ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences, Volume XI-2-2026
XXV ISPRS Congress 2026 “From Imagery to Understanding”, Commission II, 4–11 July 2026, Toronto, Canada

Wide-area Scene Reconstruction with Polyhedral Buildings featuring Recognized Regularities

Jochen Meidow

Fraunhofer IOSB, Ettlingen, Germany – jochen.meidow@iosb.fraunhofer.de

Keywords: Building model, boundary representation, reasoning, geometric constraints, manifold constraints

Abstract

The modeling of buildings suffers from a dichotomy between generic and specific representations: the lack of domain knowledge in flexible models that can represent many shapes, and the restricted geometry of pre-specified parametric building primitives. To fill this gap, we propose using general boundary representations enriched with automatically recognized and enforced geometric constraints derived from human-made regularities. The proposed reasoning process relies on the statistics of the planar point groups extracted from airborne-captured point clouds. Hence, a chosen significance level is the only process parameter. To enforce the creation of sound solids, we apply manifold constraints for the generation of the boundary representations. The feasibility and usability of the approach are demonstrated by evaluating an airborne-captured laser scan containing approximately 7,600 buildings over an area of 50 km^2 featuring both inner-city and rural landscapes.

1. Introduction

1.1 Motivation

The modeling of buildings suffers from a sharp division between two opposing, mutually exclusive concepts: generic and specific representations of human-made solids. Generic models allow for the representation of objects with almost arbitrary complexity. The model instantiation is data-driven; no building-specific knowledge is required or contained. Examples for generic representations are boundary representations and polyhedra. In contrast, specific models, e.g., parametric models, are directly related to certain building types. For model-driven reconstruction, predefined building shapes or primitives are required—usually provided in a library.

This dichotomy poses a dilemma, and the challenge is to introduce knowledge of building shapes without being overly restrictive (Heuel and Kolbe, 2001). Table 1 contrasts a specific representation of building (here a parametrized saddleback roof) with a generic representation using an unconstrained polyhedron.

specific	generic
e.g., parametric model	e.g., b-rep, polyhedron
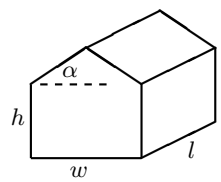	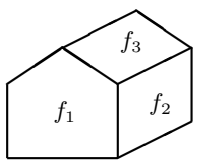
related to buildings types but too restrictive	no domain knowledge but too unspecific

Table 1. Dichotomy of typical as-built models for buildings. Left side: parametric model of symmetric gable roof with width w, height h, length l, and roof slope α. Right side: boundary representation (b-rep) with seven faces.

Thus, the goal is to provide a powerful, general-purpose representation that is capable of representing a wide variety of shapes, possibly accompanied by domain-specific knowledge in the form of recognized constraints.

Available methods for data-driven and model-driven reconstruction typically record roof shapes and building footprints with the required accuracy and detail, then use this information to generate a geometric representation of the building in a subsequent step (Haala and Kada, 2010).

To recognize human-made structures, a hypothesis-generation and testing strategy can be applied. Then, the constraints imposed by the detected geometric relations can be enforced by a subsequent adjustment procedure. Eventually, the vertices of the roof polygons are computed based on the roof planes and their intersection as proposed by (Xiong et al., 2014), for instance. A proof of concept for this approach can be found in (Meidow et al., 2016), which uses a very small data set.

1.2 Contributions

We focus on two contributions: (1) the process of geometric reasoning with uncertain planes, and (2) the enforcement of manifold constraints to obtain “watertight” model instances as required for flood simulations, for instance.

Geometric reasoning. To bridge the gap between generic and specific representations, we utilize a boundary representation to represent polyhedra that feature possibly human-made regularities. For this, we consider generic constraints such as orthogonality and parallelism, as well as problem-specific constraints such as horizontally running roof ridges. The geometric constraints are formulated for the planes belonging to the roof areas captured by airborne laser scanning, for instance. These planes are derived from the observed point cloud and are therefore uncertain. Therefore, the geometric constraints at hand are identified through hypothesis testing and enforced by a subsequent reasoning process. The procedure requires only a significance level as a parameter.

Enforcing manifold constraints. Precise outlining of roof areas is crucial to obtain proper 3D models. Given the point cloud, we first consider the boundary edges of a triangulation. These edges are snapped to possibly nearby present gutter lines and possibly nearby intersection lines of planes. Polylines not in the vicinity of straight lines are straightened by vertex decimation. The resulting straight line segments have to be joined to

This contribution has been peer-reviewed. The double-blind peer-review was conducted on the basis of the full paper.
https://doi.org/10.5194/isprs-annals-XI-2-2026-313-2026

form proper, especially non-self-intersecting, polygons representing the roof areas. As a result, each edge of the reconstruction is adjacent to exactly two faces. This manifold constraint ensures that the resulting model instances are “watertight” – a property that enables, for example, the computation of volumes. In (Nan and Wonka, 2017), manifold constraints are enforced in 3D by first intersecting all planes corresponding to observed or hypothesized faces. This yields a complete subdivision, or cell decomposition, of the space, consisting of many convex polyhedra. The faces of these polyhedra are then the candidates for a boundary representation of the solid. By applying an optimization for a binary linear program, the faces are selected which are (1) supported by the point observations and (2) feature edges that are adjacent to exactly two other faces each. Inspired by the work of (Nan and Wonka, 2017), we formulate a binary linear integer program to construct the roof area outlines in the horizontal plane. The endpoints of straight line segments and polylines are linked to form polygons, i.e., closed polylines.

