applied sciences

MDPI

Article

Robust Automatic Orientation of SLAM Point Clouds via 2D Binary-Image Cross-Correlation

Jose Juan Arranz [1], Rocio Romero-Jaren [1,*], Irving Diaz [1] and Carmen Jaren [2]

[1] Departamento de Ingeniería Topográfica y Cartografía, Escuela Técnica Superior de Ingenieros en Topografía, Geodesia y Cartografía, Universidad Politécnica de Madrid, Campus Sur-UPM, Calle Mercator 2, 28031 Madrid, Spain; josejuan.arranz@upm.es (J.J.A.); irving.diaz@alumnos.upm.es (I.D.)

[2] Departamento de Ingeniería, Escuela Técnica Superior de Ingeniería Agronómica y Biociencias, Universidad Pública de Navarra, Campus de Arrosadia, Arrosadia, s/n, 31006 Pamplona, Spain; cjaren@unavarra.es

* Correspondence: rocio.romero.jaren@upm.es

Featured Application

CORA-2D automatically orients successive SLAM surveys of a stockpile warehouse onto a wall model, without user intervention, as a preliminary step to computing the change in stockpile volume.

Abstract

SLAM-based handheld laser scanners survey a building in minutes, but each survey lies in an arbitrary coordinate system, is affected by drift and has a density and coverage that change between epochs. In stockpile warehouses, where almost all the geometry changes between surveys, existing registration methods need an initial alignment, repeatable local descriptors or training data. This work presents CORA-2D, a deterministic method that registers unoriented SLAM clouds onto a wall model. CORA-2D extracts the walls in three horizontal slices, infers their orientation under the Manhattan-world assumption and rasterises them as binary images. An exhaustive branch-and-bound cross-correlation estimates the translation and resolves the 90° ambiguity. A finer local search then refines the planimetric solution, and the mode of the floor elevation differences fixes the vertical offset. On four surveys of an operational warehouse, with initial rotations of up to 91° and offsets of up to 475 m, the median point-to-plane distance on the walls ranged from 1.5 to 5.1 cm. The solution remained unchanged when only 5% of the points were used, and a cloud covering 30% of the warehouse was registered with the same accuracy. ICP, NDT and FPFH + RANSAC did not converge without initialisation.

Keywords: SLAM; automatic orientation; point cloud registration; industrial warehouse monitoring; binary-image cross-correlation; Manhattan world

Academic Editor: Mauro Lo Brutto

Received: 22 October 2026
Revised: 17 November 2026
Accepted: 24 November 2026
Published: 28 November 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

1.1. SLAM Clouds and Multitemporal Monitoring

SLAM-based (simultaneous localisation and mapping) handheld laser scanners allow spaces without GNSS coverage to be documented quickly. They are used to map indoor spaces and manage buildings [1–3], tunnels, mines and underground infrastructure [4–7], construction sites against BIM models [8,9] and buildings with long corridors

[10,11]. They are also used under forest canopy [12–19] and in areas struck by disasters or affected by natural hazards [20–26]. In all these cases SLAM clouds are noisier than those from static scanners, and their main source of error is the drift accumulated along the trajectory, which is mostly vertical. Registering them onto control points or onto a static reference reduces this error to the centimetre level [6,27].

A particularly demanding application is the repeated survey of bulk-material warehouses. Stockpile volume has been estimated with UAV photogrammetry [28–30], with terrestrial laser scanning [31,32] and with UAV-mounted SLAM LiDAR in GNSS-denied warehouses [33]. The latter levels each cloud separately but does not resolve the planimetric rotation or the translation between epochs. Comparing epochs requires that they all share a reference system and that real changes are not mistaken for alignment errors [34,35]. On construction sites, this registration has been solved with phase correlation between photogrammetric clouds [36] or against the BIM model [37], but current methods degrade when much of the geometry changes between epochs [38]. This is the situation in a stockpile warehouse: the stockpiles occupy most of the floor and change between surveys, and only the floor and the walls remain.

1.2. Registration Methods and Their Limitations

Local methods, such as ICP [39] and NDT [40,41], converge to the nearest minimum and need an initial alignment within their basin of convergence, so they are usually used as refiners in coarse-to-fine schemes [42–44]. Global methods remove this dependence at a higher computational cost: 4PCS and Super4PCS search for congruent four-point sets [45,46] and Go-ICP obtains the global optimum by branch and bound [47]. Feature-based methods match handcrafted or learned descriptors [48,49] and estimate the transformation robustly [50]. End-to-end learning methods [51,52] achieve high accuracy on benchmark datasets, but require training and GPUs and generalise poorly to large unseen scenes [53,54]. In low-density multitemporal clouds, keypoint repeatability is the limiting factor: authors in [55] observed that around 60% of ISS keypoints [56] were not repeatable even with a 50 cm tolerance, and sparse or unbalanced sampling continues to affect correspondence-based methods [57]. The flat, repetitive surfaces of an industrial warehouse represent the least favourable case for all of them.

1.3. Two-Dimensional Registration and Closest Related Work

Reducing registration to 2D is not new. In robotics, correlative scan matching exhaustively correlates a rasterised scan with a map [58], Cartographer accelerates this search with branch and bound over multiresolution grids [59], and LiDAR scans have been localised on architectural floor plans [60]. In photogrammetry, construction-site clouds projected to images have been registered by phase correlation [61], the orientation of indoor spaces has been normalised under the Manhattan-world assumption, which leaves a 90° ambiguity [62,63], and clouds have been aligned with building models by matching lines and planes [64]. Other closely related works solve parts of the problem under assumptions that do not hold here: a preliminary alignment provided by the scanner's visual–inertial system and the manufacturer's software [65], an initial orientation and texture for matching SIFT features [66], a prior external coarse registration before detecting stable areas [67], or local descriptors that become ambiguous when the geometry repeats [68]. Table S1 in the Supplementary Materials summarises the limitations of each work and how CORA-2D addresses them.

1.4. Objective and Contributions

This work presents CORA-2D (Cross-Orthogonal Orientation and Registration

Algorithm in 2D), a deterministic method that orients multitemporal indoor SLAM clouds with respect to a stable reference model, without initial alignment, ground control points or training data. Instead of matching points or descriptors, it matches the walls, which persist between surveys, after projecting them into 2D binary images. The following two objectives guide the evaluation: (1) to recover a common reference system with low sensitivity to point density and to the area covered, including areas with point spacing of up to 5 cm (about 400 points/m^2); and (2) to reduce the computational cost by operating mainly in 2D without losing centimetre-level planimetric resolution. The contributions are:

- a deterministic four-degree-of-freedom registration workflow with all its parameters fixed in advance;
- wall extraction by means of horizontal slices at three heights, which recovers the walls even when the stockpiles partly occlude them;
- an exhaustive binary-image cross-correlation with branch and bound, which resolves the 90° ambiguity and guarantees the global maximum within the search range; and
- an evaluation on four surveys, complete and partial, of an operational warehouse, with sensitivity analysis and comparison with ICP, NDT and FPFH + RANSAC.

Section 2 describes the application, the data and the method; Section 3 presents and discusses the results; and Section 4 draws the conclusions.

2. Materials and Methods

2.1. Application and Data

CORA-2D was developed for the Acopios module of the Lot Of Points application [69], which computes the change in volume of each storage area of a warehouse for users without expertise in point cloud processing. The registration works as a black box: it must be deterministic, require no parameter tuning and tolerate clouds that cover only part of the warehouse, because it is not always necessary to survey all the stockpiles.

The reference is a wall model of the warehouse. It was surveyed with a Leica RTC360 terrestrial laser scanner (up to 2 million points per second and 1.9 mm 3D accuracy at 10 m, according to the manufacturer [70]), registered in Cyclone and exported as a cloud of 100 million points. The wall planes were fitted to it with 3DReshaper, and their 16,945 faces were sampled in MDTopX 9.22.01.016 [71] in horizontal layers every 0.2 m, giving a cloud of 359,931 points that contains only the walls, with no floor or interior content. The SLAM clouds were acquired with a Leica BLK2GO (up to 420,000 points per second and ±10 mm system accuracy indoors [72]) on four dates, with the trajectory processed in Cyclone REGISTER 360 (Table 1). Each survey took a few minutes, at normal walking pace and without targets, and has an arbitrary origin and orientation; that is, the clouds enter the algorithm unoriented. With respect to the model, S1 is rotated by 19.1° and offset by about 20 m, S2 is rotated by 91.4° and offset by about 475 m, S3 is rotated by 19.2°, and S4 is rotated by −10.7° and offset by about 39 m. S4 is a partial survey covering 30% of the model. The warehouse is divided into several rooms of different sizes and ceiling heights, in which the material is stockpiled in different ways.

Table 1. Point cloud inventory. The points "after thinning" are those of the algorithm input (0.02 m 3D grid), and the "0.1 m cells" are those of the 2D grid of step 1. Coverage is the percentage of the model's 0.1 m cells that the cloud occupies after registration.

Cloud	Sensor	Raw Points	Points After Thinning	0.1 m Cells	Plan Area (m^2)	Mean Raw Spacing (cm)	Model Coverage (%)
T	RTC360 (wall model)	359,931	359,931 (not thinned)	13,535	133.3 (wall footprint)	8.54	100
S1	BLK2GO	204,044,595	48,007,202	492,430	4926.3	0.54	83.0
S2	BLK2GO	315,153,687	59,813,071	534,178	5342.2	0.50	92.7
S3	BLK2GO	222,245,313	47,317,720	464,855	4648.9	0.52	80.2
S4	BLK2GO	27,365,540	9,396,730	162,321	1623.2	0.64	30.2 (partial)

Both scanners deliver levelled clouds thanks to their inertial measurement unit: the normal of the plane fitted to the floor deviates from the vertical by 0.04° in the model and by 0.13° in S1. CORA-2D relies on this property and estimates four degrees of freedom as follows: one rotation about the vertical and three translations. The SLAM input is the cloud thinned with a 0.02 m 3D grid, keeping the first point in each cell, which is also the cloud that is finally transformed; the model is used in full.

2.2. Method

CORA-2D estimates the four-degree-of-freedom rigid transformation that brings the SLAM cloud P_S into the frame of the model P_T through nine steps grouped into five blocks (Figure 1). Steps 1 to 5 are applied to both inputs, with the following three differences: the floor elevation is computed differently, the model's slice points are perturbed before triangulation, and the model is also rotated to its own principal orientation, a rotation that is undone at the end. All parameters are fixed in the implementation except three application settings (minimum line length and maximum stitching distance, both 0.5 m, and cell size, 0.1 m), which the user does not modify. Table S2 and Algorithm S1 of the Supplementary Materials list all the parameters and the pseudocode, and Section S2 gives the details needed to reimplement the method. The algorithm runs on a single thread.

