

Automatic reconstruction of fully volumetric 3D building models from point clouds

Sebastian Ochmann^{a,*}, Richard Vock^a, Reinhard Klein^a

^aUniversity of Bonn, Institute of Computer Science II, Endenicher Allee 19a, 53115 Bonn, Germany

ARTICLE INFO

Article history:

Received April 10, 2019

Keywords: Indoor Building Reconstruction, Point Cloud Processing, Integer Linear Programming, Building Information Modeling

ABSTRACT

We present a novel method for reconstructing parametric, volumetric, multi-story building models from unstructured, unfiltered indoor point clouds by means of solving an integer linear optimization problem. Our approach overcomes limitations of previous methods in several ways: First, we drop assumptions about the input data such as the availability of separate scans as an initial room segmentation. Instead, a fully automatic room segmentation and outlier removal is performed on the unstructured point clouds. Second, restricting the solution space of our optimization approach to arrangements of volumetric wall entities representing the structure of a building enforces a consistent model of volumetric, interconnected walls fitted to the observed data instead of unconnected, paper-thin surfaces. Third, we formulate the optimization as an integer linear programming problem which allows for an exact solution instead of the approximations achieved with most previous techniques. Lastly, our optimization approach is designed to incorporate hard constraints which were difficult or even impossible to integrate before. We evaluate and demonstrate the capabilities of our proposed approach on a variety of complex real-world point clouds.

This is an early preprint version. The final version, which was accepted for publication in the ISPRS Journal of Photogrammetry and Remote Sensing, is available at

<https://doi.org/10.1016/j.isprsjprs.2019.03.017>

1. Introduction

The challenging problem of generating high-quality, three-dimensional building models from point cloud scans has been approached in a variety of ways in recent years by the computer graphics and remote sensing communities as well as in the architecture domain. Especially for various applications in Computer Aided Design (CAD) and emerging fields such as Building Information Modeling (BIM), the reconstructed models are usually required to adhere to industry-standard specifications such as the Industry Foundation Classes (IFC). In contrast to the representation of a building in the form of e.g. an unordered point cloud, a set of unconnected surfaces, or boundary meshes, a BIM/IFC model closely resembles the physical building structure by defining buildings as semantically annotated, volumetric building entities such as walls and floor slabs, usually including additional information regarding how these elements are interconnected.

Most previous approaches focus on the reconstruction of completely separate, planar surfaces without additional information regarding how they relate to each other [1], or on representing buildings as watertight boundaries of either the whole building [2] or separate rooms [3, 4] and are thereby lacking in providing insights into the building structure. Also, assumptions such as the one that stories can be globally separated by horizontal planes are very limiting in practice. None of these approaches yields a representation which enables unhindered usage in the aforementioned scenario. While one recent approach [5] does model the measured point cloud data using volumetric building entities, the method is restricted to single-story buildings which limits its usability without laborious manual separation of the point cloud data into separate stories. Additionally, generation of resulting wall and floor slab elements is done in a post-processing step without being integrated into the used optimization framework which may result in locally implausible results. Other methods [6, 7] aiming at reconstructing true BIM models make the severe assumption that walls are positioned in a Manhattan world constellation which is often violated by real-world buildings.

Our proposed method overcomes limitations of previous approaches by alleviating the requirements on the input data and by providing a flexible optimization framework for indoor building reconstruction. Some prior methods (e.g. [4, 5]) require separate scans and scan positions to derive an initial,

*Corresponding author: e-mail: ochmann@cs.uni-bonn.de

coarse segmentation into rooms. In contrast, our fully automatic room segmentation approach, by design, does not depend on the availability of such information and does not impose particular rules for scanning (e.g. one scan per room). Furthermore, our novel integer linear programming formulation for the reconstruction problem provides flexible means to steer the reconstruction process while globally constraining the solution space to feasible solutions, thus guaranteeing a plausible model. Additional information such as manually augmented hints may optionally be incorporated by means of hard constraints in order to further guide the reconstruction process. While some previous approaches regularize the resulting model based on room boundary complexity, they fail to account for dependencies between surfaces related by volumetric wall elements, e.g. opposing surfaces between neighboring rooms. Our formulation of the solution space based on volumetric entities enables better regularization of the model with respect to the actual volumetric walls and slabs used to represent the building. In contrast to any previous approach, the result of our optimization immediately yields the complete geometry of all reconstructed walls and slabs as well as their volumetric intersections which allows for a direct generation of plausible BIM/IFC models.

In summary, the main features of our approach are:

1. Fully automatic, volumetric reconstruction including volumetric intersections between elements.
2. Flexible integration of constraints to enforce global and local properties of the resulting model.

Our main technical contributions are:

1. Automatic filtering of outliers and room segmentation of unstructured, multi-story 3D point clouds.
2. A new formulation of the indoor reconstruction task as a linear integer programming problem that can be efficiently solved using off-the-shelf software.

2. Related Work

Research on scan-to-BIM and related approaches led to a wide range of developments in recent years and still is a current topic of ongoing work. We first provide a comprehensive overview of methods dealing specifically with indoor building reconstruction which we then complement with a summary of more loosely related but complementary abstraction approaches and applications.

2.1. Indoor building reconstruction

The works presented in this section are closely related to our goal of indoor building reconstruction. Table 1 summarizes and compares key features of different approaches.

Some methods aim at the generation of 2D floor plans. Okorn et al. [16] model 2D floor plans by projecting detected structures into the horizontal plane and performing wall segment detection based on the Hough transform. Ambruş et al. [13] reconstruct floor plans including a room labeling obtained using an energy minimization approach. A deep neural architecture for automatic floor plan generation from RGBD video has been

presented by Liu et al. [7]. Using pixel-wise predictions of floor plan geometry and semantics, integer programming [17] is used to recover a vector graphics reconstruction.

Some approaches perform a reconstruction of individual rooms. Budroni et al. [8] reconstruct closed boundary representations of single rooms using plane sweep surface detection and a 2D line arrangement with a split-and-merge approach. The methods by Adán et al. [9] and Xiong et al. [10] focus on recovering detailed surface labelings, explicitly reasoning about occlusions using a ray-tracing approach. In a similar spirit, Previtali et al. [12] perform a reconstruction of single rooms as polyhedral models including ray-tracing based reasoning about occlusions and opening detection.

Certain methods aim at the reconstruction of the building as a whole without explicitly considering room topology or segmentation. Sanchez et al. [1] represent buildings as polygonal surface models including detection of smaller-scale structures such as parametric staircases. Oesau et al. [18] use a 2D cell decomposition to perform binary inside/outside labeling using a Graph-Cut based optimization. The detail level of this approach is enhanced by Oesau et al. in [2] by means of an improved line detection strategy. With a similar goal of providing simplified environment maps for e.g. navigation, Xiao et al. [19] employ constructive solid geometry (CSG) operations to generate a volumetric wall model. Room topology is not explicitly modeled.

Many recent methods approach the reconstruction problem in a 2.5D setting, including a segmentation into separate rooms. Mura et al. [20, 11] model buildings as 2.5D polyhedral meshes by means of constructing a 2D line arrangement and performing k -medoid clustering based on diffusion embeddings. Mura et al. [21] also propose a related approach which allows arbitrary wall orientations and performs recursive clustering on a constrained Delaunay tetrahedralization. The method by Turner et al. [22] provides efficient means to generate 2.5D, textured meshes for e.g. navigation purposes including a room segmentation obtained by Graph-Cut in a triangulated environment map. An extension providing enhanced texture mapping has been presented in [3]. The reconstruction method by Wang et al. [15] models outer and inner walls by means of 2D line arrangements labeled using diffusion embeddings similar to [11]. They also reconstruct doors using a simulated ray casting approach. Murali et al. [6] present a system to quickly generate BIM models from mobile devices such as Google Project Tango, Microsoft Kinect or Microsoft HoloLens, including semantic annotations and relations between reconstructed elements. The approach is currently limited to single-story, Manhattan world buildings.

Few approaches consider the more general case of slanted walls or ceilings. Mura et al. [4] reconstruct polyhedral room boundaries with arbitrary wall and ceiling orientations. Early rule-based classification of detected elements helps pruning invalid parts. Room segmentation is performed by clustering steered by visible surface overlap. Mura et al. [23] propose an extension using automatically clustered synthetic viewpoints and show applicability on complex multi-story buildings.