In contrast to the most recent approaches, we use only existing domain knowledge in the form of recognized geometric constraints. No learning of roof topologies is required as proposed in (Ren et al., 2021), for instance.

The feasibility and usability of the approach are demonstrated through a large-scale reconstruction featuring buildings with a level of detail 2 as defined by the Open Geospatial Consortium (OGC, 2021) and illustrated in Table 1. Additionally, extensions such as alcoves, large wall indentations, and roof superstructures (e.g., dormers) are included, resulting in the refined level of detail 2.2 described in (Biljecki et al., 2016). The captured scene covers about 50 km^2, spans urban and rural areas, and contains approximately 7,600 buildings and building complexes.

2. Modeling and Utilized Methods

The proposed instantiation of building models is based on observed point clouds obtained either by photogrammetry or by airborne laser scanning. Thus, the first task is to classify these points and extract those lying on the buildings’ hulls. Grouping these points by evaluating planarity and proximity yields the required point groups that represent the roof areas. The uncertain planes corresponding to the roof areas serve as the basis for a reasoning process that comprises the recognition and the enforcement of geometric constraints. Enforcing manifold constraints ensures “watertight” model instances.

2.1 Classification and Grouping of Points

For the building reconstruction, we need the point groups representing the roof areas. Given the point cloud of the laser scan, we separate the “roof points” and the “non-roof points” in the first instance. This can be accomplished conventionally by first extracting the “non-terrain” points, e.g., by using an existing terrain model or by filtering the data (Zhang et al., 2016). The roof points are then extracted using a convenient grouping process. We utilize a modified RANSAC algorithm to find planar point groups that represent the roof areas (Schnabel et al., 2007).

2.2 Representation of the Buildings and Model Assumptions

To benefit from a generic, and therefore flexible representation, we utilize a boundary representation to model buildings as polyhedra. This implies that we assume the human-made objects of interest are completely bounded by flat faces.

Furthermore, due to the scanning geometry and inevitable occlusions, we cannot capture all faces of a building. Thus, to cope with missing data and to obtain reasonable outlines, we make several model assumptions following common domain knowledge:

1. The walls of facades and the walls constituting step edges on the roofs are vertical.
2. Gutters run horizontally.

We presume that these assumptions hold for almost all buildings.

2.3 Recognition and Enforcement of Geometric Constraints

Almost all buildings exhibit some geometric regularities that should be recognized and established. We formulated these relations as constraints for the planes corresponding to the roof areas. Thereby, we distinguish between generic constraints, such as parallelism and orthogonality, and problem-specific constraints, such as horizontally running roof ridges. Table 2 lists conceivable constraints.

type	name	arity	dof
generic	two parallel planes	2	2
	two orthogonal planes	2	1
	two identical planes	2	3
	four copunctual planes	4	1
specific	a horizontal roof plane	1	1
	a horizontal ridge line	2	1
	two planes with identical slopes	2	1
	two in xy orthogonal oriented planes	2	1

Table 2. Generic and problem-specific constraints for adjacent roof areas. For each constraint, the number of involved planes (arity) and the number of independent equations (degrees of freedom, dof) is given.

Four or more planes (not pairwise identical) with a common point are called “copunctual” or “concurrent” in the following. Please note that such an intersection point may lie at infinity — a situation handled conveniently by the calculus of projective geometry. Furthermore, the use of homogeneous representations yields simple polynomial equations for all constraints. This enables the use of tools from commutative algebra, e.g., to identify functionally independent equations (Awange and Grafarend, 2005, Meidow and Hammer, 2016).

2.4 Enforcement of Manifold Constraints

Topological consistency, completeness, and integrity are essential properties of model instances (Mäntylä, 1987). To obtain accurate models, we initially aim to accurately outline the building and roof areas in the horizontal plane.

Starting with the short boundary edges of a 2D alpha shape triangulation (Edelsbrunner et al., 1983), we snap them to possibly nearby projected intersection or gutter lines based on a proximity criterion. Then, collinear intermediate points can be removed, and the remaining vertices can be decimated using the Ramer-Douglas-Peucker algorithm (Ramer, 1972, Douglas and Peucker, 1973), for instance. Given these sets of straight line segments and polylines in 2D, the task is to form polygons that represent the projected roof areas by conveniently connecting their endpoints.

Figures 1 to 4 illustrate the polygonalization of a roof area belonging to a pyramid roof. In this context, it is possible to consider the common points of concurrent, i.e., copunctual, planes as the vertices of a geometric graph, see Figure 3 for the apex of the pyramid. Note that the reasoning is performed only with the uncertain planes that correspond to the roof areas and result from observed points. The upper corners of the faced walls are copunctual by construction since the facades are extrusions from the computed gutter lines. If, furthermore, the gutter lines run horizontally as assumed, and the XY-components of the planes’ normals are orthogonal (last constraint listed in Table 2), the walls must also be pairwise orthogonal.

Thus, geometric 2-edge-connected graphs with non-intersecting edges are sought. The edges of these graphs bound the – possibly multiply-connected – faces representing the roof areas. In this context, the graph exterior face (or outer face) is the infinitely large, unbounded plane that surrounds the graph.

Binary integer program. We solve a constraint satisfaction problem (CSP) to find the pairs of vertices to be connected to form a 2-edge-connected graph. Conceptually, this corresponds to the solution of the travelling salesman problem with the special feature that many edges are already given by the – possibly snapped – triangulation edges (see Figure 3).