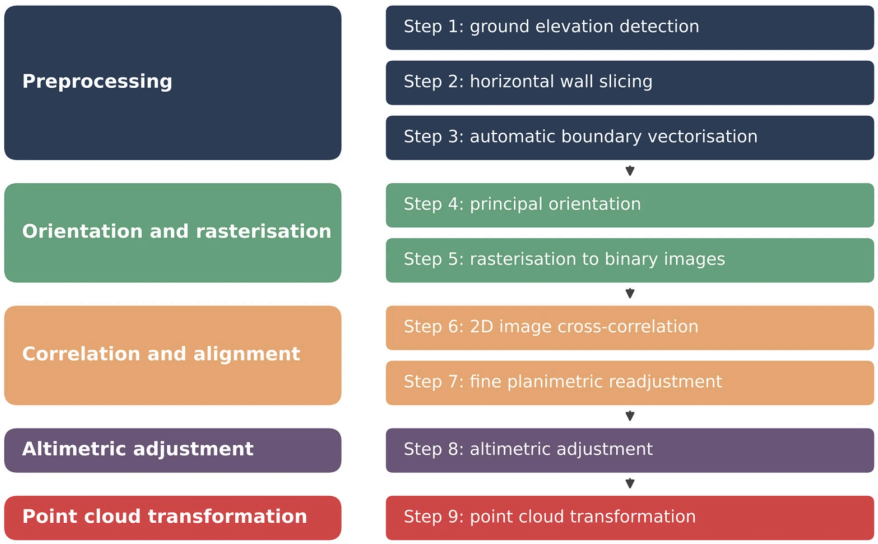

Figure 1. CORA-2D workflow in five functional blocks and nine steps.

2.2.1. Preprocessing: Floor, Wall Slices and Contours

Step 1: floor elevation. For each input, the lowest point in each cell of a 0.1 m 2D grid is kept, and a histogram of these elevations is built with 0.1 m bins. In the SLAM cloud, the floor elevation Z_ground is the mean of the points lying between 0.1 m below

and 0.2 m above the mode; in the model, which has no floor, the mode is the most populated layer of the wall base and Z_ground is that mode plus 0.05 m. The mode is used rather than the mean because it is insensitive to outliers [3]: the floor is the largest horizontal surface in the warehouse and the only one whose elevation repeats between surveys, whereas the stockpile surface changes and is spread over many bins. The points with $|z - Z_ground| \leq 0.2$ m form the floor used in step 8.

With m the mean elevation of the modal bin, the floor elevation of the SLAM cloud and the set F of floor points are

$$Z_ground = (1/|P_m|)\,\Sigma_(p \in P_m)\, z_p, \quad P_m = \{p : m - 0.1 \leq z_p \leq m + 0.2\}, \quad F = \{p : |z_p - Z_ground| \leq 0.2\} \tag{1}$$

with elevations in metres; in the model, Z_ground = m + 0.05 m.

Step 2: wall slices. Three horizontal slices are taken at 1, 2 and 3 m above Z_ground, with a half-thickness of 0.2 m, and processed separately, so that a wall visible at one height is retained even if stockpiles or machinery occlude it at another. The heights were chosen on a practical basis and were not optimised: at 1 m the slice lies above loose material on the floor, at 2 m above people walking through during the survey and at 3 m above most machinery. Higher slices were not used because they might not intersect the walls of rooms with lower ceilings. In the current implementation the height constants are given in Table S2, and Section 3.4 shows that the result depends little on them. The slice points of the model, which is synthetic and has collinear points, are shifted in X and Y by a random integer number of millimetres between −4 and 4, with a fixed seed, so that the triangulation does not degenerate.

Step 3: contour vectorisation. Each slice is triangulated in plan [73] and triangles with any edge longer than 1 m are removed. Points on the same wall are a few centimetres apart, whereas long edges span the gap among different walls, an opening, or the space between a wall and an object far from it. Removing these edges split the mesh into pieces that follow the walls, and objects resting against a wall remain in the same piece (Figure 2). The edges belonging to a single triangle are stitched into closed loops, outer and inner, which are the contours of the slice.

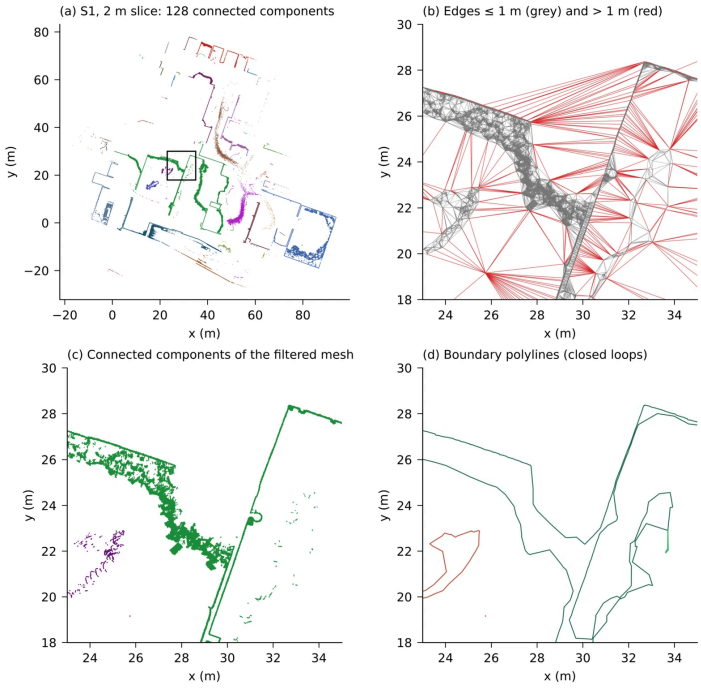

Figure 2. Contour vectorisation (step 3) on the 2 m slice of S1, in the coordinates of the unoriented

cloud. (**a**) The 128 connected components of the filtered mesh over the whole warehouse, with the detail inset. (**b**) Detail: Delaunay triangulation edges; those longer than 1 m (red) are discarded and the remaining ones (grey) form the filtered mesh. (**c**) Connected components of the filtered mesh in the detail (3). (**d**) Contour loops in the detail (4): edges belonging to a single triangle, stitched into closed outer and inner loops. The stockpile resting against the wall (centre of **b**–**d**) forms a single component with it, and its contour is merged into that of the wall.

2.2.2. Orientation and Rasterisation

Step 4: principal orientation. The orientation of each input is inferred from the azimuths of its contours under the Manhattan-world assumption [62], with no prior position or control points. The azimuth of each segment is reduced to the interval [−45°, 45°) and accumulated in a cyclic histogram of 90 bins of 1°, weighted by segment length. The mode is the weighted mean of the azimuths in the highest-weight bin and its two neighbours, which makes the estimate continuous when the wall direction falls near the boundary between two bins. The orientation is then refined iteratively, with up to 20 iterations. Each iteration (i) keeps the segments close to the mode (within 10 bins); (ii) stitches them into polylines per slice and direction, joining endpoints less than 0.5 m apart; (iii) reduces each polyline to a line by orthogonal least squares, removing residuals greater than 3.29σ; (iv) discards lines shorter than 0.5 m; and (v) rotates contours and lines according to the mode of the lines. The iterations stop when the rotation changes by less than 0.0001° between consecutive iterations. The accumulated rotation is the principal orientation (θ_S in the SLAM cloud and θ_T in the model), and the lines from the last iteration, rotated by that angle, are the geometry that is rasterised. As in the normalisation of Hübner et al. [63], the resulting orientation is ambiguous up to multiples of 90°; this ambiguity is resolved in step 6.

Formally, the azimuth α_i of each segment, measured clockwise from the +Y axis, is reduced to the interval [−45°, 45°) and accumulated in the histogram H with its length l_i as weight:

$$\tilde{\alpha}_i = \alpha_i - 90° \cdot \lfloor (\alpha_i + 45°)/90° \rfloor, \quad H(k) = \sum_{i : \lfloor \tilde{\alpha}_i + 45° \rfloor = k} l_i, \quad k = 0, \ldots, 89 \tag{2}$$

If k* is the bin with the highest weight, c* its centre and B the set of segments in k* and its two neighbours (modulo 90), the mode is

$$\mu = c^* + \sum_{i \in B} l_i \cdot \delta_i / \sum_{i \in B} l_i, \quad \delta_i = \tilde{\alpha}_i - c^* \text{ reduced to } [-45°, 45°) \tag{3}$$

Each polyline, with points p_j, is replaced by the line n·p = ρ, with n = (cos φ, sin φ), which minimises the sum of the squared orthogonal distances:

$$(\varphi^*, \rho^*) = \operatorname{argmin}_{(\varphi, \rho)} \sum_j (n(\varphi) \cdot p_j - \rho)^2 \tag{4}$$

The solution is the line through the centroid of the points with normal n that is equal to the eigenvector of the smallest eigenvalue of their covariance matrix; points with $|n \cdot p_j - \rho| > 3.29\sigma$ are removed and the fit is repeated. The principal orientation is the sum of the rotations g^(k) of each iteration, $\theta = \sum_k g^{(k)}$, and the fit stops when $|g^{(k)} - g^{(k-1)}| < 0.0001°$.

Step 5: rasterisation. The lines are rasterised into a binary image with cell size c = 0.1 m covering their bounding box, densified with one point every c; a pixel is set to 1 if a line passes through it, with no dilation or hole filling. The SLAM and model images share the cell size but have their own origin and dimensions. In both, the walls appear as one-pixel-wide lines, without the stockpiles or the rest of the warehouse contents.

2.2.3. Correlation and Alignment

Step 6: cross-correlation. Let A_S,q be the set of active pixels of the SLAM image rotated by q × 90°, with q ∈ {0, 1, 2, 3}. For each candidate translation d = (Δrow, Δcol), the score is the number of those pixels that fall on an active pixel of the model image I_T:

$$C_q(d) = \Sigma_(r,c)\in A_S,q \quad I_T(r + \Delta row, c + \Delta col) \tag{5}$$

The score is not normalised, so a cloud that covers only part of the building can achieve a high match with its own pixels. The search is exhaustive over ±½ of the rows and columns of I_T and is repeated for each of the four rotations, re-rasterising the rotated SLAM lines, and the rotation–translation pair with the highest score is retained. To speed it up, branch and bound [74] is applied, as in correlative scan matching [58,59]: the active pixels are traversed row by row and a candidate is abandoned as soon as its accumulated matches, plus the pixels still to be checked, cannot exceed the best score. This bound preserves the solution of the exhaustive search. The algorithm always returns a solution. The ratio between the scores of the best and second-best rotations is recorded as a quality indicator, since a ratio close to 1 signals symmetry or a lack of evidence. The result is the rotation giro1 = θ_S + 90q* and a coarse translation (difx, dify) in metres.