None of the aforementioned methods reconstruct volumetric wall and slab elements which are directly usable in a BIM setting. Few methods have approached this problem before. Stam-

	No scan positions	Non-Manhattan	Multiple rooms	Full 3D recons. ¹	Slanted ceilings	Volumetric walls
Budroni '10 [8]	✓	✗	✗	✗	✗	✗
Adán '11 [9]	✗	✓	✓	✗	✗	✗
Xiong '13 [10]	✗	✓	✓	✗	✗	✗
Mura '14 [11]	✗	✓	✓	✗	✗	✗
Oesau '14 [2]	✓	✓	✗	✗	✗	✗
Previtali '14 [12]	✗	✓	✗	✗	✗	✗
Turner '15 [3]	✗	✓	✓	✗	✗	✗
Mura '16 [4]	✗	✓	✓	✓	✓	✗
Ochmann '16 [5]	✗	✓	✓	✗	✗	✓
Ambruş '17 [13]	✓ ²	✓	✓	✗	✗	✗
Macher '17 [14]	✓	✓	✓	✗	✗	✓ ³
Murali '17 [6]	✓	✗	✓	✗	✗	✓
Wang '17 [15]	✗ ⁴	✓	✓	✗	✗	✗
Ours	✓	✓	✓	✓	✗	✓

Table 1. Feature comparison of recent indoor reconstruction approaches. Notes: ¹Full 3D reconstruction means that arbitrary vertical room constellations including rooms spanning multiple floors are handled correctly. ²Virtual scan positions are automatically estimated. ³Volumetric walls are generated manually in post-processing. ⁴Scanner trajectories are used.

bler et al. [24] aim to generate volumetric 3D building models using learning approaches for the classification and scoring of detected elements, and simulated annealing for optimizing the overall model. The approach makes strong assumptions about the input data, requiring both interior and exterior scans, as well as scanner positions. Thomson et al. [25] generate volumetric walls from point clouds by detecting planes using a RANSAC approach and fitting suitable IFC wall entities to the detected surfaces; room volumes and topology are not explicitly modeled. They also propose a point cloud segmentation scheme based on a corresponding IFC model. A method which explicitly represents buildings as interconnected volumetric wall elements has been presented by Ochmann et al. [5]. They construct a 2D line arrangement of wall center lines representing pairs of opposing wall surfaces and perform a room labeling of the arrangement faces by means of a Graph-Cut based multi-label energy minimization. Multi-story buildings are not supported. Macher et al. [14] propose a semi-automatic reconstruction approach by first segmenting the input data automatically and exporting the result in an interim OBJ format, and subsequently constructing an IFC file with manual intervention in a post-processing step.

To our knowledge, our approach is the first to combine general multi-story, multi-room reconstruction with fully volumetric room and wall entities.

2.2. Abstraction, segmentation, and reconstruction

We now highlight some loosely related approaches which pursue more general or complementary goals which may be beneficial for tackling the reconstruction problem on different levels. Monzpart et al. [26] represent man-made scenes (e.g. buildings) by a regular arrangement of planes, taking into account non-local inter-primitive symmetry relations. Such a regularization may be useful for various arrangement-based reconstruction approaches. A method for reconstructing lightweight, manifold, polygonal boundary models from point clouds has been presented by Nan et al. [27]. They employ an inside/outside labeling approach using binary linear programming. Jung et al. [28] generate watertight floor maps by means of skeletonization in a 2D binary occupancy map with subse-

quent labeling of separate rooms. A 3D room partitioning approach using anisotropic potential fields with subsequent unsupervised clustering has been presented by Bobkov et al. [29]. Pursuing a similar goal, Ochmann et al. [30] perform a segmentation of indoor point clouds into separate rooms using a visibility-based approach. Openings between neighboring rooms are detected to obtain a room connectivity graph. Bassier et al. [31] employ a machine learning approach to classify structural elements such as walls, floors, ceilings, and beams in point cloud data. The method by Liu et al. [17] generates topologically and geometrically consistent floor plans from 2D raster images using an integer programming approach. While the approach is designed to work on 2D data and assumes Manhattan world geometry, the idea to enforce global properties of the resulting model using integer programming is related to our work. Focusing on non-structural elements relevant to BIM models, Adán et al. [32] present an approach for detecting various important entities such as sockets, switches, signs, and safety-related items. While the method by Son et al. [33] does not explicitly model a building's room topology, they detect various important volumetric elements such as walls, slabs, columns and beams, also taking into account material properties and relations between elements.

2.3. Applications

Automated scan-to-BIM methods facilitate a range of diverse applications in different areas such as construction surveillance, facility management, or energy simulations. Garwood et al. [34] propose a framework for storing building geometry in a format suitable for e.g. energy simulation and verification tasks, and highlight the importance of fast, automated methods for obtaining suitable models. Hyland et al. [35] propose the usage of open standards and automatically derived BIM models from measurements for performing automated compliance control by comparing the as-built and as-designed states of buildings. In a similar spirit, O'Keeffe et al. [36] have developed validation approaches for determining and analyzing differences between scans and BIM models. A prototypical approach has been presented by Brodie et al. [37] who propose a cloud-based platform integrating tools for generating models from and val-

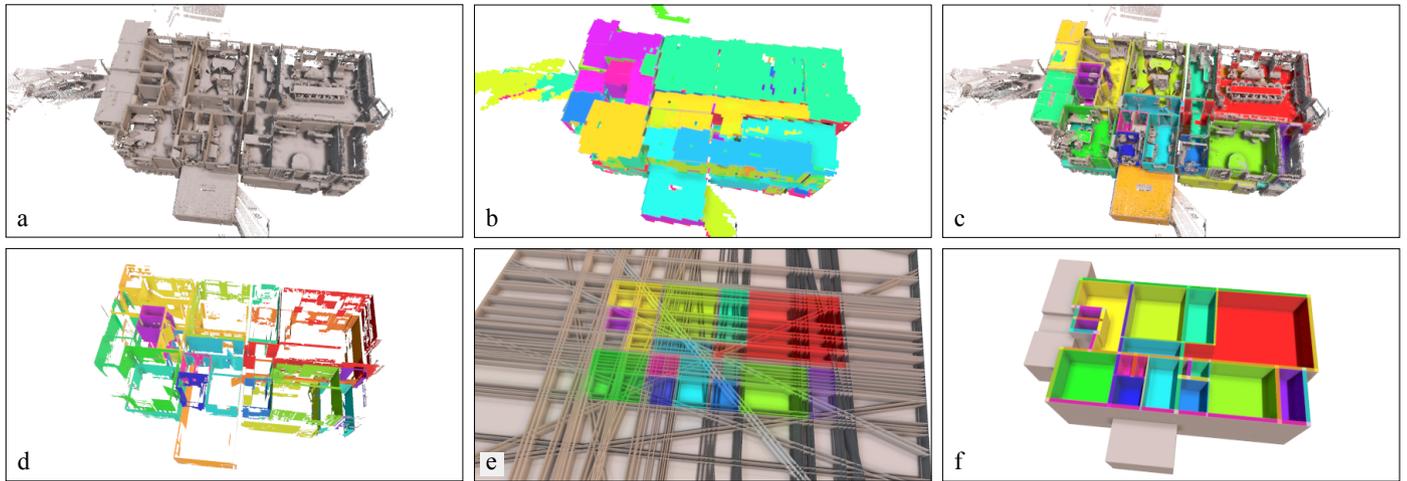


Fig. 1. Overview of the main steps. Ceiling of upper floor is hidden in (a), (c), and (f) for visualization purposes. (a) The input is a registered but otherwise unstructured and unfiltered indoor point cloud. (b) Planes are detected by means of RANSAC shape detection. (c) Outliers are automatically removed and rooms are segmented using an unsupervised clustering approach based on mutual visibility between point patches. (d) Detected planes are classified as horizontal slab surfaces and vertical wall surfaces (only latter shown). Surfaces are assigned multi-label support bitmaps. (e) A 3D plane arrangement is constructed by intersecting all planes, yielding a cell complex. Priors for rooms, outside area and surface support are estimated. (f) The final model consisting of interrelated room and wall volumes is obtained by solving an integer linear program in which cell labels are binary variables.

idating models against point clouds. Krispel et al. [38] developed a method for automatic detection of power sockets and for the generation of hypotheses for electrical lines based on automatically generated building models. An approach for integrating IFC BIM models and point cloud data in a common file format has been presented by Krijnen et al. [39]. They highlight the semantically meaningful association of both worlds for documentation, structuring, annotation, synchronization and retrieval tasks.

3. Overview

The input of our approach is a 3D indoor point cloud (Figure 1 a) with oriented normals whose “up” direction is assumed to be the z -axis. If normals are not yet available, they are estimated by local Principal Component Analysis (PCA).

We first detect planes using an efficient RANSAC implementation [40] (Figure 1 b) and compute occupancy bitmaps for each detected plane from the respective supporting points.

The detected planes are used to automatically eliminate outlier points, and to determine point clusters corresponding to individual rooms. This clustering is performed by means of Markov Clustering [41] which does not require prior information about the number of rooms and results in a labeling of the point cloud (Figure 1 c).

The resulting point labels are projected to the previously detected planes and discretized into multi-label bitmaps. Planes are pruned, rectified, clustered, and classified as candidates for vertical wall or horizontal slab surfaces (Figure 1 d; only vertical surfaces shown for visualization purposes). Since we base our reconstruction on volumetric walls and slabs instead of single surfaces, pairs of nearby, approximately parallel surfaces are grouped to wall and slab candidates.

Based on promising previous approaches (e.g. [4, 23, 5, 13]), we then derive a three-dimensional *arrangement of planes* from the set of wall and slab candidates (Figure 1 e). To this end, all surfaces are interpreted as infinite planes and intersected with each other which results in a segmentation of 3D space into convex polyhedral cells. In particular, each wall and slab candidate is represented by a set of cells located between the respective two candidate surfaces. Priors for the existence of different rooms and wall surfaces are estimated for each 3D cell and 2D face using the labeled surface candidates.

The main step of our approach is to find a *labeling* of all cells such that each cell is either assigned to a *room*, or *outside space*. Additionally, volumetric *walls* must be placed wherever a transition between inside and outside space takes place which is also modeled as part of the labeling problem. The labeling should faithfully conform to the measured data and simultaneously fulfill certain constraints (e.g. wall connectivity) to ensure a plausible resulting model (Figure 1 f).

Formulating this task as an optimization problem requires three parts: First, we define a space of possible solutions with meaningful priors to guide the solver. The geometry of this space is given by the arrangement of planes. Priors for locations of rooms and walls in the cell complex are derived from the measured data. Second, we need to define constraints to restrict the feasibility of a solution. They enforce that any solution satisfies predefined rules, e.g. a room and outside space must be separated by a wall. Third, an objective function for assessing the quality of a solution is formulated as a cost function which is minimized under the given constraints.