Starting with an initial face, we consider all V vertices that belong to a face f and determine all possible pairings. This provides a set of hypothetical edges to be added. Hypothetical pairings with an unusually large point distance of, say, more than 5 m can be removed from this set. Furthermore, each vertex with label f must have two outgoing edges labeled with f, which provides V linear manifold constraints.

For each hypothetical edge, we have to decide whether selecting it yields a feasible solution with the minimal total length of all selected edges. With the vector $\boldsymbol{x} = [x_1, x_2, \ldots, x_E]^\mathsf{T}$ of binary variables $x_e \in \{0, 1\}$ and the lengths or distances $\boldsymbol{d} = [d_1, d_2, \ldots, d_E]^\mathsf{T}$, the binary or 0-1 integer program reads

$$\min_x \left\{ \boldsymbol{d}^\mathsf{T} \boldsymbol{x} \right\} \qquad \text{s.t.} \qquad \boldsymbol{A}\boldsymbol{x} = \boldsymbol{b} \tag{1}$$

with the objective function $y = \boldsymbol{d}^\mathsf{T}\boldsymbol{x}$ and the linear manifold constraints $\boldsymbol{A}\boldsymbol{x} = \boldsymbol{b}$. The constraints are given by the V vector $\boldsymbol{b}$ with $b_v = 2$ outgoing edges for all v, and the $V \times E$ coefficient matrix $\boldsymbol{A} = (a_{ve})$. If the edge e can be an outgoing edge of vertex v, $a_{ve} = 1$ holds, $a_{ev} = 0$ otherwise.

Polygon creation. After the identification of the vertex pairs to be linked, we check if the straight lines corresponding to the adjacent straight line segments intersect in a nearby point. If a hypothetical intersection point is close to a pair of vertices to be linked, we introduce it as an additional vertex of the polygon (see Figure 4). After forming the polygon for a face, the procedure is applied to the next face and so on.

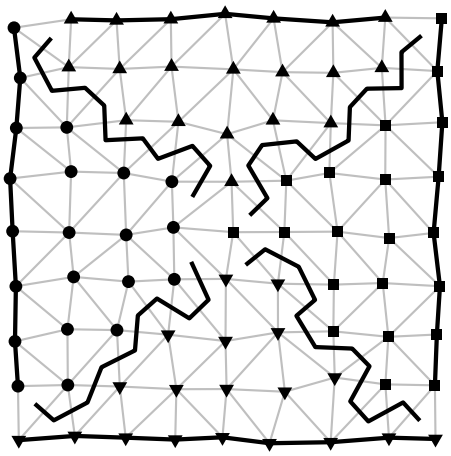

Figure 1. Illustration of the polygonalization with synthetic data: Point clusters (symbols), alpha shape triangulation (gray), boundary edges of each cluster, and linked midpoints of the triangles between two clusters each.

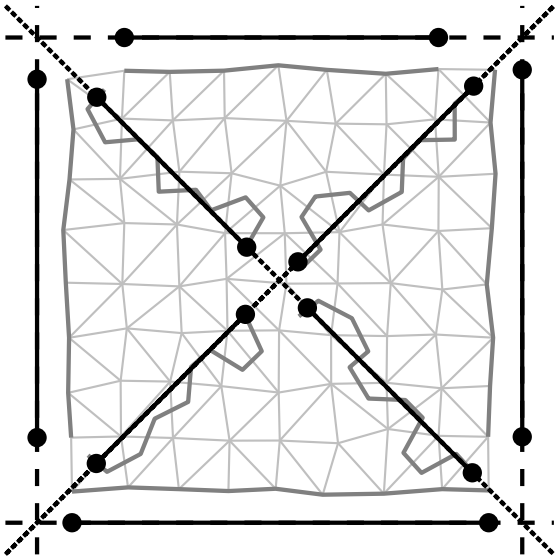

Figure 2. Gutter lines (dotted) corresponding to one height, intersection lines of planes (dashed), and the segments obtained after snapping the edges to nearby lines (hips or gutters) and removing collinear points.

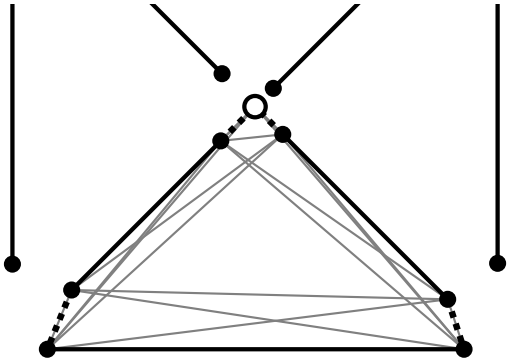

Figure 3. Polygonization of a first roof area: Hypothetical links for the vertices (gray edges), copunctual point of the four planes (circle), and solution of the binary integer program (dotted edges).

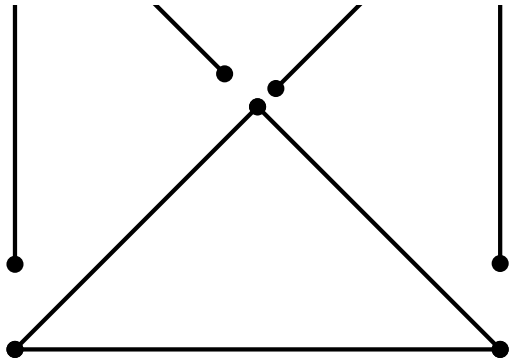

Figure 4. Polygon of a first roof area. For each pair of vertices to be linked (dotted edges in Figure 3), the intersection point of the corresponding pair of straight lines is computed.