If N is the number of active pixels in A_S,q, c_j(d), the matches accumulated after checking j of them and C* the best score found so far, candidate d is abandoned as soon as

$$c_j(d) + (N - j) < C^* \tag{6}$$

The inequality is strict, so a candidate that can still tie is evaluated in full. In the worst case, the search compares 4·(2⌊h_T/2⌋ + 1)·(2⌊w_T/2⌋ + 1)·N pixels, where h_T and w_T are the rows and columns of I_T; pruning reduces this number without changing the maximum, and because only 1% to 2% of the pixels are active (Section 3.1), N is small.

Step 7: fine refinement. With the rotation fixed, the correlation is repeated with 5 cm cells in a window of ±40 pixels (±2 m) around the position from step 6. The SLAM image is now rasterised from the points of the three horizontal slices, already rotated and translated, and the model image from its contours densified every 2.5 cm, so that the refinement uses all the wall evidence, not only the evidence that survived the fit in step 4. The quantisation error drops from ±5 to ±2.5 cm per axis, and the difference between the translations of steps 6 and 7 (difx2, dify2) serves as a robustness check.

Formally,

$$d^* = \text{argmax}_(d\in W)\ C'(d), \quad W = \{d_0 + (i, j) : |i|, |j| \le 40\} \tag{7}$$

where C′ is the score of Equation (5) with 0.05 m cells, d_0 the position from step 6 and (difx2, dify2) the difference between d* and d_0 expressed in metres.

2.2.4. Vertical Adjustment and Transformation

Step 8: vertical adjustment. The floor points of the SLAM cloud, already in the model frame, are triangulated into a digital elevation model [73]. The floor points of the model are thinned with a 0.5 m 2D grid, keeping the one with the lowest elevation, and projected onto that DEM; for each one, δz = Z_SLAM − Z_model is computed. The vertical offset t_z is the negative of the mean, of the differences in the modal bin of their histogram (1 cm bins). As in step 1, the mode prevents stockpiles and objects present in only one survey from biassing the result. In the wall model used here, the floor points of the model are the wall base, so the resulting absolute elevation is referred to that base (Section 3.3).

For each floor point i of the model, with coordinates (x_i, y_i, z_i),

$$\delta z_i = Z_SLAM(x_i, y_i) - z_i, \quad t_z = -(1/|B^*|)\ \Sigma_{(i \in B^*)}\ \delta z_i \tag{8}$$

where Z_SLAM(x_i, y_i) is interpolated in the triangle of the digital elevation model that contains the point and B* is the set of differences in the 1 cm modal bin.

Step 9: transformation. The transformation composes the rotation of the SLAM cloud, the translation and the rotation that undoes the orientation of the model:

$$p' = R(giro2) \cdot (R(giro1) \cdot p + t) \tag{9}$$

with giro1 = θ_S + 90q*, giro2 = −θ_T and t = (−(difx + difx2), −(dify + dify2), −t_z), where R(α) is an anticlockwise rotation of α degrees about the coordinate origin (0, 0). The two rotations are equivalent to a single rotation about the vertical axis followed by a translation, so the transformation is rigid and has four parameters. The planimetric translations are truncated to 1 mm, and the vertical one is rounded to 1 mm. The application writes giro1, t and giro2 in the MDTopX 9.22.01.016 transformation format and transforms the thinned cloud.

2.3. Implementation and Execution Environment

The results were obtained by running the module code on the four clouds with an instrumented branch of MDTopX 9.22.01.016 [71] that outputs the intermediate values. The transformations are identical with and without instrumentation, and seven runs on each cloud gave the same transformation. All runs used an AMD Ryzen 9 3950X (3.5 GHz) with 128 GB of memory and Windows 11 Pro; the code was compiled with Microsoft Visual C++ 19.51 and run on a single thread, as in the module. The evaluation was carried out with the help of an AI-based programming assistant (Claude Code, Anthropic) under the supervision of the authors; the CORA-2D algorithm and its C++ implementation are the authors' own, and all numbers and figures come from running that code on the data. MDTopX 9.22.01.016, Lot Of Points 1.0.11.0 and Claude 2.42304.0 (e7f60e) were used.

3. Results and Discussion

Figure 3 shows the wall model with S1 registered onto it, and Figure 4 shows three of the clouds before and after registration.

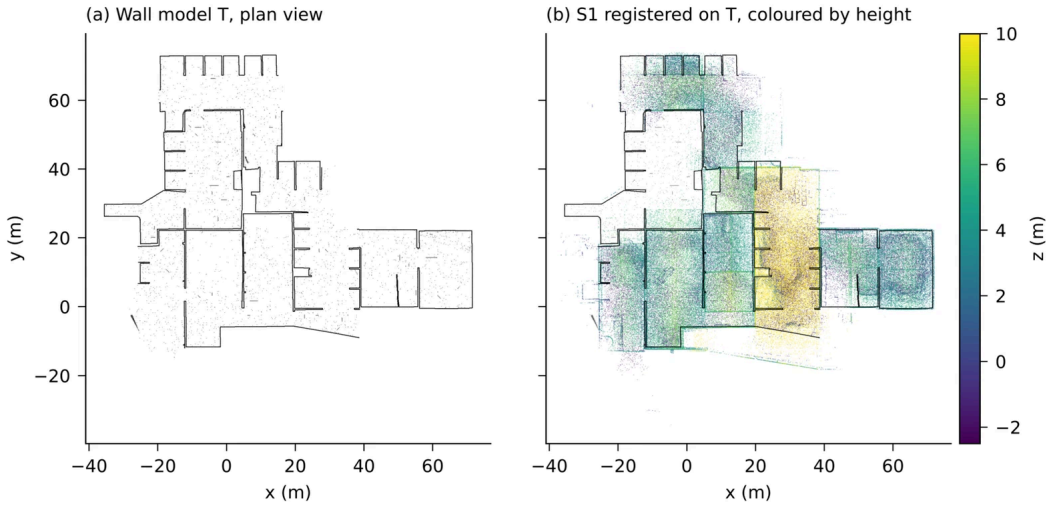

Figure 3. Test data. (**a**) Plan view of the wall model T (359,931 points), formed by layers of points every 0.2 m on the wall planes and with no floor surface. (**b**) Cloud S1 registered with CORA-2D

(sample of 250,000 points, coloured by elevation) onto T.

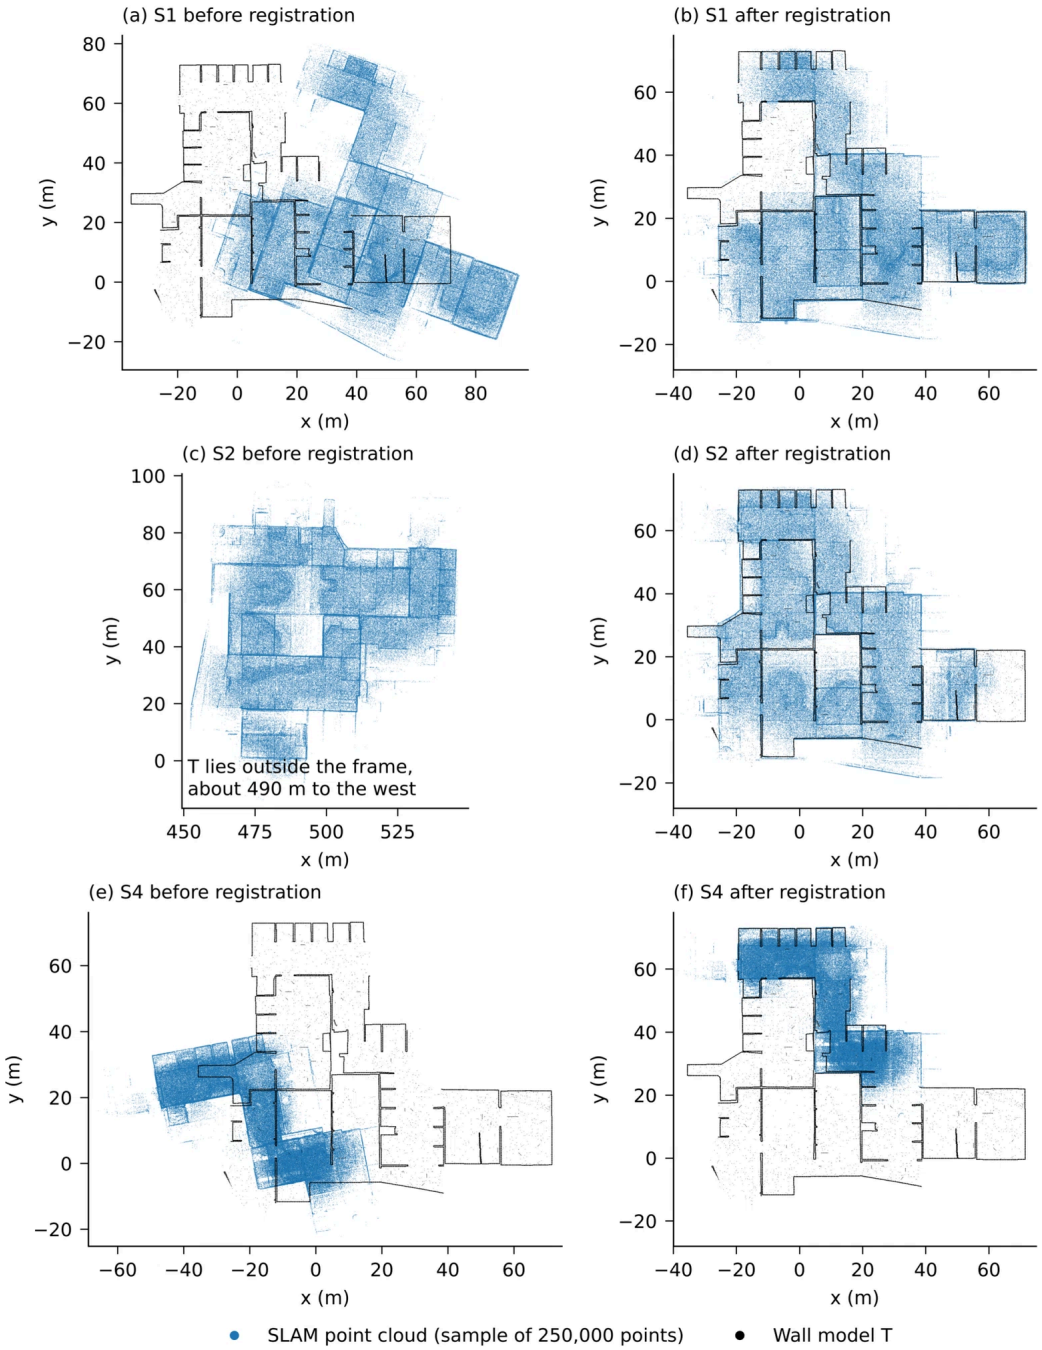

Figure 4. Clouds S1, S2 and S4 (samples of 250,000 points) before (**a**,**c**,**e**) and after (**b**,**d**,**f**) registration with CORA-2D, over the wall model T. S1 is initially rotated by 19.1° and offset by about 20 m; S2 is rotated by 91.4° and offset by about 475 m, so the model lies outside the frame in (**c**); and S4, the partial cloud, is rotated by −10.7° and offset by about 39 m.