After a solution is found it can easily be converted into a format suitable for rendering or exporting, e.g. an IFC file or a mesh, by considering the cell labeling and the boundaries between differently labeled cells.

4. Method

In this Section, we provide details regarding each of the steps involved in our approach with a focus on the formulation as an optimization problem.

4.1. Plane detection

Based on the widely used assumption that the coarse geometry of most buildings can be represented (or sufficiently approximated) by piecewise planar surfaces, a crucial first step of our approach is the detection of planes in the point cloud data. To this end, an efficient RANSAC approach [40] implemented in CGAL [42] is used. The most important parameters are maximum point-to-plane distance, normal angle threshold, minimum number of supporting points per plane, and the probability to miss the largest plane candidate. These can usually be chosen depending on point cloud data quality and used for a wide variety of datasets with similar characteristics (e.g. scanner type, density, noise level). The supporting points of each plane are projected into occupancy bitmaps on the respective plane (Figure 2 b), yielding a discretized approximation of support by measured points. Planes with low support area (estimated using the occupancy bitmaps) are pruned later (Section 4.4). Since the relatively coarse occupancy bitmaps are independent of the point cloud density, the minimum number of points for detecting a plane of the RANSAC algorithm may be set relatively low to cope with lower-resolution point clouds.

4.2. Point cloud cleaning

Real-world point clouds often contain large amounts of outlier points, often due to outside areas scanned through openings. In order to prune outlier points early in the process, we employ a simple but very effective ray casting approach similar to [5]. From each point p , n stochastically sampled rays r_i , $i = 1, \dots, n$, are cast into the hemisphere oriented into the direction of the normal at point p . Ray casting is performed against the occupancy bitmaps of the previously detected primitives. Let $h(r_i)$ be a *hit* function which is 1 if some surface was hit, and 0 otherwise. We approximate the probability that p lies inside of the building as $in(p) = \frac{1}{n} \sum_{i=1}^n h(r_i)$. If $in(p)$ is below a given threshold (in our experiments 0.5), p is removed from the point cloud and the occupancy bitmaps of the planes are updated. This process is iterated a small number of times.

4.3. Point cloud labeling

Priors for the locations of rooms and outside area in three-dimensional space are vital for the later optimization step, even if they are coarse estimations. We formulate the estimation of priors as a point cloud labeling problem where each label represents either a room, or the outside area.

Our proposed automatic labeling approach is based on the idea that regions of the point cloud with high mutual visibility form clusters which correspond to rooms of the building. We implement this by performing visibility tests by means of ray casting between point patches on detected surfaces which yields a visibility graph. Nodes of this graph are then clustered

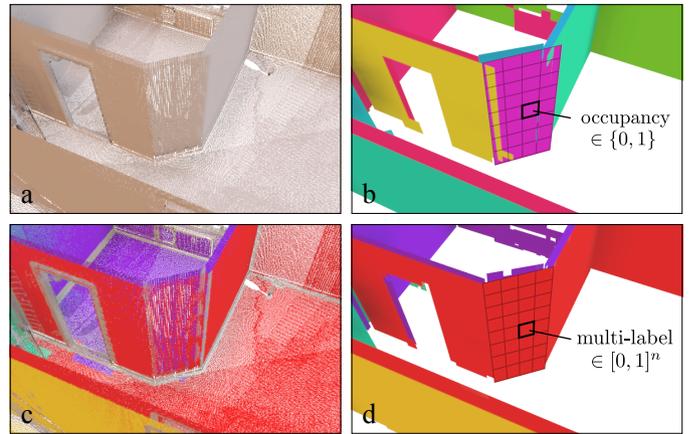


Fig. 2. Detected planes are the basis for our reconstruction. Different kinds of bitmaps (i.e. grids) on the planes are used throughout the approach. (a) Unlabeled input point cloud. (b) Binary occupancy bitmaps on the detected planes (Section 4.1) are used as a lightweight representation of support by the point cloud. Different planes are shown in different colors. (c) Ray casting against the occupancy bitmaps and clustering yields a segmentation of the point cloud into rooms (Section 4.3). Different room labels are shown in different colors. (d) The point cloud labeling is projected into multi-label bitmaps where each pixel contains a soft assignment $[0, 1]^n$ to the n different room labels as defined in Section 4.3. This is used for estimating locations of rooms in 3D space as described in Section 4.7. The multi-label bitmaps are shown using the same label colors as in (c). In (b) and (d), only vertical planes are shown for clarity.

by means of the Markov Clustering algorithm [41] which determines natural clusters within the graph by flow simulation.

Point patches are constructed by generating coarse occupancy bitmaps for each plane and considering each occupied pixel as a patch with a normal identical to the respective plane normal. In our experiments, a patch size of $40\text{cm} \times 40\text{cm}$ was used. We use patches instead of all points to drastically reduce the number of nodes in the visibility graph which makes the computation feasible. Let p_i be the i -th patch with center position c_i and normal n_i . For each pair p_i, p_j , $i \neq j$, ray casting between the points $c_i + \varepsilon n_i$ and $c_j + \varepsilon n_j$, with $\varepsilon := 10\text{cm}$ in our experiments, is performed. If no surface is hit, the visibility between p_i, p_j is set to 1, otherwise it is set to 0. This yields a visibility graph whose nodes are clustered using the Markov Clustering algorithm. The computed visibility is interpreted as flow between node pairs corresponding to the respective point patches. The main advantage of this method is that it is unsupervised and thus does not require a manual specification of the number of occurring labels.

As a result, we obtain n disjoint clusters of patches which belong to different rooms and define the set of room labels $\mathcal{R} := \{r_1, \dots, r_n\}$ which will be used throughout the remainder of the reconstruction process. Each point of the point cloud is assigned the room label of the respective point patch. Note that the number of room labels n may be larger than the number of rooms that will actually be contained in the final reconstruction.

4.4. Surface candidates

The detected planes usually include many surfaces which are not part of walls, floors and ceilings. Even correctly detected surfaces will generally not be perfectly vertical or horizontal.

We thus apply a pruning, classification and rectification step to extract two sets of candidates for wall and slab surfaces. The occupancy bitmaps are used to estimate the support area of each surface independently of point cloud density. Planes with support below an area threshold as well as planes which are not approximately vertical or horizontal are discarded. The remaining planes are classified as wall or slab surface candidates depending on their normal direction, and adjusted to be perfectly horizontal or vertical.

As a prerequisite for later room prior estimation (Section 4.7), each surface is also assigned a multi-label support bitmap with continuous values in $[0, 1]$ for each room label in \mathcal{R} (Figure 2 d). This provides a soft-assignment of different regions of each surface to different room labels. The label bitmaps are generated by projecting all supporting points onto the respective surface and averaging the previously determined point labels within each pixel.

Furthermore, we *dilate* the support bitmaps. The rationale is that reconstructed walls with no surface support by the point cloud data are penalized by a cost function defined later in Section 4.8. Since we reconstruct wall intersections volumetrically, placing the respective wall entities in between rooms would cause high costs since surface support is naturally restricted to regions that are visible to the scanner (Figure 3, middle row). By slightly extending the surface support, we encourage construction of intersecting wall entities in regions with nearby surface support (Figure 3, bottom row).

4.5. Wall and slab candidates

Since our approach is based on the notion of *volumetric* walls and slabs instead of single surfaces, the next step is to determine pairs of opposing surfaces forming potential building elements. To this end, a simple pairing procedure is employed. For each surface, we search a matching, approximately parallel surface with opposing normal orientation within a user-defined distance

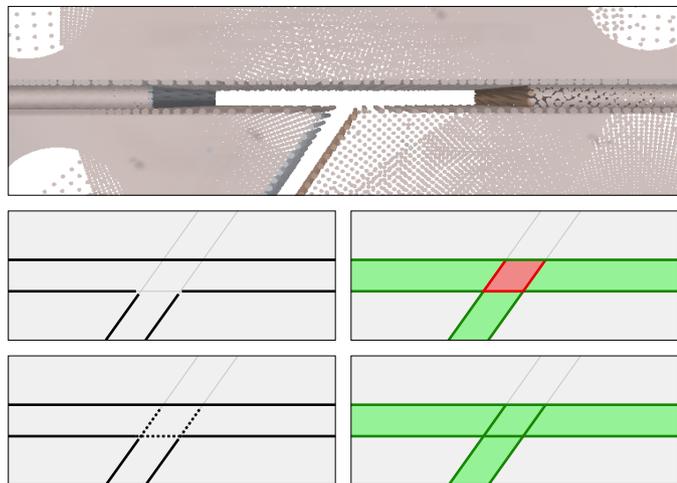


Fig. 3. Dilation of surface support. **Top:** Example point cloud viewed from above. **Middle:** Surface support (left, thick lines) is naturally restricted to parts visible to the scanner. This leads to high costs for the wall intersection since wall surfaces are not supported by points (right). **Bottom:** Dilating the surface support (left, dotted lines) extends support into the interiors of walls, encouraging placement of volumetric intersections (right).

and angle threshold. If a match is found, the two surfaces are paired to form a wall or slab candidate. It should be noted that a single surface may thus be part of multiple pairs. For surfaces without any matching counterpart, *virtual* surfaces with a user-defined distance are added to the set of surfaces. This is usually the case for outside walls for which only the inner side has been scanned. This augmentation is important since interior spaces are required to be bounded by volumetric walls or slabs. The m generated candidates constitute the set of wall/slab labels $\mathcal{W} := \{w_1, \dots, w_m\}$. Which of these walls and slabs are contained in the final model is decided by the optimization described in Section 4.8.