For a pyramid roof, one obtains a geometrically perfect pyramid if the slopes of the four faces are recognized to be identical or if the normal vectors of adjacent faces are orthogonal in projection onto the horizontal XY-plane (see Table 2).

2.5 Determination of Gutter Lines

The proposed reconstruction relies in particular on the information obtained from slanted roof areas. Hip lines, roof valley, and ridges are easily obtained by the intersection of the planes corresponding to adjacent faces. However, identifying gutter lines or eaves is somewhat more challenging, as locally, only one point group or roof area is involved.

To determine the eaves' heights, we consider the boundary points of the 3D triangulations for each group of adjacent slanted roof areas. In the vicinity of the eaves, the heights of these points are approximately identical. These heights can be grouped, e.g., by estimating the components of a Gaussian mixture model. To initiate this process, we first apply a kernel density estimation with a bandwidth of, say, 5 cm, which reflects the expected variation in the observed heights. The positions of the K resulting peaks are then used as initial values for estimating K mixture components. Eventually, these estimates are filtered based on the mixing proportions, which reflect their relevance.

Figure 5 shows an example of a model selection with a histogram and estimated probability density functions. Of course, the actual gutter heights are slightly lower as we observe only points on the roof areas right above the gutters.

Finally, the gutter lines are obtained by intersecting the slanted roof planes with the horizontal planes corresponding to the gutter heights. The projections of these straight 3D lines onto the horizontal XY-plane are used to form the outlines of the roof areas as illustrated in Figure 2.

3. Reasoning with uncertain Planes

The reasoning process for the given uncertain planes representing the roof areas is the core of the proposed workflow. Therefore, we provide more details for completeness.

Enforcing recognized constraints yields slightly adjusted planes for the roof areas. Therefore, the adjustment must be made before the fixation of the roof layout in the horizontal plane by outlining the roof areas and hence the building outline, as proposed in Section 2.4 above.

The reasoning comprises the following steps:

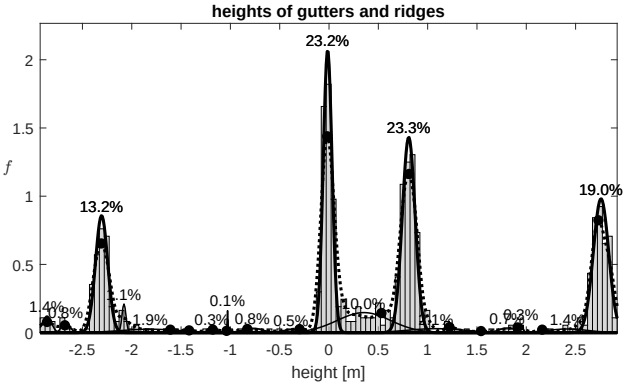

Figure 5. Model selection for gutter heights. Kernel smoothing (····) provides local maxima (●) suitable for the initialization of the estimation of the components of a Gaussian mixture model (—). The selected dominant components are drawn in bold (**—**).

1. Buildup of the roof topology graph or region adjacency graph based on the point groups representing the roof areas

2. Reasoning with the corresponding uncertain planes of the roof areas:

   (a) Identification of geometric relations by hypothesis generation and testing

   (b) Selection of sets with independent and consistent equations, i.e., non-redundant and non-contradictory constraints

   (c) Enforcement of the recognized constraints via adjustment

The steps are explained in more detail in the following paragraphs.

Buildup of the roof topology graph. Given the point clusters for each roof area, we can check for geometric adjacencies in 2D and in 3D. Two areas that are adjacent in 2D but not in 3D are separated by a step edge, i.e., a vertical wall on the roof. To speed up computations, we use the axis-aligned bounding boxes of the point clusters for culling. Two clusters are considered adjacent if the shortest distance in 2D and 3D, respectively, between their boundary points is below a heuristic, data-specific threshold, say 1.0 meters, which depends on the point density.

Reasoning. For each connected component of the 3D graph, a reasoning process is performed with the uncertain planes for the roof areas.

1. For the *hypotheses generation*, the maximal cliques in the roof topology graph are identified and evaluated for the unary, binary, ternary, and quaternary constraints listed in Table 2.

   We do not check for orthogonal roof planes since this rarely occurs. The non-observed facades and step edges are vertical by construction, and horizontal planes are considered as a problem-specific constraint. A symmetric saddleback roof, for instance, is represented by two roof areas with identical slopes and a horizontally running ridge line (Förstner and Wrobel, 2016, p. 429).

2. For *hypothesis testing*, heuristic pre-checks are performed. If passed, a conventional statistical hypothesis test is performed at the specified significance level. This yields a potentially huge set of conditional equations for the recognized constraints.

3. For the *selection* of independent and consistent conditional equations, we apply a greedy algorithm that evaluates numerical criteria, e.g., the condition numbers of the equation systems or the estimated ranks of the coefficient matrices (Meidow and Hammer, 2016).

4. Eventually, for the strict *enforcement* of the recognized constraints, an adjustment of the planes is performed (see below for details).

The result of this reasoning process is a set of consistent, i.e., non-contradicting and non-redundant, conditional equations. The adjusted parameter values for the planes satisfy these equations, and the intersection points and lines can be used to support the outlining of roof areas. This is demonstrated in Figure 3, where the common point of the four planes constitutes the apex of the pyramid.