3.1. Preprocessing, Orientation and Rasterisation

In all four clouds the modal bin of the elevation histogram corresponds to the floor, and the floor elevation differs from the mode by less than 3 cm (Table S3 in the Supplementary Materials). The histogram of S1 is not unimodal: it has two secondary peaks, with 78% and 52% of the count of the modal bin, which step 1 ignores (Figure S1). A histogram with several levels of similar support is a possible failure mode in other warehouses. The triangulation produces between 362 and 529 contour

loops in the complete clouds, 150 in the partial one and 136 in the model. The 1 m slice has a greater contour length than the 3 m slice (2120 versus 1915 m in S1) because it cuts through stockpiles, machinery and vehicles, whereas at 3 m mainly walls and pillars remain. The higher slices thus compensate for occlusion in the lower ones, and step 4 must discard the interior contours. For this reason, only between 11.9% and 16.0% of the contour length of the SLAM clouds lies within 1° of the mode of the first iteration, compared with 57.8% in the model (Figure S2). Even so, the iterative fit converges in five iterations in S1, S3 and S4 and in eight in S2 and corrects the initial mode by between 0.04° and 0.36° (Table S4). In the model it converges to −0.2220°, within 0.01° of the wall direction computed directly from the faces of the DXF. The binary images at 0.1 m range from 625 × 825 to 1047 × 1129 pixels, of which only between 1.1% and 2.0% are active (0.9% in the model), which explains the efficiency of the correlation.

3.2. Correlation and Alignment

Table 2 summarises the correlation. In all four clouds the correct rotation was identified by a clear margin: the best score is 4.37, 4.48, 4.32 and 2.91 times the second in S1 to S4, and in S2 the chosen rotation is k = 1, with a total rotation of 91.6°. In the complete clouds, 31–33% of the model pixels are matched in step 6, because the SLAM image also contains lines from the interior contents, and 42–54% in step 7, whose slice-point image is much denser. In S4, which covers a third of the warehouse, the values are 12.2% and 16.0%. The search evaluates 3,645,684 candidate translations (4 × 851 × 1071), and pruning abandons between 35% and 71% of them before traversing all the pixels. Step 6 took between 97 and 121 s in the complete clouds and 30 s for the partial one, whereas step 7, with 6561 candidates, took less than 5 s. The score map of the correct rotation has a single sharp peak, with two ridges along the axes corresponding to translations parallel to the long walls, and the other three rotations have no comparable peak (Figure S3 in the Supplementary Materials). Figure 5 shows the images being correlated and the matching pixels of the solution.

Table 2. Planimetric parameters of steps 6 and 7 per SLAM cloud. The step 7 translation is the accumulated one; the match refers to the active pixels of the model image; times are medians of five runs.

Cloud	Step	Cell (m)	Δx (m)	Δy (m)	k (×90°)	Matching Pixels	Match (%)	Best/2nd Rotation	Pruned Candidates (%)	Time (s)
S1	6	0.10	−16.608	−11.127	0	2598	31.6	4.37	35.3	96.9
S1	7	0.05	−16.620	−11.144	=	10,349	42.7	—	—	2.9
S2	6	0.10	68.537	−470.374	1	2578	31.4	4.48	37.2	120.8
S2	7	0.05	68.475	−470.422	=	12,976	53.5	—	—	4.1
S3	6	0.10	−13.233	4.867	0	2723	33.2	4.32	48.1	108.1
S3	7	0.05	−13.241	4.858	=	10,100	41.7	—	—	2.6
S4	6	0.10	23.534	31.675	0	1003	12.2	2.91	70.7	30.1
S4	7	0.05	23.536	31.689	=	3882	16.0	—	—	0.6

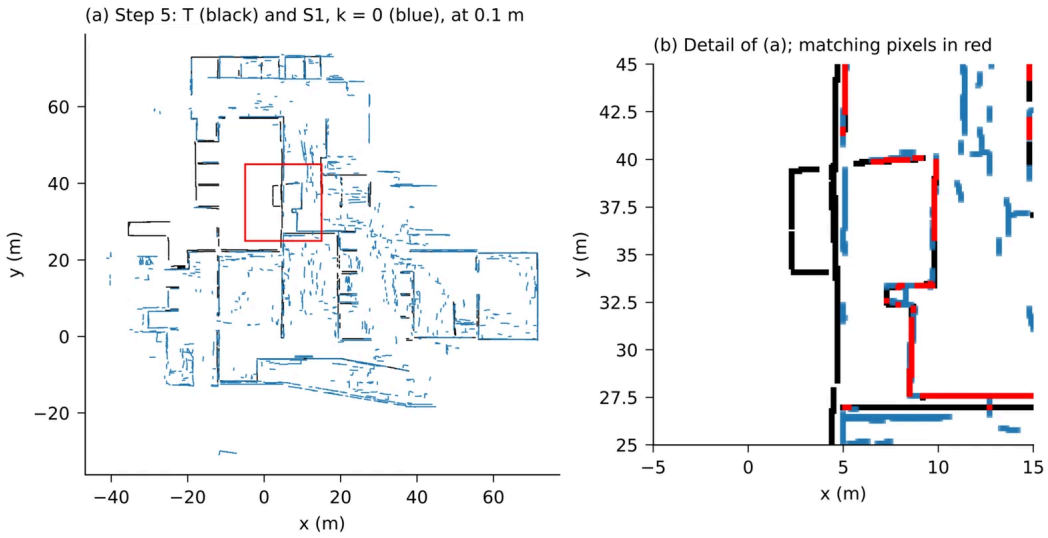

Figure 5. Binary images from step 5 and correlation result from step 6 in S1. (**a**) 0.1 m image of the wall model T (black, 8211 active pixels) and of S1 (blue, 17,356) in the canonical frame, with S1 rotated and translated according to the best solution of step 6 (k = 0); the red box marks the detail. (**b**) Detail: in red, the pixels of S1 that fall on a model pixel, whose number (2598 in the whole image) is the score of Equation (5). The score maps of the four rotations and the step 7 window are shown in Figure S3 in the Supplementary Materials.

The fine refinement changed the coarse translation by 2.1, 7.8, 1.2 and 1.4 cm in S1 to S4 (less than one coarse cell) without changing the rotation. The search at 0.1 m is therefore stable, and step 7 halves the quantisation error at a cost of less than 5 s. This error matters for the application, because on the flank of a stockpile with slope α a horizontal offset d produces an apparent elevation difference in $d \cdot \tan \alpha$.

3.3. Final Transformation and Accuracy

The total rotations were 19.30°, 91.64°, 19.41° and −10.46° in S1 to S4 (Table S6). The complete process, including reading the two clouds and writing the oriented cloud, took 221, 315 and 222 s for the complete clouds (48, 60 and 47 million points) and 70 s for the partial one (9.4 million), with a peak of 8.9 GB of memory. The coarse correlation, the vertical adjustment (50–99 s, dominated by the triangulation of 3.7 to 5.9 million floor points) and the triangulation of the horizontal slices (46–67 s) account for more than 90% of the time.

References for accuracy. There are no control points measured in both epochs, so accuracy was assessed in two ways. The first compared the registered clouds with the complete RTC360 survey (100 million points, with floor) brought into the model frame. This transformation is not independent: it was computed with the same algorithm (in MDTopX 9.22.01.016) and against the same model, although with a different configuration (maximum triangle edge length of 0.2 m and single-bin mode in step 4). The residuals therefore measure the difference between two registrations against the same model and cannot detect an error common to both. The following two checks bound this error: the walls of the transformed RTC360 lie 7.7 cm (median) from those of the model, and a point-to-plane ICP shifts them by 3.5 cm planimetrically with a rotation of 0.06°. The second used the wall model itself: the point-to-plane distance from the SLAM walls to the faces of the DXF from which the model was sampled, and the correction that a four-degree-of-freedom diagnostic ICP still applies to the registered cloud. Stable points were selected with the same rules in all clouds: floor with a near-vertical normal ($|n_z| \geq 0.95$) in cells that also have floor in the RTC360, which excludes the stockpiles of either epoch,

and walls with a near-horizontal normal ($|n_z| \leq 0.2$) within 0.3 m of an RTC360 wall, which excludes the stockpile slopes. Normals were estimated with Open3D (radius of 0.1 m and 30 neighbours).

Against the RTC360, the point-to-plane median on the walls is 1.5, 4.0, 5.1 and 4.1 cm in S1 to S4, with between 40% and 84% of the wall points within 5 cm (Table 3). Against the model, it is 3.6 to 4.2 cm, with a residual rotation of 0.006° to 0.067° and a residual offset of 1.4 to 3.8 cm (Table 4). The partial cloud S4 is registered with the same accuracy as the complete ones. The 95th percentile of the C2C distance (22–49 cm) is due to drift. The error does not grow continuously; instead, some sectors of the warehouse are displaced relative to others, and these sectors differ between epochs (Figure 6). This pattern is consistent with drift between segments of the trajectory, which a rigid transformation cannot absorb. In addition, a six-degree-of-freedom ICP shows that S1 is tilted by 0.27° relative to the RTC360 (from 0.09° to 0.12° in the others); CORA-2D does not estimate the tilt, and at 50 m from the origin this tilt amounts to an elevation difference of 23 cm.

Table 3. Accuracy of the CORA-2D registration against the RTC360 survey brought into the model frame (see the text regarding the independence of this reference). Walls: C2C and point-to-plane distances; floor: vertical difference per 0.1 m cell (negative = SLAM cloud below).