4.6. Arrangement of planes

The geometry of the search space for finding an optimal constellation of rooms, walls and slabs is modeled as an *arrangement of planes* and the *3D cell complex* induced thereby. It is constructed by intersecting all (infinite) planes of the wall and slab candidate surfaces with each other. Since vertical walls and horizontal slabs are treated identically, we will hereafter refer to both simply as *walls*.

Cells of the arrangement are convex, three-dimensional subsets of the space inside and outside of the building. Each cell belongs either to a room, or the outside area. Additionally, walls may be placed in cells that are part of the outside area. Constraints such as that a cell may belong to at most *one* room, or that walls may *only* occur in the outside area (e.g. between rooms) are formulated as constraints for the optimization problem in Section 4.8.

Faces between neighboring cells are convex, two-dimensional subsets of regions on the planes of wall surfaces. Each face may separate different regions (e.g. a room and a wall) from each other.

4.7. Volume and surface priors

For guiding the optimization, two kinds of priors are estimated from the data. First, volumetric priors for the existence of different rooms as well as outside area are estimated for each 3D cell of the arrangement. Second, support by the point cloud data is estimated for each 2D face between neighboring cells.

Preparations. The arrangement consists of cells $\mathcal{C} := \{c_1, \dots, c_p\}$. For two cells, the notation $c_a + c_b$ means that c_a, c_b are neighboring and the normal of the separating oriented face f_{c_a, c_b} points towards c_a (Figure 4 a). The set of all oriented faces is denoted as $\mathcal{F} := \{f_{c_a, c_b} \mid c_a + c_b\}$. For brevity, we write f instead of f_{c_a, c_b} if the specific incident cells are irrelevant.

The set of room labels is $\mathcal{R} := \{r_1, \dots, r_n\}$ with n being the number of room clusters as introduced in Section 4.3, and the set of wall labels is $\mathcal{W} := \{w_1, \dots, w_m\}$ with m being the number of generated wall candidates as introduced in Section 4.5. We furthermore define an additional *outside* label $\mathcal{O} := \{o\}$. As detailed later, the outside label is used for cells that are not the interior space of a room. The union of rooms and outside labels is denoted $\mathcal{R}_o := \mathcal{R} \cup \mathcal{O}$; the set of all labels is $\mathcal{L} := \mathcal{R} \cup \mathcal{O} \cup \mathcal{W}$.

Let $\mathcal{C}_w \subseteq \mathcal{C}$ be the set of cells that are *contained* in wall candidate w , i.e. all cells that are located between the two surfaces

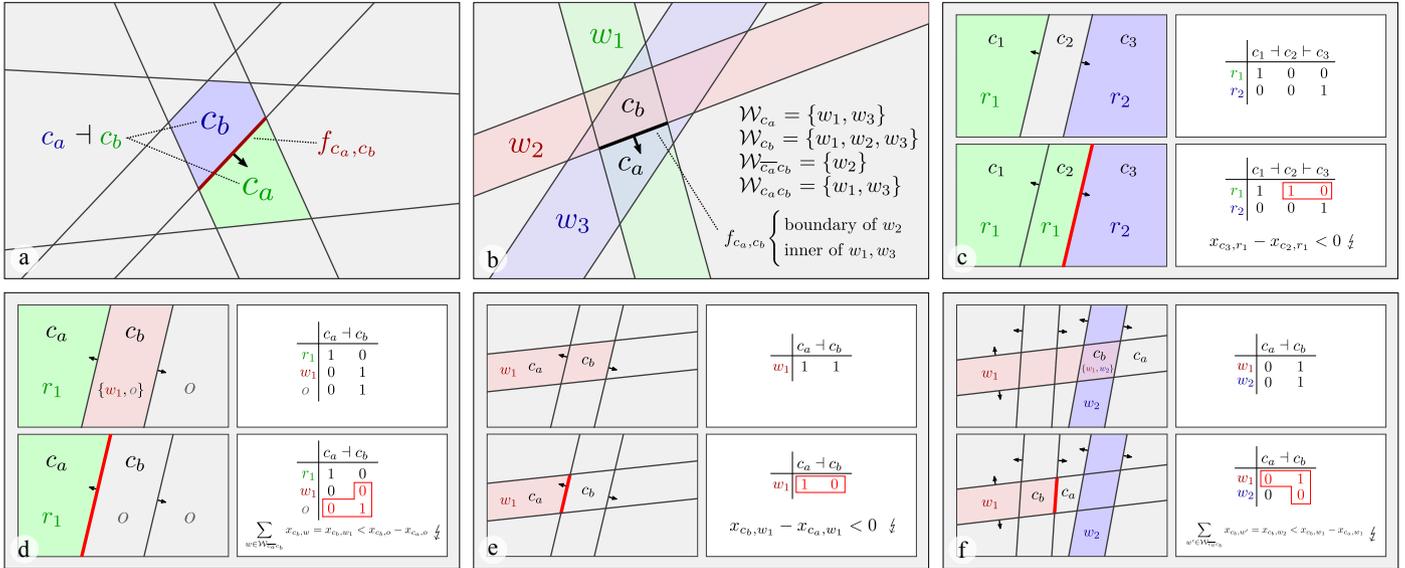


Fig. 4. Explanation of notation and constraints. (a) Neighboring cells are considered as ordered pairs $c_a \dashv c_b$ with respect to normal orientation of the separating face f_{c_a, c_b} . (b) Considering a cell c , \mathcal{W}_c is the set of wall candidates enclosing cell c . A face f_{c_a, c_b} may be a boundary or an inner face of a set of walls. The set of boundary walls with respect to that face is $\mathcal{W}_{c_a c_b}^-$, the set of inner walls is $\mathcal{W}_{c_a c_b}$. (c) Transitions between interior and exterior area at a face f may only occur with the room label being on the positive side of f (Constraint 2). (d) If a face f is the boundary of a room, some walls need to be active on the negative side of f (Constraint 4). (e) A wall may end at an inner face f only if the wall label is on the negative side of f (Constraint 5). (f) If a wall w_1 ends at an inner face, this face must be a boundary face of at least one other active wall w_2 . This enforces connectedness between walls.

of w . Conversely, $\mathcal{W}_c := \{w \in \mathcal{W} \mid c \in C_w\}$ is the set of walls that *contain* cell c .

For a particular cell pair c_a, c_b , we define the set $\mathcal{W}_{c_a c_b}^-$ of walls that are contained in cell c_b but *not* in c_a , i.e.

$$\mathcal{W}_{c_a c_b}^- := \mathcal{W}_{c_b} \setminus \mathcal{W}_{c_a}.$$

The separating face f_{c_a, c_b} is called a *boundary* face of the walls in $\mathcal{W}_{c_a c_b}^-$. Analogously, we define the set $\mathcal{W}_{c_a c_b}$ of walls that are contained in *both* c_a and c_b , i.e.

$$\mathcal{W}_{c_a c_b} := \mathcal{W}_{c_a} \cap \mathcal{W}_{c_b}.$$

The separating face f_{c_a, c_b} is called an *inner* face of the walls in $\mathcal{W}_{c_a c_b}$. These definitions are exemplified in Figure 4 b.

It should be noted that an inner face of a wall w is always the boundary face of another wall (which is often approximately perpendicular to w). As an example, in Figure 4 b, face f_{c_a, c_b} is an inner face of wall w_1 and a boundary face of the intersecting wall w_2 . This will become important for the definition of the optimization constraints in Section 4.8.

Room and outside priors. To estimate probabilities where different rooms and outside area are located in 3D space, we estimate a volumetric prior function $p_C(c, l) : C \times \mathcal{R}_o \rightarrow [0, 1]$ which returns a high value iff a label l is likely to occur within a cell c . To this end, we perform stochastic ray casting from sampled points in 3D space and average previously computed room labels on surfaces visible from each point. For each cell c , k random points are sampled within c . To draw enough samples for narrow cells, which are very common due to parallel surfaces, k is chosen proportional to

$$\max(\text{volume}(c), \text{diameter}(c)).$$

Centered at each sampled point, d rays are cast into random directions. $p_C(c, l_i)$, $i = 1, \dots, n$, is then set to the average over all observed room labels. Rays hitting the back side of surfaces, as well as rays without surface intersections, are counted as *outside*.

Face support priors. In addition to the volumetric room and outside prior function, we estimate a face support function $p_{\mathcal{F}}(f) : \mathcal{F} \rightarrow [0, 1]$ which returns a high value iff a face f is supported by the point cloud. This function is later used for selecting probable wall candidates and regularizing the optimization result. To estimate $p_{\mathcal{F}}(f)$ for a face f , we first sample k random points within f where k is proportional to

$$\max(\text{area}(f), \text{diameter}(f)).$$

Subsequently, all sampled points are projected onto the surface from which face f was generated in the arrangement. $p_{\mathcal{F}}(f)$ is then set to the ratio between the number of sampled points lying within the support approximated by the occupancy bitmap of the respective surface to the total number of sampled points.

4.8. Cell complex optimization

For finding an optimal labeling of all cells, we employ a 0-1 integer linear programming approach in which binary variables for each cell are interpreted as room, outside, and wall label assignments to cells. This approach has the advantage that a set of rules to be fulfilled by any feasible solution can be formulated as hard constraints. Approximate multi-label methods based on e.g. Graph Cuts [43] are more restricted regarding the family of objective functions and constraints that can be used and may fail to find good solutions if the objective is not sufficiently smooth. We first discuss the set of constraints imposed on our model before defining the objective function.