3.1 Adjustment Model

We interpret the estimated parameters of the planes as observations that are accompanied by covariance matrices. With the recognized set of consistent constraints, an adjustment with observations only is applied:

The parameters of the planes are compiled in the vector of observations $\boldsymbol{y}$ and the corresponding covariance matrices in the block-diagonal matrix $\boldsymbol{\Sigma}_{yy}$. The constraints $\boldsymbol{c}(\widehat{\boldsymbol{y}}) = \boldsymbol{0}$ hold for the adjusted, i.e., constrained, observations $\widehat{\boldsymbol{y}}$. Minimizing the objective function

$$\Omega = (\widehat{\boldsymbol{y}} - \boldsymbol{y})^\mathsf{T} \boldsymbol{\Sigma}_{yy}^{-1} (\widehat{\boldsymbol{y}} - \boldsymbol{y}) \qquad \text{s.t.} \qquad \boldsymbol{c}(\widehat{\boldsymbol{y}}) = \boldsymbol{0} \tag{2}$$

with the linearized constraints yield the updates for the observations $\boldsymbol{y}$, see model “C” in (Förstner and Wrobel, 2016, p. 172) or (Koch, 1999) for instance.

This model cannot be applied directly when using homogeneous coordinates to represent geometric entities since the overparametrization leads to singular covariance matrices. However, to obtain numerically unique results, unit-1 constraints can be applied to all homogeneous vectors. When interpreting these vectors as points, they reside on the unit sphere, and an optimization on this curved manifold can be performed. By applying the concept of reduced homogeneous coordinates, we obtain a minimal representation with regular covariance matrices (Förstner and Wrobel, 2016, p. 371), and minimal rotations can be utilized to define vector bases spanning the tangent spaces (Meidow and Hammer, 2025).

4. Experiments

4.1 Data Set: Airborne captured Laserscan

A test dataset was captured in September 2017 by airborne laser scanning. The scan covers the city of Ettlingen, Germany, and its surroundings, encompassing approximately 7,600 buildings over an area of approximately 50 km^2. The center of the town is characterized by dense development, with historic buildings of rather complex shapes, and some trees. The surroundings include agricultural areas, residential areas in neighboring villages, and some suburban areas of Karlsruhe.

The test area was covered by several strips captured by an airborne laser scanner. Inside an individual strip, the average point density is approximately 20 pts/m^2. The captured point cloud consists of approximately 1,536 billion geo-referenced points, compiled into tiles. Points on the buildings’ hulls usually represent roof areas, rarely facades. Multiple echoes and intensities were recorded, and the points are accompanied by a classification in terrain and non-terrain points provided by the company that conducted the surveying flight.

4.2 Preprocessing of the Point Cloud

The points representing the building are almost exclusively located on the roof areas. Thus, we utilized a digital terrain model in raster representation to determine the ground levels of the buildings. To specify a terrain model, the points classified as “terrain” can be used. If such a classification is not available, the terrain model can be computed as usual from scratch, e.g., using the method proposed in (Zhang et al., 2016).

Given the original point cloud, we performed the following – quite conventional – preprocessing steps:

1. *Extraction of non-terrain points.* For each tile, the non-terrain points are extracted, and the sets are merged into a single point cloud.

2. *Filtering of non-terrain points.* To exclude points on the voluminous vegetation, i.e., trees and bushes, we delete isolated points and points resulting from rays with multiple returns.

3. *Clustering of non-terrain points in 2D (connected components) necleting the height coordinates.* This step yields clusters of points, each presumably representing a building or a building complex.

4. *Extraction of planar point groups.* For each connected component, we apply the RANSAC method to identify groups of adjacent, planar scan points, each representing a roof area (Schnabel et al., 2007).

As a result, we obtain point clusters, each representing a building or a building complex, with labels indicating the various roof areas. These groups, along with the derived or given digital terrain model, serve as input for the building reconstruction as described in Section 3 above.

4.3 Results

The significance level of 0.05 was chosen for hypothesis testing in the experiments. Not recognizing relations that are actually present (type I error) does not prevent reconstruction.

Selected buildings and building complexes. Figure 6 shows a reconstruction with eight roof areas and the corresponding planes being copunctual. The copunctuality constraint is formulated for a minimum set of four copunctual planes (Three planes are always copunctual, possibly with an intersection point at infinity). Thus, in principle, $\binom{8}{4} = 70$ constraints could be inferred for noiseless observations, since all roof areas are adjacent. From this set of constraints, the greedy algorithm

identified five functionally independent conditional equations. Apart from this, the roof areas feature identical slopes and horizontal ridge lines in the reconstruction. Some of the adjacent facades are almost coplanar. These faces can be merged during a post-processing stage that has not yet been implemented.

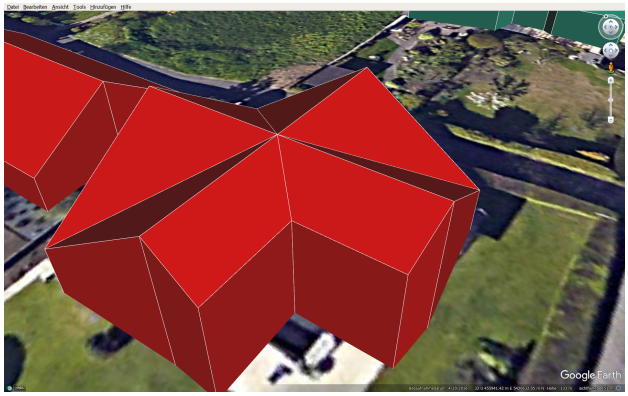

Figure 6. Building with eight copunctual planes. The greedy algorithm identified five independent constraints, each enforcing the copunctuality of four planes. Furthermore, the eight roof slopes are identical.