Cloud	Wall Points ($\times 10^6$)	Walls C2C: Median/p95 (cm)	Walls C2C < 5 cm (%)	Walls Point-to-Plane: Median (cm)	Floor Cells ($\times 10^3$)	Floor Δz: Median (cm)	Floor Δz: IQR (cm)	Planimetric Correction Left by a Diagnostic ICP: Median/Max. (cm)
S1	5.19	2.1/22.4	83.7	1.5	114.3	−20.6	7.6	2.4/6.3
S2	5.57	5.5/48.9	46.5	4.0	112.9	−29.4	12.1	3.5/6.5
S3	4.44	6.9/41.0	39.9	5.1	65.2	−34.7	7.6	5.1/8.3
S4 (partial)	1.18	5.0/29.6	50.4	4.1	44.4	−7.4	4.4	7.1/12.3

Table 4. Accuracy of the registration against the wall model T itself (planes of the DXF from which it was sampled), on 200,000 wall points of each transformed cloud: point-to-plane median, and residual planimetric rotation and offset left by a four-degree-of-freedom diagnostic ICP against T.

Cloud	Walls Point-to-Plane Against T: Median (cm)	Residual Rotation Against T (°)	Residual Planimetric Offset: Median/Max. (cm)
S1	3.7	0.006	1.4/2.1
S2	4.1	0.022	1.9/3.8
S3	4.2	−0.026	3.6/6.0
S4 (partial)	3.6	0.067	3.8/9.3

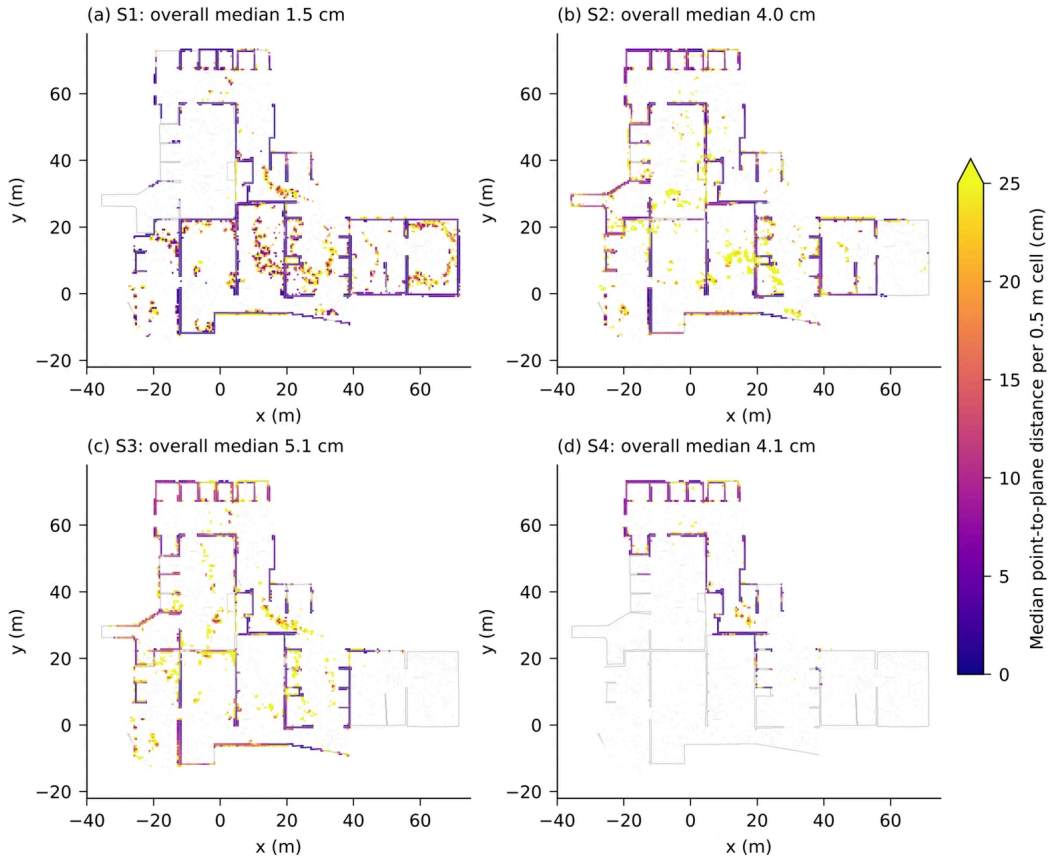

Figure 6. Map of planimetric residuals after registration: median, per 0.5 m cell, of the point-to-plane distance between the wall points of each transformed SLAM cloud and the planes of the RTC360 survey (sample of 200,000 points per cloud, of which 104,000–147,000 are wall points). **(a)**SLAM cloud S1 (overall median 1.5 cm);**(b)**SLAM cloud S2 (overall median 4.0 cm);**(c)**SLAM cloud S3 (overall median 5.1 cm);**(d)**SLAM cloud S4 (overall median 4.1 cm). The colour scale is saturated at 25 cm. Grey lines show the walls of the model. The isolated interior cells correspond to vertical surfaces in the RTC360 that are not walls of the model (machinery, stockpile sides) and that change between epochs.

Vertical. The transformed SLAM floor lies between 7.4 and 34.7 cm below the RTC360 floor. This is a consequence of applying step 8 to a model without a floor. The elevation is matched to the wall base, whose elevations in the model range from −1.49 to −1.14 m (5th and 95th percentiles), so the differences form a broad distribution, with standard deviations of 11.2 to 13.5 cm (Table S5 and Figure S4), in which the 1 cm single-bin mode is fragile: in S4 the chosen bin exceeds the next one by a single difference, which determines 7 cm of t_z. This bias does not affect volumes, because the Acopios module computes the volume of each stockpile relative to a plane specific to its area and its epoch, and it would disappear if the model included the floor, as step 8 assumes.

Comparison with other methods. The same pair of inputs was registered with point-to-plane ICP, NDT and FPFH + RANSAC from Open3D 0.20, with the clouds thinned to 0.1 m and the parameters given in Section S5 of the Supplementary Materials. A method is considered to have converged if the mean C2C distance on the walls is 20 cm or less and at least half of the points that CORA-2D retains within the wall mask are also retained. Without initialisation, none of them converges on any cloud (Table 5): in S1, S3 and S4, ICP and NDT become trapped in local minima with wall means of 34 to 58 cm, and FPFH + RANSAC finds no hypothesis that passes the geometric checks (fitness 0) and returns the identity on all four clouds. When initialised with the CORA-2D solution,

ICP and NDT converge in 1 to 5 s, but do not improve it: the wall mean decreases only in S3 (from 12.5 to 11.8 and 11.4 cm), and the floor mean worsens by between 6 and 31 cm on all clouds, because the model has no floor and the vertical component of these methods remains unconstrained. These results confirm what was expected: local methods require an initialisation within their basin of convergence, and FPFH descriptors do not yield reliable correspondences in a warehouse of flat, repetitive surfaces.

Table 5. Comparison with other methods on the same pair of inputs (unoriented SLAM cloud against the wall model). Mean C2C in cm on walls/floor against the RTC360; "no" = does not converge according to the criterion given in the text. FPFH + RANSAC returns the identity on all four clouds, so its means in S1, S3 and S4 are those of the unmoved cloud; in S2, which lies about 475 m from the model, uninitialised ICP, NDT and FPFH + RANSAC return the identity and no points remain within the masks.

Cloud	CORA-2D	Point-to-Plane ICP, No Init.	ICP, CORA-2D Init.	NDT, No Init.	NDT, CORA-2D Init.	FPFH + RANSAC, No Init.
S1	4.8/16.9	52.2/26.3 (no)	7.9/30.2	53.5/13.9 (no)	6.9/48.3	51.4/114.6 (no)
S2	13.1/28.9	identity (no)	13.6/35.0	identity (no)	14.0/46.3	identity (no)
S3	12.5/35.2	47.9/41.5 (no)	11.8/48.8	58.4/7.0 (no)	11.4/45.5	57.9/11.5 (no)
S4	8.5/7.7	33.8/41.4 (no)	9.2/30.2	50.0/45.9 (no)	8.7/31.0	48.2/77.6 (no)
Time on S1 (s)	208.2 (algorithm)	37.8	3.6	4.5	4.6	110.5

Formulation of the methods. Table 6 compares the formulation of the four methods and explains these results. Point-to-plane ICP minimises the sum of the squared distances from each point to the tangent plane of its nearest neighbour in the other cloud, and NDT maximises the likelihood of the transformed points with respect to one Gaussian per voxel of the reference cloud; both are local optimisations in six degrees of freedom that converge to the minimum closest to the initial position. FPFH + RANSAC matches 33-dimensional local descriptors and generates random hypotheses from three correspondences, so it is global, but it depends on the descriptors being repeatable and distinctive, and its result varies with the seed. CORA-2D reduces the problem to four degrees of freedom, because the clouds are levelled, and maximises the number of matching wall pixels (Equation (5)) over all the translations in the range and the four rotations left by the Manhattan assumption, without point-to-point correspondences or descriptors. It therefore needs no initial alignment and is deterministic, at the cost of assuming an approximately Manhattan floor plan and of not estimating tilt or scale.

Table 6. Formulation of the compared methods. p_i and q_i are corresponding points of the two clouds, n_i the normal at q_i, (R, t) the transformation, and μ_k and Σ_k the mean and covariance of the Gaussian of the voxel k containing the transformed point.

Aspect	Point-to-Plane ICP	NDT	FPFH + RANSAC	CORA-2D
Objective function	min Σ_i ((R p_i + t − q_i)·n_i)^2	max Σ_i exp(−½ (R p_i + t − μ_k)^T Σ_k^{-1} (R p_i + t − μ_k))	Max number of correspondences consistent with the hypothesis	max C_q(d), Equation (5)
Correspondences	Point-to-plane with the nearest	Point to the Gaussian of the	Between 33-dimensional	None: overlap of wall pixels

Aspect	Point-to-Plane ICP	NDT	FPFH + RANSAC	CORA-2D
	neighbour, recomputed at each iteration	voxel containing it	FPFH descriptors	
Search	Local, iterative	Local, iterative	Global, random (hypotheses from three correspondences)	Global, exhaustive with branch and bound (4 rotations × all translations in the range)
Degrees of freedom	6	6	6	4 (rotation about the vertical and three translations)
Initial alignment	Required	Required	Not required	Not required
Determinism	Yes, for a given initialisation	Yes, for a given initialisation	No: depends on the seed	Yes
Main assumption	Initial position within the basin of convergence	Initial position within the basin of convergence	Repeatable and distinctive descriptors	Levelled clouds, approximately Manhattan floor plan and visible walls

3.4. Sensitivity

The registration of the four clouds was repeated, varying one factor at a time and leaving the rest at their default values (151 runs; Table 7). Subsampling the cloud to 50, 25, 10 and 5% of its points, either randomly or by grid, changes the solution by less than 7 cm and 0.1°. The raster cell size (from 0.05 to 0.4 m) does not change it (≤ 0.1 cm), because it is fixed by step 7; a single horizontal slice or a half-thickness of ±0.1 or ±0.4 m changes it by less than 8 cm in the complete clouds and by less than 10 cm in the partial one; the maximum triangle edge length (from 0.2 to 2 m) changes it by less than 7 cm; and the initial position (rotations from 15° to 180° and a translation of 100 m) by less than 6 cm.