Preparations. Each of the binary variables

$$x_{c,l} \in \{0, 1\}, \quad c \in C, \quad l \in \mathcal{L},$$

of our optimization is a binary assignment of a label l to a cell c . A value of 1 means that the label is *assigned* or *active*. It should be noted that a cell is not necessarily assigned only a single label. In particular, a cell can be assigned the outside label *and* a nonempty set of wall labels at the same time as defined by the constraints below. Also, cells where walls intersect are assigned *all* labels of the intersecting walls. We also use the notion of *inner* and *boundary* faces as defined in Section 4.7.

Constraint 1. Each cell c must be assigned exactly one label from \mathcal{R}_o , i.e.

$$\forall c \in C : \sum_{r \in \mathcal{R}_o} x_{c,r} = 1. \quad (1)$$

Constraint 2. At boundary faces of room interiors, the room label may only occur on the *positive* side of the separating face, i.e.

$$\forall f_{c_a,c_b} \in \mathcal{F} \quad \forall r \in \mathcal{R} : x_{c_a,r} - x_{c_b,r} \geq 0, \quad (2)$$

as shown in Figure 4 c. Note that this constraint implies that two different room labels $r_p \neq r_q$ cannot be directly neighboring since this would violate the constraint for one of the room labels. As a consequence, this avoids ‘‘paper thin’’ walls between rooms since they must be separated by outside area, thereby following the physical nature inherent to walls.

Constraint 3. Wall labels may only occur in cells which are assigned the outside label, i.e.

$$\forall c \in C \quad \forall w \in \mathcal{W}_c : x_{c,w} \leq x_{c,o}. \quad (3)$$

Constraint 4. The boundary faces of room interiors must also be the boundary faces of an active wall, i.e.

$$\forall f_{c_a,c_b} \in \mathcal{F} : \sum_{w \in \mathcal{W}_{c_a c_b}} x_{c_b,w} \geq x_{c_b,o} - x_{c_a,o}, \quad (4)$$

as illustrated Figure 4 d. This constraint implies that there cannot be a transition between room interior and outside area without activating a wall at all faces where the transition occurs.

Constraint 5. At wall boundaries which occur at inner faces, the wall label must be on the negative side of the respective faces, i.e.

$$\forall f_{c_a,c_b} \in \mathcal{F} \quad \forall w \in \mathcal{W}_{c_a c_b} : x_{c_b,w} - x_{c_a,w} \geq 0, \quad (5)$$

as exemplified Figure 4 e. This constraint is a prerequisite for Constraint 6 as well as the objective function which require the left-hand side expression to be nonnegative.

Constraint 6. A wall may end at an inner face only if this face is a boundary face of at least one other active wall, i.e.

$$\forall f_{c_a,c_b} \in \mathcal{F} \quad \forall w \in \mathcal{W}_{c_a c_b} : \sum_{w' \in \mathcal{W}_{c_a c_b}} x_{c_b,w'} \geq x_{c_b,w} - x_{c_a,w}, \quad (6)$$

as depicted Figure 4 f. This constraint enforces that walls are interconnected at their endpoints since it disallows that a wall ends at an inner face without it coinciding with a boundary face of an active wall.

Objective function. To determine the optimal labeling, we define a cost function F_C for a solution over the for cell complex C of the form

$$F_C := -R_C + \alpha(W_{\mathcal{F}_b} + W_{\mathcal{F}_i}), \quad (7)$$

consisting of the following terms. The volumetric room and outside area fitness term R_C rewards the assignment of the most likely labels for each cell $c \in C$ and is defined as

$$R_C := \sum_{c \in C} \sum_{r \in \mathcal{R}_o} x_{c,r} \cdot p_C(c, r) \cdot \text{volume}(c), \quad (8)$$

where $x_{c,r}$ denotes the binary variable for the assignment of label r to cell c and $p_C(c, r)$ represents the volumetric room and outside prior (Section 4.7), weighted by the volume of cell c . Note that this term is included with a negative sign within F_C such that its value is being maximized. The wall face cost terms $W_{\mathcal{F}_b}$ and $W_{\mathcal{F}_i}$ penalize placement of walls in terms of the required boundary and inner face areas, respectively. This penalty is attenuated for faces with high support. The terms are defined as

$$W_{\mathcal{F}_b} := \sum_{f_{c_a,c_b} \in \mathcal{F}} \sum_{w \in \mathcal{W}_{c_a c_b}} x_{c_b,w} \cdot (1 - p_{\mathcal{F}}(f_{c_a,c_b})) \cdot \text{area}(f_{c_a,c_b}), \quad (9)$$

and

$$W_{\mathcal{F}_i} := \sum_{f_{c_a,c_b} \in \mathcal{F}} \sum_{w \in \mathcal{W}_{c_a c_b}} (x_{c_b,w} - x_{c_a,w}) \cdot (1 - p_{\mathcal{F}}(f_{c_a,c_b})) \cdot \text{area}(f_{c_a,c_b}), \quad (10)$$

where $x_{c_a,w}$, $x_{c_b,w}$ are the binary variables for the assignment of the wall label w to the cells c_a , c_b respectively, $p_{\mathcal{F}}(f_{c_a,c_b})$ is the face support prior (Section 4.7), and $\text{area}(f_{c_a,c_b})$ is the area of face f_{c_a,c_b} . It should be noted that $(x_{c_b,w} - x_{c_a,w}) \in \{0, 1\}$ due to Constraint 5. Also note that in Equation 9, it suffices to consider $x_{c_b,w}$ since for a boundary face of wall w , $x_{c_a,w}$ does not exist (i.e. $x_{c_a,w}$ can be considered to be zero).

We then minimize F_C s.t. Constraints 1-6 using the Gurobi Optimizer [44].

Note that in our experiments, we added the following constraint which gave a small performance improvement although it is already implied by Constraints 1-2. At boundary faces of outside area, the outside label may only occur on the *negative* side of the separating face, i.e.

$$\forall f_{c_a,c_b} \in \mathcal{F} : x_{c_a,o} - x_{c_b,o} \leq 0. \quad (11)$$

We attribute this slight performance improvement to heuristics used by the particular optimizer implementation.

4.9. Optimization result

The result is an assignment of each cell to either one room, or the outside area. Cells which are assigned the outside area may also be assigned a nonempty set of walls. On the one hand this provides a dense segmentation of space into rooms and outside space. Volumes to which multiple walls are assigned are (volumetric) intersections of the respective walls. Since the underlying data structure provides adjacency information between all

cells, semantic information like room adjacency and wall incidence is immediately available, e.g. for navigation or simulation purposes. On the other hand this information is closely related to the definition of building elements in BIM formats like IFC. This enables immediate transfer of the results into standard architecture software and integration into existing BIM pipelines.

5. Implementation details

Input point clouds were subsampled to a minimum point distance of 2 cm. Plane detection was performed using a plane distance threshold of 1 cm, a point cluster epsilon of 20 cm, a normal threshold of about 6° (18° for the “Case study 2” dataset), minimum support of 1000 points and miss probability of 0.001. Multi-label bitmaps had a resolution (pixel size) of 10 cm, occupancy bitmaps had a resolution of 20 cm. Three ray casting iterations were performed for point cloud cleaning. For automatic labeling, MCL was used with default parameters (inflation set to 2.0) in multi-threaded mode. The surface cost weight α in Equation 7 was empirically chosen as 0.04. We used PCL 1.8.1 [45], CGAL 4.12 [42, 46], MCL 14-137 [41], Gurobi 8.0.1 [44], and NVIDIA OptiX 5.0 for GPU-based ray casting under Linux on a 6-core Intel i7 CPU and a NVIDIA GeForce GTX 980 GPU.

6. Evaluation

We evaluate the reconstruction quality and performance of our approach on a variety of datasets and show comparisons with groundtruth IFC and related work. Furthermore, we exemplify the flexibility of our integer linear programming approach by specifying additional constraints to modify and guide the resulting reconstruction in an intuitive manner.

Datasets. We used a variety of real-world datasets and one synthetic dataset for our evaluation. Table 2 shows six multi-story point clouds measured using terrestrial laser scanners. These datasets were provided by The Royal Danish Academy of Fine Arts Schools of Architecture, Design and Conservation (CITA). The Table lists properties of the input data including the number of points and scans, as well as quantities derived during reconstruction such as the number of room labels, extracted surfaces, wall candidates, etc. It also shows runtime measurements of the main processing steps. We also tested our approach on publicly available datasets provided by other research groups. Figure 5 shows the dataset “synth3” by the Visualization and MultiMedia Lab at University of Zurich, Figure 10 depicts the dataset “Case study 2” from the ISPRS Benchmark on Indoor Modeling [47], and Figure 9 shows the dataset “Area 3” from the Stanford 3D Large-Scale Indoor Spaces Dataset [48]. We used the latter two for demonstrating different parameters and interactive modification as described below.

Reconstruction quality. Our reconstruction approach generally worked well on the test datasets without any dataset-specific tuning. Automatic outlier removal reliably ignored even large-scale clutter scanned through windows in e.g. Datasets 1, 3, and 6. In some cases, particularly thick walls (bottom region of

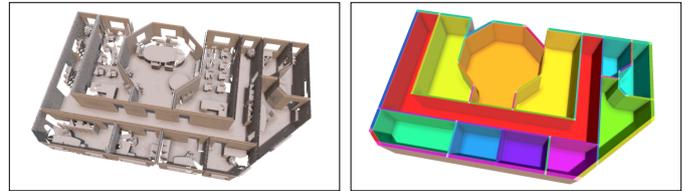


Fig. 5. Our reconstruction result on the synthetic dataset “synth3” by the Visualization and MultiMedia Lab at University of Zurich. Rooms and walls are accurately reconstructed.