Figure 7 shows buildings bounded by curved surfaces. The two buildings – presumably indoor tennis courts – are cylindrical. Thus, a parametric representation would be more appropriate, as it is less complex. However, automatic model selection using information-theoretic measures appears challenging since only parts of the geometry shapes are often observed.

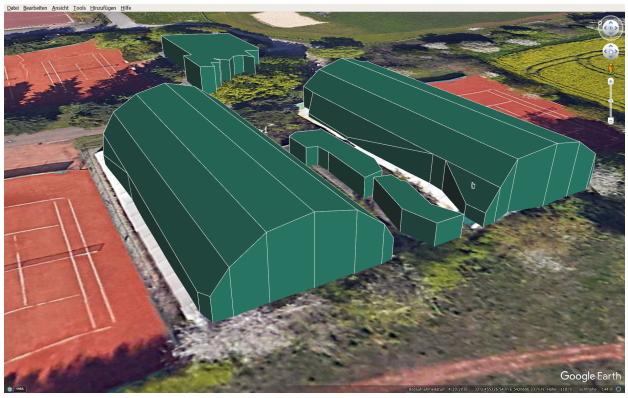

Figure 7. Examples of buildings bounded by curved surfaces (presumably indoor tennis courts).

Figure 8 shows examples with roof areas being multiply-connected, i.e., with holes because of the courtyards. This demonstrates the reconstruction process’s ability to handle arbitrary nested buildings.

Figure 9 shows a reconstruction of an inner-city block with old buildings. Human-made regularities are present but not omnipresent; some roofs are crooked and awry.

Eventually, Figure 10 shows examples of buildings with unusual roof shapes that the boundary representation can also represent.

Assessment. The assessment of the reconstruction results covers multiple aspects, particularly the effects of imposing

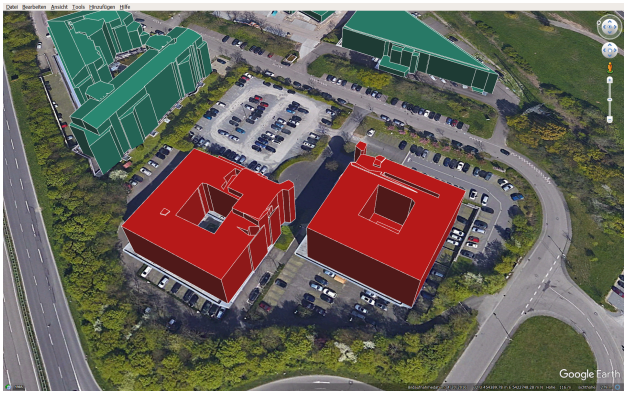

Figure 8. Examples of buildings with multiply-connected polygonal faces representing the roof areas.

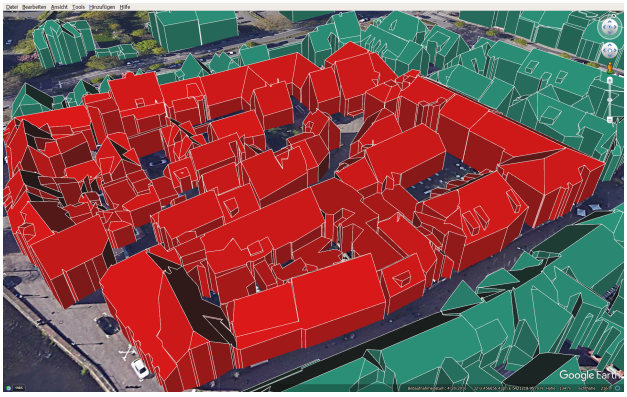

Figure 9. Reconstruction of an inner-city block with historical buildings featuring regularities in parts.

constraints during reconstruction and the completeness and geometric accuracy.

Imposing recognized geometric constraints on the unconstrained boundary representation does not necessarily change the model complexity with respect to the number of required faces: Only the copuntuality constraints will reduce the number of required vertices for a polyhedron. Furthermore, the constraints will only slightly affect the geometry of the reconstructed shapes, since they are recognized only when they are nearly satisfied. Hence, constrained and unconstrained reconstructions will not differ considerably with respect to geometry and completeness.

The completeness and geometric accuracy of the reconstructions depend on many factors: the accuracy of the chosen building model (specific or generic), the georeferencing (registration), the classification of the scan points, the clustering of the scan points in planar groups (RANSAC), the quality of the terrain model, the validity of model assumptions for the case of missing data, etc.

The agreement on an independent reconstruction serving as a reference with superior quality is challenging since, for a fair comparison, it should be based on the same preprocessed data set. For the test data set at hand, open geodata is available from the Landesamt für Geoinformation und Landesentwicklung (LGL) Baden-Württemberg. Among other things, the data comprises digital surface models and building reconstructions featuring LoD 2 encoded in CityGML. The buildings are obtained by considering various sensor data (laser scans, photogrammet-

This contribution has been peer-reviewed. The double-blind peer-review was conducted on the basis of the full paper.
https://doi.org/10.5194/isprs-annals-XI-2-2026-313-2026 | © Author(s) 2026. CC BY 4.0 License.