Table 7. Sensitivity: largest mean planimetric difference from the default solution (cm) by factor and cloud, excluding gross errors, and number of gross errors (rotation > 1° or offset > 1 m) out of the number of variants.

Factor (Variants)	S1	S2	S3	S4
Subsampling to 50/25/10/5%, random and by grid (8)	5.5 (0/8)	6.9 (0/8)	5.8 (0/8)	4.5 (0/8)
Band and corner cropping to 75/50/25% (6)	22.3 (0/6)	51.0 (0/6)	8.5 (0/6)	50.3 (1/6)
Raster cell size 0.05/0.2/0.4 m (3)	0.1 (0/3)	0.1 (0/3)	0.1 (0/3)	0.1 (0/3)
Single horizontal slice: 1, 2 or 3 m (3)	3.5 (0/3)	4.3 (0/3)	5.4 (0/3)	9.5 (0/3)
Slice half-thickness ±0.1/±0.4 m (2)	7.6 (0/2)	4.2 (0/2)	4.0 (0/2)	4.5 (0/2)

Factor (Variants)	S1	S2	S3	S4
Maximum triangle edge length 0.2/0.5/2 m (3)	6.2 (0/3)	4.8 (0/3)	6.2 (0/3)	4.7 (0/3)
Initial rotation 15/45/90/135/180° and translation of 100 m (6)	4.8 (0/6)	2.0 (0/6)	4.1 (0/6)	6.0 (0/6)
Model perturbation seed (5)	1.1 (0/5)	0.9 (0/5)	4.7 (0/5)	1.0 (0/5)

Coverage is the factor that determines where the method fails and the only one that produced a gross error (rotation greater than 1° or offset greater than 1 m) in the 151 runs: when S4, which is already partial, is band-cropped to 25%, it is registered with a 90° rotation error and a ratio between rotations of 1.00. The complete clouds, whether band-cropped or corner-cropped, degrade without failing down to 20–25% of their area, with offsets of up to 51 cm. In S1 cropped at one corner, the point-to-plane median rises from 1.5 cm to 4.7 cm at 75% and to 6.7–7.7 cm between 50% and 20%, and the ratio between rotations falls from 4.37 to 1.28 without the correlation selecting a wrong rotation (Table S7). This ratio is the indicator that best signals a weak registration: it is 1.00 for the only gross error, 1.11 and 1.99 for the two cropped clouds offset by about 50 cm, and from 2.9 to 4.5 for the complete clouds. The vertical adjustment is the most sensitive to coverage: in S1 cropped to 50% and 40%, t_z changes by 48 and 52 cm.

The mode estimator in step 4 is decisive (Table S8 of the Supplementary Materials). With the single-bin mode, the wall orientation of S2 falls next to the boundary between two bins and the adjustment alternates between two values 0.9° apart without converging. Parabolic interpolation over three bins exhausts the 20 iterations on S2 and on the model itself with four of the five seeds and leaves a residual rotation against the model of −0.07° to −0.12°. The weighted mean of the modal bin and its two neighbours, which is the estimator used by CORA-2D, converges on all clouds and seeds, leaves a residual rotation of 0.006° to 0.067° and estimates the orientation of the model to within 0.01° of that computed from the faces of the DXF. The choice matters because the rotations are applied about the origin: at 500 m from the origin, where S2 lies, an error of 0.1° amounts to almost 0.9 m. The single-bin modes in steps 1 and 8 have the same discontinuity, and it would be advisable to replace them with estimators that average several bins.

3.5. Advantages and Limitations

CORA-2D provides a global and deterministic planimetric search, independent of radiometry and of learned descriptors, which explicitly resolves the 90° ambiguity and uses stable linear structures instead of the changing surfaces of the stockpiles. However, it requires a stable reference, a sufficient amount of visible wall, an approximately Manhattan floor plan, levelled clouds and predominantly rigid differences between the epoch and the reference. It does not estimate scale, tilt or local deformations, and the translation range of step 6, limited to half of the model image, prevents the registration of a partial cloud located at the opposite end of the warehouse. The results add the following three limitations: the vertical reference is whatever the model treats as floor, the tilt of the SLAM cloud is not corrected and, because the rotations are applied about the origin, a small angular error becomes an offset proportional to the distance from the origin. Finally, the internal correlation is not sufficient to assess the fit, because a repetitive floor plan can yield a strong but incorrect maximum, and generalisation must be verified in other warehouses.

The slice heights are constants in the current implementation (1, 2 and 3 m above the detected floor). In taller warehouses, or with stockpiles higher than 3 m, they could

be occluded, although in the warehouse tested a single slice at any of the three heights changes the solution by less than 10 cm (Table 7). A slice close to the maximum height would contain less non-wall content, but in a warehouse with rooms of different heights, such as the one tested, it would lie above the ceiling of the lower rooms, and near the roof there are trusses, crane rails or light fittings that also generate contours. Choosing the heights from the elevation distribution of each cloud, or of each room, would solve both problems.

Applicability to other scenarios. CORA-2D only requires a stable reference with visible vertical walls, an approximately Manhattan floor plan and levelled clouds, conditions that are also met by other indoor spaces, such as logistics and industrial warehouses, car parks, office buildings or buildings under construction compared with their design model once the walls are built. The reference does not have to be a wall model either: it can be a SLAM survey that has already been oriented. The method is not directly applicable where there are no vertical walls or orthogonal dominant directions, such as tunnels, mines or outdoor spaces; for non-Manhattan floor plans, step 4 could be generalised to arbitrary dominant directions and step 6 to a finer sampling of the rotation, at a higher cost. The tests were carried out in a single warehouse, albeit with rooms of different sizes and heights and different stockpiling arrangements, so its applicability to other buildings must be verified.

4. Conclusions

This work has presented CORA-2D, a deterministic algorithm that orients multitemporal indoor SLAM clouds by reducing 3D registration to the cross-correlation of 2D binary images of the walls.

With regard to the first objective, low sensitivity to density and to the area covered, the four clouds, acquired in four epochs, with initial rotations of −11° to 91° and offsets of up to 475 m, were registered with a median point-to-plane distance on the walls of 1.5 to 5.1 cm against the RTC360 survey and of 3.6 to 4.2 cm against the model, with residual rotations of less than 0.07°. The partial cloud, which covers 30% of the model, was registered with the same accuracy. The solution remains unchanged with only 5% of the points, and the complete clouds degrade without failing down to 20–25% of their area. ICP, NDT and FPFH + RANSAC did not converge without initialisation on these clouds.

With regard to the second objective, computational cost, the reduction to binary images and the pruning make the exhaustive search affordable: the complete process took between 221 and 315 s on clouds of 47 to 60 million points and 70 s on the partial one, on a single thread and without the user adjusting any parameter.

The evaluation also showed that the vertical adjustment depends on the model containing the floor, that the tilt of the cloud is not corrected and that the ratio between the best and the second-best rotation signals weak registrations. These findings point to several lines of future work. To strengthen the validation, an independent reference should be obtained with targets measured in both epochs, and the floor should be included in the model. The algorithm could be extended to estimate the tilt and allow arbitrary dominant directions; to make the slice heights configurable and choose them from the elevation distribution of each warehouse or room; to correct drift with a piecewise adjustment guided by the wall polylines; to use an already oriented SLAM survey as the reference for subsequent ones; and to return the ratio between rotations to the user and replace the single-bin modes in steps 1 and 8. Finally, the evaluation should be extended to other warehouses, to other indoor spaces with vertical walls and to Super4PCS and TEASER++.

Supplementary Materials: The following supporting information can be downloaded at https://www.mdpi.com/article/10.3390/app162311846/s1, Table S1: Closest related works and how CORA-2D addresses them; Table S2: CORA-2D parameters; Algorithm S1: CORA-2D pseudocode; Table S3: Preprocessing per cloud; Figure S1: Steps 1 and 2 on real data; Table S4: Principal orientation (step 4); Figure S2: Azimuth histogram of the first iteration of step 4; Figure S3: Correlation score on S1 (steps 6 and 7); Table S5: Vertical adjustment per SLAM cloud; Figure S4: Elevation differences in step 8; Table S6: Final transformation, time and memory per cloud; Table S7: S1 cropped at one corner; Table S8: Mode estimators in step 4. S1: Parameters of the Compared Methods.

Author Contributions: Conceptualization, J.J.A.; methodology, J.J.A.; software, J.J.A.; validation, J.J.A. and R.R.-J.; formal analysis, R.R.-J.; investigation, R.R.-J. and I.D.; resources, R.R.-J. and I.D.; data curation, J.J.A., writing—original draft preparation, R.R.-J.; writing—review and editing, J.J.A. and C.J.; visualisation, J.J.A.; supervision, R.R.-J. and C.J.; project administration, R.R.-J. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The datasets presented in this article are not readily available because they are sensitive for the company that owns the facilities and cannot be distributed. The Supplementary Materials include the complete pseudocode of the method (Algorithm S1) and all its parameters (Table S2).

Acknowledgments: The authors thank Leica Geosystems España for access to the data and Digi21.net for implementing this library in its software. During the preparation of this manuscript, the authors used Claude (Anthropic, version 2.42304.0 (e7f60e)) for the purposes of unifying and restructuring drafts written by the authors, improving the wording and checking the bibliographic data. The authors have reviewed and edited the output and take full responsibility for the content of this publication.

Conflicts of Interest: CORA-2D was developed for the Acopios module of the Lot Of Points application, by Dreaming with Objects S.L. The author J.J.A. has a professional development relationship with this company and has no financial interest in it beyond that relationship. The other authors declare that they have no financial interests or personal relationships that could have influenced this work.

