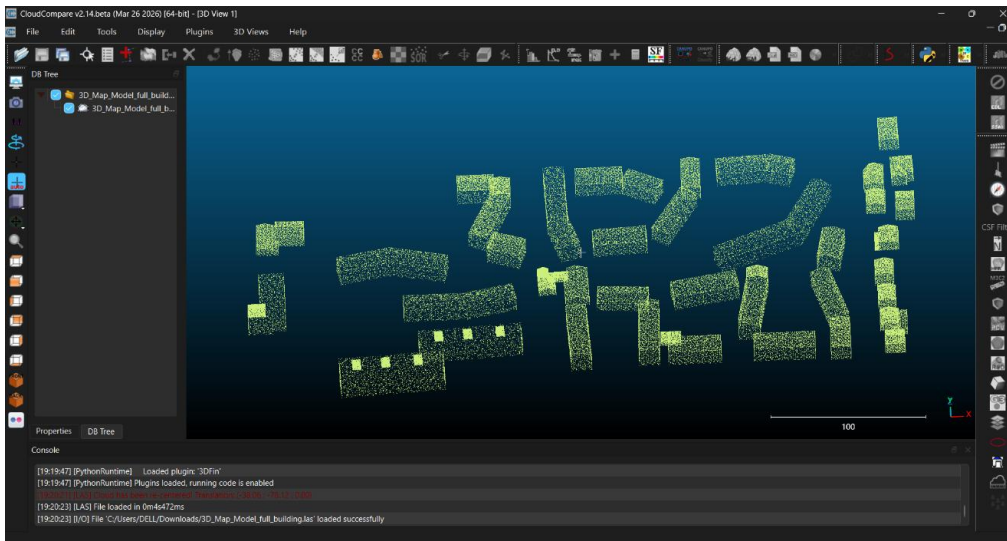


Figure 33: Clean Pointcloud Generation via CANUPO

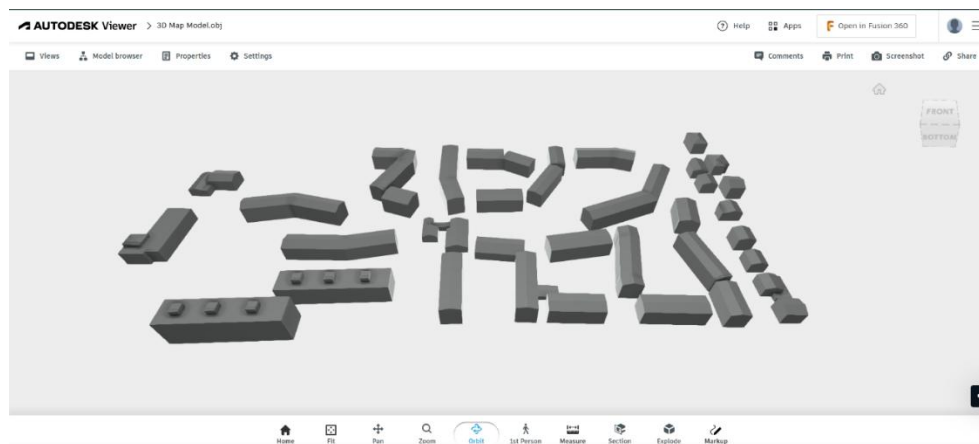
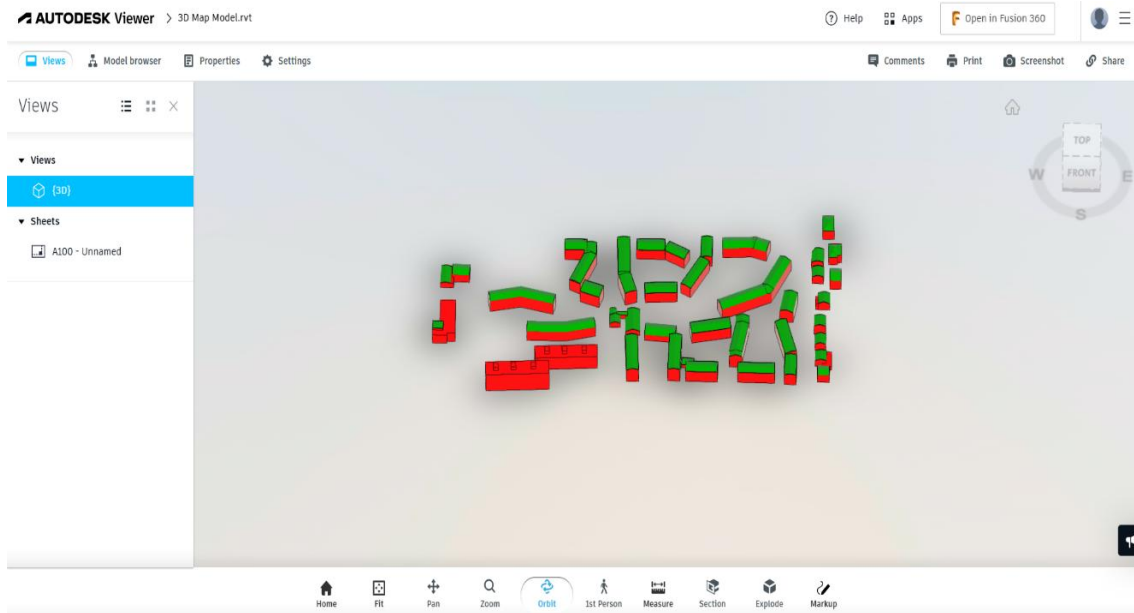


Figure 34: 3D Model View in Autodesk Viewer**Figure 35: Classification of Sloped Roofs and Walls**