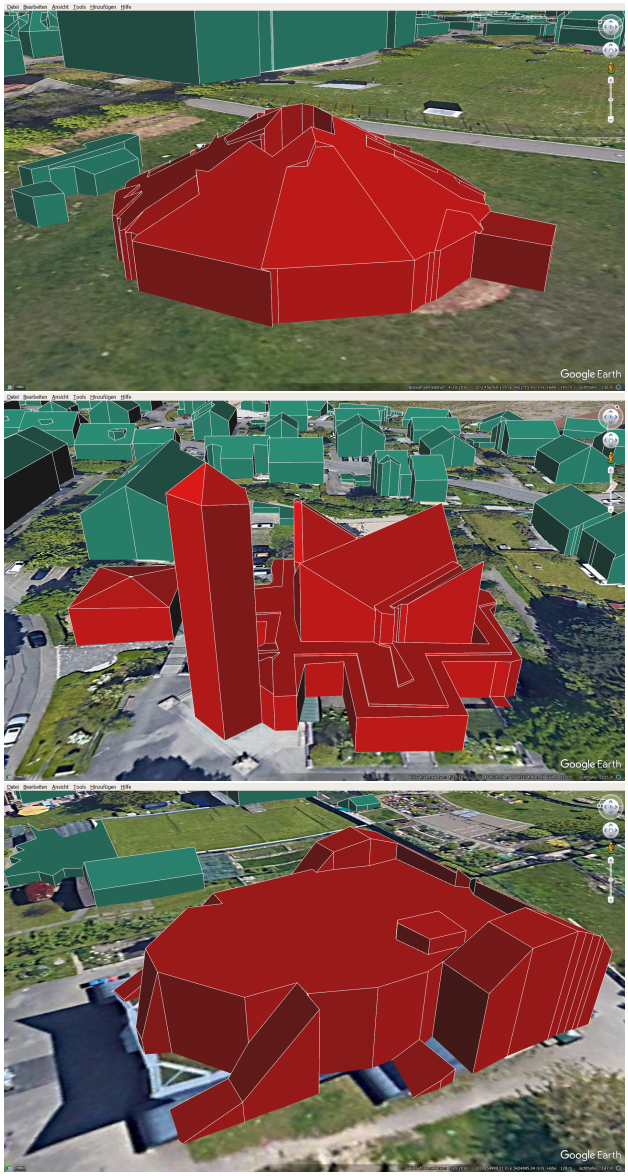

Figure 10. Examples of buildings featuring unusual architectures: A big top (circus tent), a church with a spire, and a “cat” used as a kindergarten. The reconstruction of the “cat” lacks the “ears”, as visible in the cast shadow.

ric point clouds), and building outlines from the cataster data. These geo products can be used together with the laser scan to perform at least a visual comparison; see Figures 11, 12 and 13.

5. Summary and Outlook

Summary. Boundary representations are suitable for describing objects featuring almost arbitrary complex shapes. However, they are too unspecific as they convey no domain knowledge. Therefore, we enrich these representations with automatically recognized and enforced constraints that reflect the regularities of buildings. This approach avoids using restrictive parametric models that represent only predefined shapes.

To detect human-made regularities, we pursue a hypothesis-generation-and-testing strategy. In doing so, we consider generic constraints, such as orthogonality or parallelism, as well as problem-specific constraints, e.g., horizontally running ridge

Figure 11. Proposed reconstruction and surface model. Surface model and orthopoto: LGL, www.lgl-bw.de,dl-de/by-2-0.

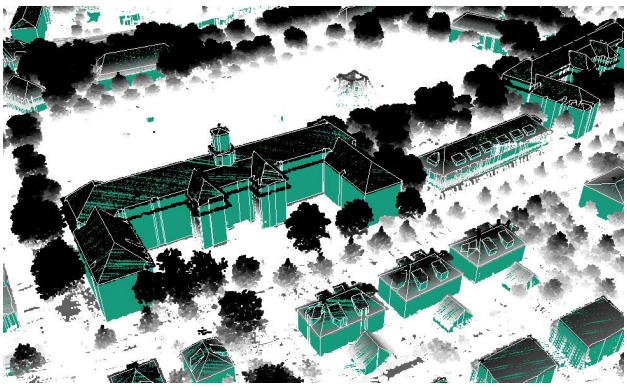

Figure 12. Proposed reconstruction and point cloud. Buildings: LGL, www.lgl-bw.de,dl-de/by-2-0.

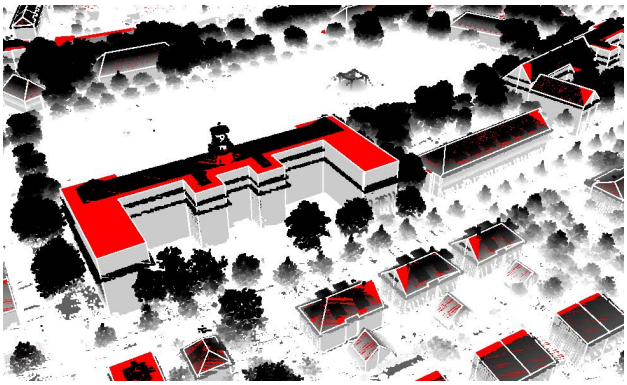

Figure 13. Reference data and point cloud. Buildings and orthopoto: LGL, www.lgl-bw.de,dl-de/by-2-0.

lines or roof areas with identical slopes. Especially the recognition of roof areas with corresponding copunctual planes, i.e., four or more planes with a common point, yields building models that are not only nicer but also less complex.

Quantifying uncertainty enables the rigorous application of statistical methods. This refers to testing hypotheses and to enforcing recognized constraints via adjustment. Thus, the only and essential parameter is the chosen significance level – a quantity with a clear interpretation that controls the sensitivity with respect to the recognition of regularities.