Abbreviations

The following abbreviations are used in this manuscript:

2D/3D	Two-dimensional/three-dimensional
4PCSs	4-points congruent sets
AI	Artificial intelligence
BIM	Building information modelling
C2C	Cloud-to-cloud (distance)
CORA-2D	Cross-Orthogonal Orientation and Registration Algorithm in 2D
DEM	Digital elevation model
DXF	Drawing exchange format
FPFH	Fast point feature histograms
GNSS	Global navigation satellite system

GPU	Graphics processing unit
ICP	Iterative closest point
IQR	Interquartile range
ISSs	Intrinsic shape signatures
LiDAR	Light detection and ranging
NDT	Normal distributions transform
p95	95th percentile
RANSAC	Random sample consensus
SIFT	Scale-invariant feature transform
SLAM	Simultaneous localisation and mapping
UAV	Unmanned aerial vehicle

References

1. Cantoni, S.; Vassena, G. Fast indoor mapping to feed an indoor DB for building and facility management. *Int. Arch. Photogramm. Remote Sens. Spat. Inf. Sci.* **2019**, *XLII-2/W9*, 213–217. [Google Scholar] [CrossRef]
2. Kaijaluoto, R.; Hyyppä, A. Precise indoor localization for mobile laser scanner. *Int. Arch. Photogramm. Remote Sens. Spat. Inf. Sci.* **2015**, *XL-4/W5*, 1–6. [Google Scholar] [CrossRef]
3. Romero-Jarén, R.; Arranz, J.J. Automatic segmentation and classification of BIM elements from point clouds. *Autom. Constr.* **2021**, *124*, 103576. [Google Scholar] [CrossRef]
4. Jones, E.; Sofonia, J.; Canales, C.; Hrabar, S.; Kendoul, F. Applications for the Hovermap autonomous drone system in underground mining operations. *J. South. Afr. Inst. Min. Metall.* **2020**, *120*, 49–56. [Google Scholar] [CrossRef]
5. Fahle, L.; Holley, E.A.; Walton, G.; Petruska, A.J.; Brune, J.F. Analysis of SLAM-based lidar data quality metrics for geotechnical underground monitoring. *Min. Metall. Explor.* **2022**, *39*, 1939–1960. [Google Scholar] [CrossRef]
6. Štroner, M.; Urban, R.; Křemen, T.; Braun, J.; Michal, O.; Jiřikovský, T. Scanning the underground: Comparison of the accuracies of SLAM and static laser scanners in a mine tunnel. *Measurement* **2025**, *242*, 115875. [Google Scholar] [CrossRef]
7. Ma, A.; Lin, X.; Yao, W.; Ma, B.; Gao, Y.; Yang, X.; Yan, T.; Wang, X. Terrain-adaptive LiDAR-IMU SLAM with multi-feature fusion for complex underground environments. *ISPRS J. Photogramm. Remote Sens.* **2026**, *238*, 816–833. [Google Scholar] [CrossRef]
8. Ibrahimkhil, M.H.; Shen, X.; Barati, K.; Wang, C.C. Dynamic progress monitoring of masonry construction through mobile SLAM mapping and as-built modeling. *Buildings* **2023**, *13*, 930. [Google Scholar] [CrossRef]
9. Vassena, G.P.M.; Perfetti, L.; Comai, S.; Mastrolembo Ventura, S.; Ciribini, A.L.C. Construction progress monitoring through the integration of 4D BIM and SLAM-based mapping devices. *Buildings* **2023**, *13*, 2488. [Google Scholar] [CrossRef]
10. Gharineiat, Z.; Tarsha Kurdi, F.; Henny, K.; Gray, H.; Jamieson, A.; Reeves, N. Assessment of NavVis VLX and BLK2GO SLAM scanner accuracy for outdoor and indoor surveying tasks. *Remote Sens.* **2024**, *16*, 3256. [Google Scholar] [CrossRef]
11. Chrbolková, A.; Štroner, M.; Urban, R.; Michal, O.; Křemen, T.; Braun, J. A comparative study of indoor accuracies between SLAM and static scanners. *Appl. Sci.* **2025**, *15*, 8053. [Google Scholar] [CrossRef]
12. Ryding, J.; Williams, E.; Smith, M.J.; Eichhorn, M.P. Assessing handheld mobile laser scanners for forest surveys. *Remote Sens.* **2015**, *7*, 1095–1111. [Google Scholar] [CrossRef]
13. Oveland, I.; Hauglin, M.; Giannetti, F.; Kjørsvik, N.S.; Gobakken, T. Comparing three different ground based laser scanning methods for tree stem detection. *Remote Sens.* **2018**, *10*, 538. [Google Scholar] [CrossRef]
14. Chen, S.; Liu, H.; Feng, Z.; Shen, C.; Chen, P. Applicability of personal laser scanning in forestry inventory. *PLoS ONE* **2019**, *14*, e0211392. [Google Scholar] [CrossRef]
15. Gollob, C.; Ritter, T.; Nothdurft, A. Forest inventory with long range and high-speed personal laser scanning (PLS) and simultaneous localization and mapping (SLAM) technology. *Remote Sens.* **2020**, *12*, 1509. [Google Scholar] [CrossRef]
16. Hyyppä, E.; Yu, X.; Kaartinen, H.; Hakala, T.; Kukko, A.; Vastaranta, M.; Hyyppä, J. Comparison of backpack, handheld, under-canopy UAV, and above-canopy UAV laser scanning for field reference data collection in boreal forests. *Remote Sens.* **2020**, *12*, 3327. [Google Scholar] [CrossRef]
17. Wu, Y.; Zhong, S.; Ma, Y.; Zhang, Y.; Liu, M. Application of SLAM-based mobile laser scanning in forest inventory: Methods, progress, challenges, and perspectives. *Forests* **2025**, *16*, 920. [Google Scholar] [CrossRef]

18. Kükenbrink, D.; Marty, M.; Rehush, N.; Abegg, M.; Ginzler, C. Evaluating the potential of handheld mobile laser scanning for an operational inclusion in a national forest inventory—A Swiss case study. *Remote Sens. Environ.* **2025**, *321*, 114685. [Google Scholar] [CrossRef]
19. Zhao, Y.; Zhao, S.; Zhou, Y.; Shen, B.; Wang, J. A lightweight LiDAR-SLAM system integrated with semantic segmentation for the acquisition of forest biomass parameters. *Ann. For. Sci.* **2026**, *83*, 21. [Google Scholar] [CrossRef]
20. Chiabrando, F.; Sammartano, G.; Spanò, A. A comparison among different optimization levels in 3D multi-sensor models. A test case in emergency context: 2016 Italian earthquake. *Int. Arch. Photogramm. Remote Sens. Spat. Inf. Sci.* **2017**, *XLII-2/W3*, 155–162. [Google Scholar] [CrossRef]
21. Dewez, T.J.B.; Yart, S.; Thuon, Y.; Pannet, P.; Plat, E. Towards cavity-collapse hazard maps with Zeb-Revo handheld laser scanner point clouds. *Photogramm. Rec.* **2017**, *32*, 354–376. [Google Scholar] [CrossRef]
22. Sammartano, G.; Spanò, A. Point clouds by SLAM-based mobile mapping systems: Accuracy and geometric content validation in multisensor survey and stand-alone acquisition. *Appl. Geomat.* **2018**, *10*, 317–339. [Google Scholar] [CrossRef]
23. Shin, D.Y.; Cho, D.H.; Hong, J.W.; Han, S.H.; Cho, S.B.; Lim, E.T.; Jung, Y.H.; Kim, S.S. LiDAR-based multi-sensor module development and assessment of applicability for damage investigation of indoor disaster. *Int. Arch. Photogramm. Remote Sens. Spat. Inf. Sci.* **2022**, *XLIII-B3-2022*, 1167–1174. [Google Scholar] [CrossRef]
24. Wang, T.; Lu, F.; Qin, J.; Huang, T.; Kong, H.; Shen, P. AscDAMs: Advanced SLAM-based channel detection and mapping system. *Nat. Hazards Earth Syst. Sci.* **2024**, *24*, 3075–3094. [Google Scholar] [CrossRef]
25. LaSala, B.; Sankey, T.T.; Nebel, M.; Springer, A.E.; Mildice, A. Three-dimensional characterization of caves within the Grand Canyon's deep karst aquifer. *Sci. Rep.* **2025**, *15*, 31962. [Google Scholar] [CrossRef]
26. Li, H.; Zhang, J.; Li, Y.; Zhang, L.; Ming, Z.; Chen, R. Debris flow vulnerability assessment in the Eastern Tibetan Plateau using integrated UAV-SLAM 3D reconstruction: Shuihaizigou Gully case study. *Landslides* **2026**, *23*, 2193–2221. [Google Scholar] [CrossRef]
27. Keitaanniemi, A.; Rönnholm, P.; Kukko, A.; Vaaja, M.T. Drift analysis and sectional post-processing of indoor simultaneous localization and mapping (SLAM)-based laser scanning data. *Autom. Constr.* **2023**, *147*, 104700. [Google Scholar] [CrossRef]
28. Tucci, G.; Gebbia, A.; Conti, A.; Fiorini, L.; Lubello, C. Monitoring and computation of the volumes of stockpiles of bulk material by means of UAV photogrammetric surveying. *Remote Sens.* **2019**, *11*, 1471. [Google Scholar] [CrossRef]
29. Alsayed, A.; Yunusa-Kaltungo, A.; Quinn, M.K.; Arvin, F.; Nabawy, M.R.A. Drone-assisted confined space inspection and stockpile volume estimation. *Remote Sens.* **2021**, *13*, 3356. [Google Scholar] [CrossRef]
30. Alsayed, A.; Nabawy, M.R.A. Stockpile volume estimation in open and confined environments: A review. *Drones* **2023**, *7*, 537. [Google Scholar] [CrossRef]
31. Hu, X.; Xia, T.; Yang, L.; Wu, F.; Tian, Y. 3D modeling and volume measurement of bulk grains stored in large warehouses using bi-temporal multi-site terrestrial laser scanning data. *J. Agric. Eng.* **2024**, *55*, 1555. [Google Scholar] [CrossRef]
32. Gejdoš, M.; Lieskovský, M.; Ferenčík, M. Analysis and comparison of methods for determining small piles of wood chips using laser scanning technology. *BioResources* **2025**, *20*, 1807–1819. [Google Scholar] [CrossRef]
33. Wang, Y.; Xie, D.; Ren, S.; Sun, Z.; Yi, C.; Xiao, R.; Zhong, R. Indoor coal inventory algorithm using UAV-based LiDAR SLAM. *IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens.* **2025**, *18*, 26219–26234. [Google Scholar] [CrossRef]
34. Kharroubi, A.; Poux, F.; Ballouch, Z.; Hajji, R.; Billen, R. Three dimensional change detection using point clouds: A review. *Geomatics* **2022**, *2*, 457–485. [Google Scholar] [CrossRef]
35. Stilla, U.; Xu, Y. Change detection of urban objects using 3D point clouds: A review. *ISPRS J. Photogramm. Remote Sens.* **2023**, *197*, 228–255. [Google Scholar] [CrossRef]
36. Huang, R.; Xu, Y.; Hoegner, L.; Stilla, U. Temporal comparison of construction sites using photogrammetric point cloud sequences and robust phase correlation. *Autom. Constr.* **2020**, *117*, 103247. [Google Scholar] [CrossRef]
37. Meyer, T.; Brunn, A.; Stilla, U. Change detection for indoor construction progress monitoring based on BIM, point clouds and uncertainties. *Autom. Constr.* **2022**, *141*, 104442. [Google Scholar] [CrossRef]
38. Sun, T.; Hao, Y.; Huang, S.; Savarese, S.; Schindler, K.; Pollefeys, M.; Armeni, I. Nothing stands still: A spatiotemporal benchmark on 3D point cloud registration under large geometric and temporal change. *ISPRS J. Photogramm. Remote Sens.* **2025**, *220*, 799–823. [Google Scholar] [CrossRef]
39. Besl, P.J.; McKay, N.D. A method for registration of 3-D shapes. *IEEE Trans. Pattern Anal. Mach. Intell.* **1992**, *14*, 239–256. [Google Scholar] [CrossRef]