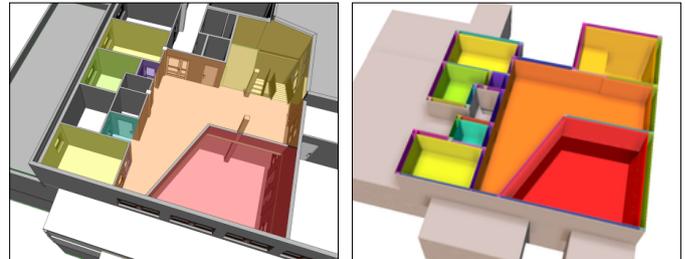
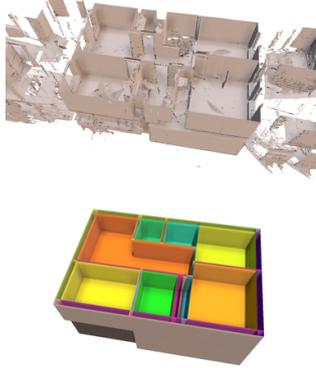
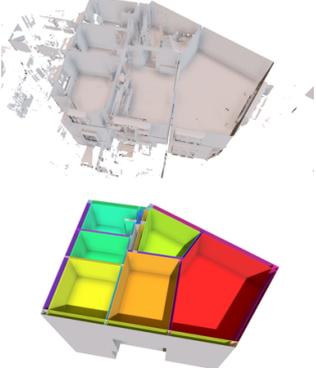
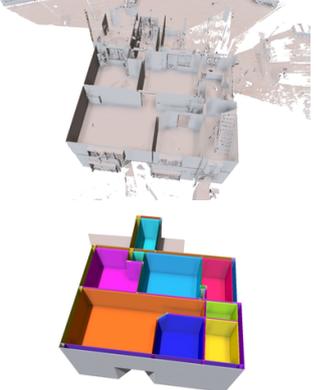


Fig. 6. Comparison of a hand-crafted BIM model (left) and our reconstruction (right) of Dataset 5 (see Table 2). Reconstructed room labels were manually overlaid on the BIM model for reference.

Dataset 1, top region of Dataset 2) were reconstructed as two thinner, parallel wall elements which may be a matter of interpretation. Increasing the maximum thickness of generated wall candidates in these cases can help recognizing such cases as single walls. A few cases of room-oversegmentation can be observed. In Dataset 4, the large central room is split into a larger L-shaped part (orange) and a smaller room (green, to the right of the building) without a real wall separating the reconstructed rooms in the point cloud data. In Dataset 5, indentations of the central room (orange) were reconstructed as small, separate rooms (cyan, purple). Since our approach currently only considers horizontal ceilings, the slanted ceiling of the staircase in Dataset 6 (yellow, elongated room) is reconstructed as a horizontal structure (see also Limitations below).

Runtime. Total runtime for the reconstruction of the test datasets lies in the range of one minute (Datasets 1, 2) to 10 minutes (Dataset 6). The runtime of primitive detection is mainly dependent on the CGAL implementation, and the time for solving the optimization problem is the runtime of the Gurobi optimizer. The runtime for auto labeling contains the time for our raycasting and clustering using the Markov Cluster Algorithm. Runtime of the optimization mainly depends on the complexity of the plane arrangement, which in turn depends on the number of detected surfaces since every surface introduces global splits in the cell complex. Therefore a tradeoff between reconstructing details (i.e. small surfaces) and computational feasibility must be made. In our experiments, we thus chose a minimum estimated area of 2 m^2 for vertical surfaces and 5 m^2 for horizontal surfaces.

Comparison to IFC. For Dataset 5 a corresponding, professionally made BIM model in IFC format was available. Figure 6 shows a comparison between our reconstruction and the BIM

	Dataset 1	Dataset 2	Dataset 3
Input #scans / #points / #pts. cleaned	12 / 3168600 / 2702813	21 / 5151388 / 4723219	29 / 7688111 / 5874557
#Entities Room labels / Surfaces ¹ / Walls ¹ Cells / Variables / Constraints Nonzeros	27 / 42+5 / 37+5 17666 / 594748 / 1775298 4174634	30 / 34+7 / 28+5 12749 / 459373 / 1334699 3134503	39 / 51+5 / 39+4 17334 / 781794 / 2261980 5315578
Runtime (seconds) Plane detection Cleaning (3 iterations) Auto labeling Arrangement + Priors Optimization	18.2 14.1 6.3 14.6 20.9	20.1 21.9 7.2 11.5 7.4	73.9 32.7 11.4 15.5 9.8
			

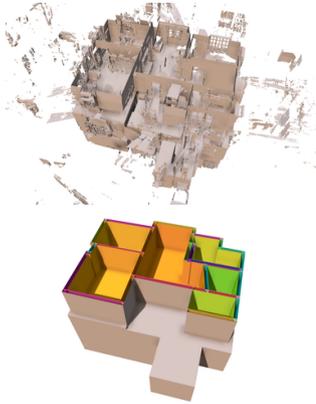
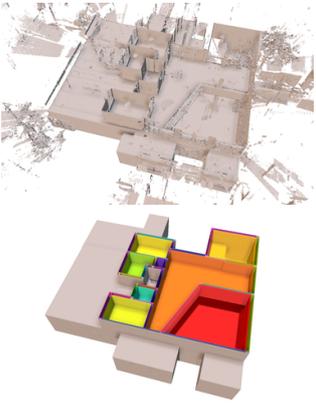
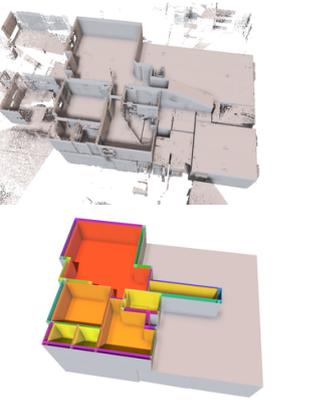
	Dataset 4	Dataset 5	Dataset 6
Input #scans / #points / #pts. cleaned	21 / 6452193 / 5627781	13 / 12409443 / 10688132	39 / 34964707 / 33687751
#Entities Room labels / Surfaces ¹ / Walls ¹ Cells / Variables / Constraints Nonzeros	49 / 61+7 / 58+5 42262 / 2366866 / 7035794 16535885	29 / 48+8 / 46+7 35196 / 1254848 / 3757832 8837356	108 / 72+5 / 66+5 52701 / 6041984 / 17610275 41350862
Runtime (seconds) Plane detection Cleaning (3 iterations) Auto labeling Arrangement + Priors Solving	51.1 28.2 16.1 38.9 85.4	82.6 59.0 56.8 29.1 42.9	218.6 170.9 84.5 53.6 182.8
			

Table 2. Evaluation results on various real-world datasets. The number of scans in the point cloud is shown for reference only and is not used in our approach. The number of points is the point count after subsampling. ¹The number of surfaces and walls is given as *vertical+horizontal*.

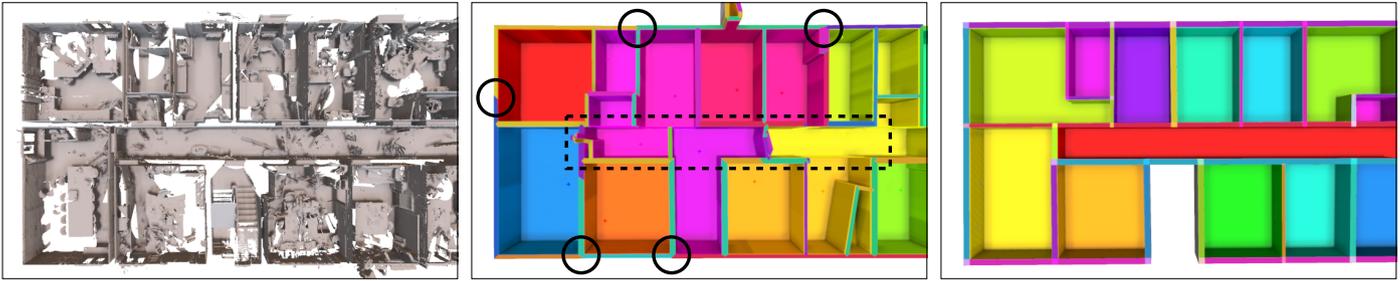


Fig. 7. Comparison between different reconstruction approaches. Left: Input point cloud viewed from above. Center: The method described in [5] may fail to regularize chains of almost coplanar walls, leading to segmented walls (circles). Also, reliance on separate scans as initial room labeling may lead to oversegmented rooms (dashed rectangle). Right: Our approach overcomes these issues by incorporating costs for all surfaces of volumetric wall elements, and room segmentation that is independent of scan positions.