The recognized geometric relations are formulated using sets of equations. Then, a proposed greedy algorithm is applied to identify sets of non-redundant, i.e., functionally independent,

This contribution has been peer-reviewed. The double-blind peer-review was conducted on the basis of the full paper.
https://doi.org/10.5194/isprs-annals-XI-2-2026-313-2026 | © Author(s) 2026. CC BY 4.0 License.

and consistent equations.

Occlusions, non-reflecting surfaces, and the scanning geometry cause missing data. Thus, model assumptions are required to obtain solutions. The proposed procedure relies on the assumptions for non-observed vertical walls, vertical step edges, and horizontally running eaves. Furthermore, we apply so-called manifold constraints to obtain “watertight” model instances: each edge is adjacent to exactly two planar faces.

The usability and feasibility of the approach have been demonstrated by processing an airborne captured laser scan with approximately 1.5 billion points, covering an area of approximately 50 km^2. The scene contains thousands of buildings in inner-city blocks and in rural areas, featuring a wide variety of roof shapes.

Outlook. Especially the reconstruction of small roof structures, such as dormers, is challenging, as only a few points are available and locally used. In this context, enabling generalization is a major unsolved problem. Furthermore, small buildings with flat roofs, e.g., garages, are difficult to reconstruct, because their outlines can currently be inferred only from triangulation boundaries with a point decimation algorithm. In this case, a separate 2D reasoning process looking for orthogonality and parallelism would be advisable.

In rare cases, the adjustment of planes can change the topology of edges and vertices obtained by the intersection of these planes. This can lead to inconsistencies, especially in the form of overlapping polygon edges.

Last but not least, the quality of the reconstruction depends on the proper classification and clustering of the laser points.

Acknowledgment

The author thanks the anonymous reviewers for their careful work and valuable recommendations, which certainly improved the manuscript.

References

Awange, J. L., Grafarend, E. G., 2005. *Solving Algebraic Computational Problems in Geodesy and Geoinformatics*. Springer.

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.

Douglas, D., Peucker, T., 1973. Algorithms for the reduction of the number of points required to represent a digitized line or its caricature. *The Canadian Cartographer*, 10(2), 112–122.

Edelsbrunner, H., Kirkpatrick, D. G., Seidel, R., 1983. On the shape of a set of points in the plane. *IEEE Transactions on Information Theory*, 29(4), 551–559. https://doi.org/10.1109/TIT.1983.1056714.

Förstner, W., Wrobel, B. P., 2016. *Photogrammetric Computer Vision*. Springer.

Haala, N., Kada, M., 2010. An update on automatic 3D building reconstruction. *ISPRS Journal of Photogrammetry and Remote Sensing*, 65(6), 570–580. https://www.sciencedirect.com/science/article/pii/S0924271610000894.

Heuel, S., Kolbe, T., 2001. Building reconstruction: The dilemma of generic versus specific models. *KI – Künstliche Intelligenz*, 3, 75–62.

Koch, K.-R., 1999. *Parameter Estimation and Hypothesis Testing in Linear Models*. 2nd edn, Springer.

Mäntylä, M., 1987. *An Introduction to Solid Modeling*. Computer Science Press, Inc., New York.

Meidow, J., Hammer, H., 2016. Algebraic reasoning for the enhancement of data-driven building reconstructions. *ISPRS Journal of Photogrammetry and Remote Sensing*, 114, 179–190. https://10.1016/j.isprsjprs.2016.02.002.

Meidow, J., Hammer, H., 2025. Minimal rotations in arbitrary dimensions with applications to hypothesis testing and parameter estimation. *ISPRS Open Journal of Photogrammetry and Remote Sensing*, 15, 100085. https:/doi.org/10.1016/j.ophoto.2025.100085.

Meidow, J., Hammer, H., Pohl, M., Bulatov, D., 2016. Enhancement of generic building models by recognition and enforcement of geometric constraints. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, III-3, 333–338. https://isprs-annals.copernicus.org/articles/III-3/333/2016/.

Nan, L., Wonka, P., 2017. PolyFit: Polygonal Surface Reconstruction from Point Clouds. *2017 IEEE International Conference on Computer Vision (ICCV)*, 2353–2361. https://doi.org/10.1109/ICCV.2017.258.

OGC, 2021. OGC City Geography Markup Language (CityGML) Part 1: Conceptual Model Standard. Version: 3.0.0, http://www.opengis.net/doc/IS/CityGML-1/3.0.

Ramer, U., 1972. An iterative procedure for the polygonal approximation of plane curves. *Computer Graphics and Image Processing*, 1(3), 244–256.

Ren, J., Zhang, B., Wu, B., Huang, J., Fan, L., Ovsjanikov, M., Wonka, P., 2021. Intuitive and Efficient Roof Modeling for Reconstruction and Synthesis. *ACM Transactions on Graphics*, 40(6), 1–17. https://doi.org/10.1145/3478513.3480494.

Schnabel, R., Wahl, R., Klein, R., 2007. Efficient RANSAC for point-cloud shape detection. *Computer Graphics Forum*, 26(2), 214–226.

Xiong, B., Elberink, S. O., Vosselman, G., 2014. Building modeling from noisy photogrammetric point clouds. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, II-3, 197–204.

Zhang, W., Qi, J., Wan, P., Wang, H., Xie, D., Wang, X., Yan, G., 2016. An Easy-to-Use Airborne LiDAR Data Filtering Method Based on Cloth Simulation. *Remote Sensing*, 8(6), article 501.