40. Biber, P.; Straßer, W. The normal distributions transform: A new approach to laser scan matching. In Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Las Vegas, NV, USA, 27–31 October 2003; Volume 3, pp. 2743–2748. [Google Scholar] [CrossRef]
41. Magnusson, M.; Lilienthal, A.; Duckett, T. Scan registration for autonomous mining vehicles using 3D-NDT. *J. Field Robot.* **2007**, *24*, 803–827. [Google Scholar] [CrossRef]
42. Tian, Y.; Yue, X.; Zhu, J. Coarse–fine registration of point cloud based on new improved whale optimization algorithm and iterative closest point algorithm. *Symmetry* **2023**, *15*, 2128. [Google Scholar] [CrossRef]
43. Luo, L.; Li, N.; Huang, K. DenseReg: Point cloud registration based on dense feature refinement. In Proceedings of the International Conference on Advanced Robotics, Automation Engineering and Machine Learning (ARAEML), Hangzhou, China, 28–30 June 2024; pp. 49–55. [Google Scholar] [CrossRef]
44. Mu, X.; Chen, H.; Wang, J.; Qin, H.; Zhu, Z. AUV SLAM method based on SO-CFAR and ADT feature extraction. *Sci. Prog.* **2024**, *107*, 00368504241286969. [Google Scholar] [CrossRef]
45. Aiger, D.; Mitra, N.J.; Cohen-Or, D. 4-points congruent sets for robust pairwise surface registration. *ACM Trans. Graph.* **2008**, *27*, 85. [Google Scholar] [CrossRef]
46. Mellado, N.; Aiger, D.; Mitra, N.J. Super 4PCS: Fast global pointcloud registration via smart indexing. *Comput. Graph. Forum* **2014**, *33*, 205–215. [Google Scholar] [CrossRef]
47. Yang, J.; Li, H.; Campbell, D.; Jia, Y. Go-ICP: A globally optimal solution to 3D ICP point-set registration. *IEEE Trans. Pattern Anal. Mach. Intell.* **2016**, *38*, 2241–2254. [Google Scholar] [CrossRef]
48. Rusu, R.B.; Blodow, N.; Beetz, M. Fast Point Feature Histograms (FPFH) for 3D registration. In Proceedings of the IEEE International Conference on Robotics and Automation (ICRA), Kobe, Japan, 12–17 May 2009; pp. 3212–3217. [Google Scholar] [CrossRef]
49. Gojcic, Z.; Zhou, C.; Wegner, J.D.; Wieser, A. The perfect match: 3D point cloud matching with smoothed densities. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Long Beach, CA, USA, 15–20 June 2019; pp. 5540–5549. [Google Scholar] [CrossRef]
50. Yang, H.; Shi, J.; Carlone, L. TEASER: Fast and certifiable point cloud registration. *IEEE Trans. Robot.* **2021**, *37*, 314–333. [Google Scholar] [CrossRef]
51. Huang, S.; Gojcic, Z.; Usvyatsov, M.; Wieser, A.; Schindler, K. PREDATOR: Registration of 3D point clouds with low overlap. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Nashville, TN, USA, 20–25 June 2021; pp. 4265–4274. [Google Scholar] [CrossRef]
52. Qin, Z.; Yu, H.; Wang, C.; Guo, Y.; Peng, Y.; Ilic, S.; Hu, D.; Xu, K. GeoTransformer: Fast and robust point cloud registration with geometric transformer. *IEEE Trans. Pattern Anal. Mach. Intell.* **2023**, *45*, 9806–9821. [Google Scholar] [CrossRef]
53. Chen, L.; Feng, C.; Ma, Y.; Zhao, Y.; Wang, C. A review of rigid point cloud registration based on deep learning. *Front. Neurorobotics* **2024**, *17*, 1281332. [Google Scholar] [CrossRef]
54. Zhang, Y.-X.; Gui, J.; Yu, B.; Cong, X.; Gong, X.; Tao, W.; Tao, D. Deep learning-based point cloud registration: A comprehensive survey and taxonomy. *Int. J. Comput. Vis.* **2026**, *134*, 168. [Google Scholar] [CrossRef]
55. Isfort, S.; Elias, M.; Maas, H.-G. Development and evaluation of a two-staged 3D keypoint based workflow for the co-registration of unstructured multi-temporal and multi-modal 3D point clouds. *ISPRS Ann. Photogramm. Remote Sens. Spat. Inf. Sci.* **2024**, *X-2-2024*, 113–120. [Google Scholar] [CrossRef]
56. Zhong, Y. Intrinsic shape signatures: A shape descriptor for 3D object recognition. In Proceedings of the IEEE 12th International Conference on Computer Vision Workshops (ICCVW), Kyoto, Japan, 27 September–4 October 2009; pp. 689–696. [Google Scholar] [CrossRef]
57. Tian, Y.; Li, X.; Yin, J. From pseudo- to non-correspondences: Robust point cloud registration via thickness-guided self-correction. *Comput. Graph.* **2026**, *134*, 104528. [Google Scholar] [CrossRef]
58. Olson, E.B. Real-time correlative scan matching. In Proceedings of the IEEE International Conference on Robotics and Automation (ICRA), Kobe, Japan, 12–17 May 2009; pp. 4387–4393. [Google Scholar] [CrossRef]
59. Hess, W.; Kohler, D.; Rapp, H.; Andor, D. Real-time loop closure in 2D LIDAR SLAM. In Proceedings of the IEEE International Conference on Robotics and Automation (ICRA), Stockholm, Sweden, 16–21 May 2016; pp. 1271–1278. [Google Scholar] [CrossRef]
60. Boniardi, F.; Caselitz, T.; Kümmerle, R.; Burgard, W. Robust LiDAR-based localization in architectural floor plans. In Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Vancouver, BC, Canada, 24–

28 September 2017; pp. 3318–3324. [Google Scholar] [CrossRef]
61. Huang, R.; Ye, Z.; Boerner, R.; Yao, W.; Xu, Y.; Stilla, U. Fast pairwise coarse registration between point clouds of construction sites using 2D projection based phase correlation. *Int. Arch. Photogramm. Remote Sens. Spat. Inf. Sci.* **2019**, *XLII-2/W13*, 1015–1020. [Google Scholar] [CrossRef]
62. Coughlan, J.M.; Yuille, A.L. Manhattan world: Compass direction from a single image by Bayesian inference. In Proceedings of the 7th IEEE International Conference on Computer Vision (ICCV), Kerkyra (Corfu), Greece, 20–27 September 1999; Volume 2, pp. 941–947. [Google Scholar] [CrossRef]
63. Hübner, P.; Weinmann, M.; Wursthorn, S.; Hinz, S. Pose normalization of indoor mapping datasets partially compliant with the Manhattan world assumption. *Remote Sens.* **2021**, *13*, 4765. [Google Scholar] [CrossRef]
64. Kaiser, T.; Clemen, C.; Maas, H.-G. Automated alignment of local point clouds in digital building models. *Int. Arch. Photogramm. Remote Sens. Spat. Inf. Sci.* **2019**, *XLII-5/W2*, 35–39. [Google Scholar] [CrossRef]
65. Akshulakov, K.; Mustabekov, A.; Sidorov, A.; Dosbossyn, A. Terrestrial laser scanning–based 3D modeling of a corridor-type building interior. *Int. J. Comput. Methods Exp. Meas.* **2026**, *14*, 229–237. [Google Scholar] [CrossRef]
66. Elias, M.; Kehl, C.; Schneider, D. Photogrammetric water level determination using smartphone technology. *Photogramm. Rec.* **2019**, *34*, 198–223. [Google Scholar] [CrossRef]
67. Isfort, S.; Elias, M.; Maas, H.-G. A workflow for the automatic extraction of glacier contours from 4D point clouds. *Int. Arch. Photogramm. Remote Sens. Spat. Inf. Sci.* **2026**, *XLIX-B2-2026*, 1035–1041. [Google Scholar] [CrossRef]
68. Liebold, F.; Wagner, F.; Sardemann, H.; Maas, H.-G. Automatic registration of voxel datasets using feature descriptor matching. *e-J. Nondestruct. Test.* **2026**, *31*, 1–7. [Google Scholar] [CrossRef]
69. Lot of Points; Software. Microsoft Store. 2026. Available online: https://apps.microsoft.com/detail/9p2bgtrjf7s7 (accessed on 2 September 2026).
70. Leica Geosystems AG. Leica RTC360 Product Specifications. Available online: https://shop.leica-geosystems.com/reality-capture/rtc/leica-rtc360-laser-scanner/buy (accessed on 2 September 2026).
71. Digi21. MDTopX; Software. 2026. Available online: https://www.digi21.net/MDTopX (accessed on 18 October 2026). [Google Scholar]
72. Leica Geosystems AG. Leica BLK2GO Technology and Specifications. Available online: https://shop.leica-geosystems.com/reality-capture/blk2go/technology (accessed on 2 September 2026).
73. Delaunay, B. Sur la sphère vide. A la mémoire de Georges Voronoï. *Bull. Acad. Sci. URSS Cl. Sci. Math. Nat.* **1934**, *6*, 793–800. [Google Scholar]
74. Land, A.H.; Doig, A.G. An automatic method of solving discrete programming problems. *Econometrica* **1960**, *28*, 497–520. [Google Scholar] [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.