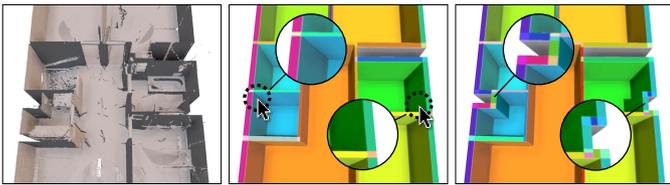


Fig. 8. Additional constraints may be added to interactively steer the reconstruction. Small indentations of wall surfaces (left) are initially lost in the reconstruction. By forcing regions to be outside area (center), our method finds an alternative wall placement under these constraints (right).

model. Colors of the reconstructed rooms were manually overlaid on the IFC model on the left-hand side. All rooms that were part of the scans mostly match the groundtruth BIM. The upper story of the building is connected to the lower story through a large horizontal opening. These areas were reconstructed as two separate rooms (red and orange) and the railing at the edge of the gallery was reconstructed as walls. The small cyan and purple rooms are an oversegmentation of the upper floor, probably due to the dilated surface support. However, this error can easily be fixed manually.

Comparison to related work. A comparison between reconstructions by the approach described in [5] and our method is shown in Figure 7. In addition to fundamental advantages of our approach such as reconstruction of multiple stories, two crucial differences are particularly notable. First, our approach results in stronger regularization of wall elements where using multiple different, similar walls to represent the building would be unnecessary. The approach in [5] leads to jumps between different, almost coplanar walls (Figure 7, center, black circles) instead of using longer, continuous walls. This can be explained by the principle that the approach separated rooms by wall center lines in 2D such that jumping from one wall to an almost coplanar wall resulted in almost no penalty in the cost function. In our case, a fully volumetric wall element would need to be added to the model to connect the parallel walls, resulting in relatively high costs. Second, the approach in [5] relies on given, separate scans and their positions for estimating an initial room segmentation. This leads to an oversegmentation of the hallway (Figure 7, center, dashed rectangle) since it tries to reconstruct one room per scan. Our method works independently of separate scans and estimates a room segmentation by unsupervised

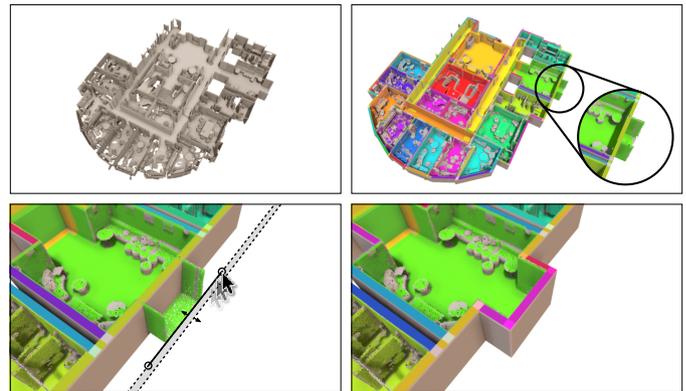


Fig. 9. Manual addition of a wall demonstrated on the dataset “Area 3” from the Stanford 3D Large-Scale Indoor Spaces Dataset. Top right: A hallway ends without a terminating wall surface such that no wall candidate is available for enclosing the protrusion. Reconstruction and point cloud are shown overlaid. Bottom left: A wall candidate can easily be added by drawing a line, adding two opposing “virtual” wall surfaces. Bottom right: The algorithm now encloses the protrusion, using the perpendicular walls with real support in the point cloud.

clustering.

Interactive modification. Our linear programming approach allows for additional constraints to be easily added. One example for manual post-processing of the reconstruction results by interactively adding hard constraints is shown in Figure 8. In this case the indentation of the wall surfaces on the left and right sides of the building were lost by the regularization of the model as can be seen in Figure 8, center. The user has the option to add constraints such as forcing inside area, outside area, wall, no wall, etc. by clicking at the desired location. In this case, the highlighted locations were forced to be outside area. The algorithm then finds the next best option, placing new walls that fulfill all constraints as shown in Figure 8, right. Another example is shown in Figure 9 where a hallway ends without any terminating wall surface in the input data. Since the algorithm has no wall candidate available, it cannot enclose the protruding room area. By adding a virtual wall candidate by means of simply drawing a line, the algorithm is able to include the protrusion in the reconstructed model, automatically using the perpendicular wall surfaces that are present in the input data. Different choices of the wall surface penalty weight α in Equa-

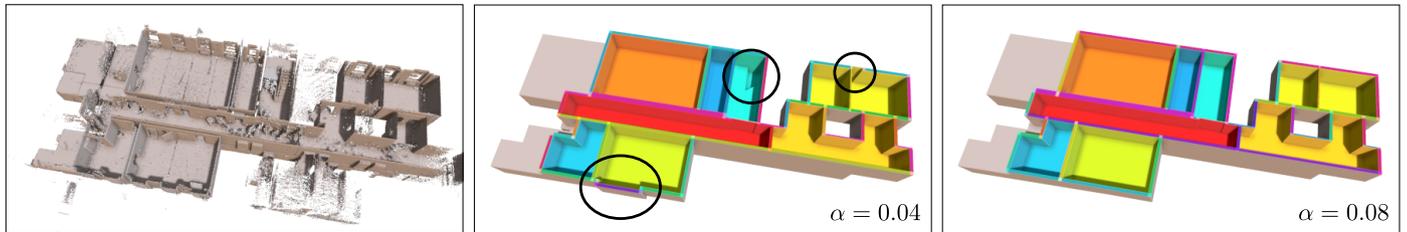


Fig. 10. Different settings for the wall surface cost parameter α in Equation 7 demonstrated on the dataset “Case study 2” from the ISPRS Benchmark on Indoor Modeling. Center: Our default setting of $\alpha = 0.04$ results in some walls to be fitted to windows which have high point support in this dataset. Also, a slab has a hole since floor support in staircases is often complex. Right: Increasing to $\alpha = 0.08$ leads to stronger regularization of walls and slabs.

tion 7 control global regularization strength. Figure 10, center shows a reconstruction where some wall and slab elements are slightly misplaced due to relatively strong surface support at windows. Increasing α from our default of 0.04 to 0.08 leads to stronger regularization as shown in Figure 10, right.

Limitations. One technical limitation of our current implementation is that slanted walls, floors or ceilings are not taken into account although this is not an inherent limitation of our approach. The reason for our decision not to include these elements is that the construction of the 3D cell complex needs to be exact to guarantee the integrity of the data structure (e.g. cell neighborhood). Unfortunately, computing the cell complex in 3D induces numerical problems and currently we do not have a stable implementation for this task at our disposal. We thus opted to use the numerically stable implementation of 2D arrangements in CGAL [46] and extend it to 3D by stacking 2D arrangements separated by horizontal planes. A numerically stable extension of arrangements supporting general slanted planes would be an interesting direction for future research which we consider to be outside the scope of this paper. Processing of very large datasets may also require further optimizations to make them computationally feasible. In particular, using a global plane arrangement results in a large increase of cells and thus variables in the optimization model with every additional detected surface. More sophisticated selection of potential surfaces, and improved optimization methods, e.g. splitting the problem into smaller subproblems, are targets for further research. Last but not least, our current algorithm is not able to identify and include important architectural structures overarching the whole building like the pillars that are included in the hand-crafted model in Figure 6. Automatically identifying such structural elements and incorporating them into the automatic reconstruction is also an interesting direction for future research.

7. Conclusion and future work

We have presented a novel approach to tackle the indoor building reconstruction problem from point clouds using integer linear programming. In contrast to previous methods, our approach reconstructs fully volumetric, interconnected wall entities and room topology on multi-story buildings with weak assumptions on the input data. The resulting models are very close to the requirements needed for Building Information

Modeling tasks including volumetric representations of room spaces and wall entities, and their interrelations. Additional hard constraints such as forcing or avoiding certain entities at chosen locations may simply be added as constraints of the optimization problem. We demonstrated our approach on a variety of real-world datasets.

Future work for our proposed method includes the extension of the plane arrangement data structure to support slanted surfaces and possibly non-planar primitives. Strategies for reducing computational complexity by e.g. pruning invalid surface and wall candidates early in the process would improve applicability to larger-scale datasets. Also, connecting our reconstruction methodology with e.g. opening and object detection approaches would further enrich the resulting models.

Acknowledgments

We acknowledge the Visualization and MultiMedia Lab at University of Zurich (UZH) and Claudio Mura for the acquisition of the 3D point clouds, and UZH as well as ETH Zürich for their support to scan the rooms represented in these datasets. Their datasets were used in our evaluation (Figure 5). We also used datasets provided by The Royal Danish Academy of Fine Arts Schools of Architecture, Design and Conservation (CITA) (Table 2), from The ISPRS Benchmark on Indoor Modeling [47] (Figure 10), and from the Stanford 3D Large-Scale Indoor Spaces Dataset [48] (Figure 9). This work was supported by the DFG projects KL 1142/11-1 (DFG Research Unit FOR 2535 Anticipating Human Behavior) and KL 1142/9-2 (DFG Research Unit FOR 1505 Mapping on Demand).

References

- [1] Sanchez, V, Zakhor, A. Planar 3D modeling of building interiors from point cloud data. In: Image Processing (ICIP), 2012 19th IEEE International Conference on. IEEE; 2012, p. 1777–1780.
- [2] Oesau, S, Lafarge, F, Alliez, P. Indoor scene reconstruction using feature sensitive primitive extraction and graph-cut. ISPRS Journal of Photogrammetry and Remote Sensing 2014;90:68–82.
- [3] Turner, E, Cheng, P, Zakhor, A. Fast, automated, scalable generation of textured 3D models of indoor environments. IEEE Journal of Selected Topics in Signal Processing 2015;9(3):409–421.
- [4] Mura, C, Mattausch, O, Pajarola, R. Piecewise-planar reconstruction of multi-room interiors with arbitrary wall arrangements. In: Computer Graphics Forum; vol. 35. Wiley Online Library; 2016, p. 179–188.
- [5] Ochmann, S, Vock, R, Wessel, R, Klein, R. Automatic reconstruction of parametric building models from indoor point clouds. Computers & Graphics 2016;54:94–103. URL: <http://www.sciencedirect.com/science/article/pii/S0097849315001119>. doi:doi:10.1016/j.cag.2015.07.008; special Issue on CAD/Graphics 2015.

- [6] Murali, S, Speciale, P, Oswald, MR, Pollefeys, M. Indoor scan2bim: Building information models of house interiors. In: Intelligent Robots and Systems (IROS), 2017 IEEE/RSJ International Conference on. IEEE; 2017, p. 6126–6133.
- [7] Liu, C, Wu, J, Furukawa, Y. Floornet: A unified framework for floorplan reconstruction from 3D scans. arXiv preprint arXiv:180400090 2018;.
- [8] Budroni, A, Boehm, J. Automated 3D reconstruction of interiors from point clouds. *International Journal of Architectural Computing* 2010;8(1):55–73.
- [9] Adan, A, Huber, D. 3D reconstruction of interior wall surfaces under occlusion and clutter. In: 3D Imaging, Modeling, Processing, Visualization and Transmission (3DIMPVT), 2011 International Conference on. IEEE; 2011, p. 275–281.
- [10] Xiong, X, Adan, A, Akinci, B, Huber, D. Automatic creation of semantically rich 3D building models from laser scanner data. *Automation in Construction* 2013;31:325–337.
- [11] Mura, C, Mattausch, O, Villanueva, AJ, Gobbetti, E, Pajarola, R. Automatic room detection and reconstruction in cluttered indoor environments with complex room layouts. *Computers & Graphics* 2014;44:20–32.
- [12] Previtali, M, Barazzetti, L, Brumana, R, Scaioni, M. Towards automatic indoor reconstruction of cluttered building rooms from point clouds. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences* 2014;2(5):281.
- [13] Ambruş, R, Claiaci, S, Wendt, A. Automatic room segmentation from unstructured 3-d data of indoor environments. *IEEE Robotics and Automation Letters* 2017;2(2):749–756.
- [14] Macher, H, Landes, T, Grussenmeyer, P. From point clouds to building information models: 3D semi-automatic reconstruction of indoors of existing buildings. *Applied Sciences* 2017;7(10):1030.
- [15] Wang, R, Xie, L, Chen, D. Modeling indoor spaces using decomposition and reconstruction of structural elements. *Photogrammetric Engineering & Remote Sensing* 2017;83(12):827–841.
- [16] Okorn, B, Xiong, X, Akinci, B, Huber, D. Toward automated modeling of floor plans. In: *Proceedings of the symposium on 3D data processing, visualization and transmission*; vol. 2. 2010;.
- [17] Liu, C, Wu, J, Kohli, P, Furukawa, Y. Raster-to-vector: Revisiting floorplan transformation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2017, p. 2195–2203.
- [18] Oesau, S, Lafarge, F, Alliez, P. Indoor scene reconstruction using primitive-driven space partitioning and graph-cut. In: *Eurographics workshop on urban data modelling and visualisation*. 2013;.
- [19] Xiao, J, Furukawa, Y. Reconstructing the world's museums. *International journal of computer vision* 2014;110(3):243–258.
- [20] Mura, C, Mattausch, O, Villanueva, AJ, Gobbetti, E, Pajarola, R. Robust reconstruction of interior building structures with multiple rooms under clutter and occlusions. In: *Computer-Aided Design and Computer Graphics (CAD/Graphics)*, 2013 International Conference on. IEEE; 2013, p. 52–59.
- [21] Mura, C, Villanueva, AJ, Mattausch, O, Gobbetti, E, Pajarola, R. Reconstructing complex indoor environments with arbitrary wall orientations. In: *Eurographics (Posters)*. 2014, p. 19–20.
- [22] Turner, E, Zakhor, A. Floor plan generation and room labeling of indoor environments from laser range data. In: *Computer Graphics Theory and Applications (GRAPP)*, 2014 International Conference on. IEEE; 2014, p. 1–12.
- [23] Mura, C, Pajarola, R. Exploiting the room structure of buildings for scalable architectural modeling of interiors. In: *ACM SIGGRAPH 2017 Posters*. ACM; 2017, p. 4.
- [24] Stambler, A, Huber, D. Building modeling through enclosure reasoning. In: *3D Vision (3DV)*, 2014 2nd International Conference on; vol. 2. IEEE; 2014, p. 118–125.
- [25] Thomson, C, Boehm, J. Automatic geometry generation from point clouds for BIM. *Remote Sensing* 2015;7(9):11753–11775.
- [26] Monszpart, A, Mellado, N, Brostow, GJ, Mitra, NJ. Rapter: rebuilding man-made scenes with regular arrangements of planes. *ACM Trans Graph* 2015;34(4):103–1.
- [27] Nan, L, Wonka, P. Polyfit: Polygonal surface reconstruction from point clouds. In: *The IEEE International Conference on Computer Vision (ICCV)*. 2017;.
- [28] Jung, J, Stachniss, C, Kim, C. Automatic room segmentation of 3D laser data using morphological processing. *ISPRS International Journal of Geo-Information* 2017;6(7):206.
- [29] Bobkov, D, Kiechle, M, Hilsenbeck, S, Steinbach, E. Room segmentation in 3D point clouds using anisotropic potential fields. In: *Multimedia and Expo (ICME)*, 2017 IEEE International Conference on. IEEE; 2017, p. 727–732.
- [30] Ochmann, S, Vock, R, Wessel, R, Tamke, M, Klein, R. Automatic generation of structural building descriptions from 3D point cloud scans. In: *Computer Graphics Theory and Applications (GRAPP)*, 2014 International Conference on. IEEE; 2014, p. 1–8.
- [31] Bassier, M, Van Genechten, B, Vergauwen, M. Classification of sensor independent point cloud data of building objects using random forests. *Journal of Building Engineering* 2018;.
- [32] Adán, A, Quintana, B, Prieto, SA, Bosché, F. Scan-to-bim for 'secondary' building components. *Advanced Engineering Informatics* 2018;37:119–138.
- [33] Son, H, Kim, C. Semantic as-built 3D modeling of structural elements of buildings based on local concavity and convexity. *Advanced Engineering Informatics* 2017;34:114–124.
- [34] Garwood, TL, Hughes, BR, O'Connor, D, Calautit, JK, Oates, MR, Hodgson, T. A framework for producing gbxml building geometry from point clouds for accurate and efficient building energy modelling. *Applied Energy* 2018;224:527–537.
- [35] Hyland, N, O'Keeffe, S, Dore, C, Brodie, S. Automatic open standard reporting for dimensional control compliance 2017;.
- [36] O'Keeffe, S, Hyland, N, Dore, C, Brodie, S. Automatic validation of as-is and as-generated ifc bims for advanced scan-to-bim methods 2017;.
- [37] Brodie, S, Hyland, N, Dore, C, O'Keeffe, S. The bim & scan@ platform: A cloud-based cyber-physical system for automated solutions utilising real & virtual worlds. 2017;.
- [38] Krispel, U, Evers, HL, Tamke, M, Ullrich, T. Data completion in building information management: electrical lines from range scans and photographs. *Visualization in Engineering* 2017;5(1):4.
- [39] Krijnen, T, Beetz, J. An ifc schema extension and binary serialization format to efficiently integrate point cloud data into building models. *Advanced Engineering Informatics* 2017;33:473–490.
- [40] Schnabel, R, Wahl, R, Klein, R. Efficient ransac for point-cloud shape detection. *Computer Graphics Forum* 2007;26(2):214–226.
- [41] Dongen, S. A cluster algorithm for graphs. *Tech. Rep.*; Amsterdam, The Netherlands, The Netherlands; 2000.
- [42] Oesau, S, Verdier, Y, Jamin, C, Alliez, P, Lafarge, F, Giraudot, S. Point set shape detection. In: *CGAL User and Reference Manual*; 4.12 ed. CGAL Editorial Board; 2018, URL: <https://doc.cgal.org/4.12/Manual/packages.html#PkgPointSetShapeDetection3Summary>.
- [43] Boykov, Y, Veksler, O, Zabih, R. Fast approximate energy minimization via graph cuts. *IEEE Transactions on pattern analysis and machine intelligence* 2001;23(11):1222–1239.
- [44] Gurobi Optimization, I. Gurobi optimizer reference manual. 2016. URL: <http://www.gurobi.com>.
- [45] Rusu, RB, Cousins, S. 3D is here: Point cloud library (pcl). In: *Robotics and automation (ICRA)*, 2011 IEEE International Conference on. IEEE; 2011, p. 1–4.
- [46] Wein, R, Berberich, E, Fogel, E, Halperin, D, Hemmer, M, Salzmann, O, et al. 2D arrangements. In: *CGAL User and Reference Manual*; 4.12 ed. CGAL Editorial Board; 2018, URL: <https://doc.cgal.org/4.12/Manual/packages.html#PkgArrangement2Summary>.
- [47] Khoshelham, K, Vilarinho, LD, Peter, M, Kang, Z, Acharya, D. The isprs benchmark on indoor modelling. *International Archives of the Photogrammetry, Remote Sensing & Spatial Information Sciences* 2017;42.
- [48] Armeni, I, Sax, A, Zamir, AR, Savarese, S. Joint 2D-3D-Semantic Data for Indoor Scene Understanding. *ArXiv e-prints* 2017;arXiv:1702.01105.