Evaluation of the Robustness of Planar-Patches based 3D-Registration using Marker-based Ground-Truth in an Outdoor Urban Scenario

Kaustubh Pathak, Dorit Borrmann, Jan Elseberg, Narunas Vaskevicius, Andreas Birk, and Andreas Nüchter

Abstract— The recently introduced Minimum Uncertainty Maximum Consensus (MUMC) algorithm for 3D scene registration using planar-patches is tested in a large outdoor urban setting without any prior motion estimate whatsoever. With the aid of a new overlap metric based on unmatched patches, the algorithm is shown to work successfully in most cases. The absolute accuracy of its computed result is corroborated for the first time by ground-truth obtained using reflective markers. There were a couple of unsuccessful scan-pairs. These are analyzed for the reason of failure by formulating two kinds of overlap metrics: one based on the actual overlapping surface-area and another based on the extent of agreement of range-image pixels. We conclude that neither metric in isolation is able to predict all failures, but that both taken together are able to predict the difficulty level of a scan-pair vis-à-vis registration by MUMC.

I. INTRODUCTION

A 3D registration algorithm [1] based on matching large planar-patches extracted from "point-clouds" sampled from 3D sensors was recently introduced by some of the authors. The algorithm is termed *Minimum Uncertainty Maximum Consensus* (MUMC) and as the name suggests, it tries to find a set of correspondences between planar-patches from the two scans being matched which minimizes the uncertainty-volume of the registration result as measured by the determinant of the 6×6 covariance matrix of the computed pose. A closed-form least-squares solution of the registration was also presented along with explicit expressions for its uncertainty.

The MUMC algorithm was tested on a variety of sensors and compared to point based methods like Iterative Closest Point (ICP), in both its point-to-point [2] and point-to-plane [3] incarnations, and 3D Normal Distribution Transform (3D NDT). It was found that the algorithm had a bigger convergence radius than its competitors because it does a global search and does not depend on a local attraction to the nearest locally optimum solution. The goodness-of-fit was measured by the quality of alignment of scans, which is a relative measure. The algorithm was also used— embedded in a pose-graph for Simultaneous Localization and Mapping (SLAM)— to generate 3D maps of disaster scenarios [4]. In that work, clearly visible ground-truth structures were used for a qualitative evaluation. The main reason for the use of these somewhat subjective criteria for evaluation was that ground-truth in 3D is hard to come by. In the field of mobile robot navigation, ground-truth comparison either requires a simulation-study or precise GPS [5]. In [6] surveying data of buildings, stored in vector format, and available from government land registration offices was used as the source of ground-truth. Satellite images can also be used to obtain a rough ground-truth.

In this paper, we present a comparison of the MUMC algorithm results to the ground-truth for the first time. The ground-truth

The authors are with the Dept. of EECS, Jacobs University Bremen, 28751 Bremen, Germany. kaust@ieee.org, a.birk@jacobs-university.de
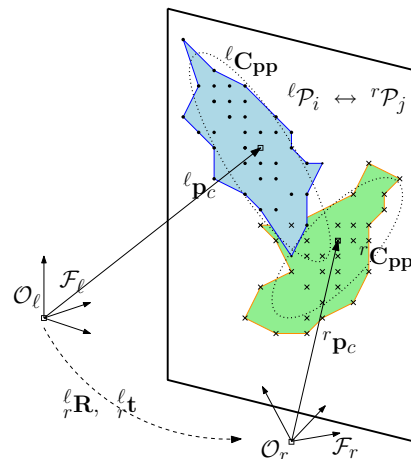
Fig. 1. The sensor in two coordinate-frames observes the same physical plane. Only a patch of the plane is visible from any pose; the polygonized patches are shown in color, with the sampled points shown as dots and crosses from the left and the right frame respectively.

was obtained based on the commercial reflective marker-based solution available on the high-end RIEGL VZ-400 3D laser scanner. The dataset consists of a fairly big outdoor urban scenario, viz. the old city-center of Bremen, Germany. The data was originally obtained as point-clouds of about 22.5 million points per scan, which were subsequently sub-sampled to range-images of half a million points per scan, from which planar-patches were extracted using a region-growing method described in [7]. This method also involves computation of the uncertainties of the plane-parameters using a sensor range error-model [8].

A. Notation Overview

The nomenclature used is briefly reviewed in this section. An infinite plane $\mathcal{P}(\hat{\mathbf{m}}, \rho)$ is given by the equation $\hat{\mathbf{m}} \cdot \mathbf{p} = \rho$, where ρ is the signed distance from the origin in the direction of the unit plane normal $\hat{\mathbf{m}}$. We see that $\mathcal{P}(\hat{\mathbf{m}}, \rho) \equiv \mathcal{P}(-\hat{\mathbf{m}}, -\rho)$. To achieve a consistent sign convention, we define planes as $\mathcal{P}(\hat{\mathbf{n}}, d)$, where, $d \triangleq |\rho| \geq 0$, and $\hat{\mathbf{n}} \triangleq \sigma(\rho) \hat{\mathbf{m}}$, where, $\sigma(\rho) = -1$ if $\rho < 0$ and $+1$ otherwise. If $\rho = 0$, then we choose the maximum component of $\hat{\mathbf{n}}$ to be positive.

For registration, we consider two robot-frames as shown in Fig. 1: a left one denoted as $\mathcal{F}_\ell$ with origin $\mathcal{O}_\ell$ from which the indexed plane-set $^\ell\mathcal{P}$ is observed, and a right one $\mathcal{F}_r$ with origin $\mathcal{O}_r$ from which the indexed plane-set $^r\mathcal{P}$ is observed. The equations of the planes are

$$^\ell\hat{\mathbf{n}}_i \cdot {}^\ell\mathbf{p} = {}^\ell d_i, \qquad {}^r\hat{\mathbf{n}}_i \cdot {}^r\mathbf{p} = {}^r d_i. \qquad (1)$$

An indexed set $^k\mathcal{P}$ of planar-patches is extracted [7] from a point-cloud associated with the k-th robot-frame $\mathcal{F}_k$ by segmentation of the range-image using region-growing followed by polygonization. As shown in Fig. 1, a patch has a set of points

$\mathbf{p}_j$ associated with it: the weighted scatter matrix associated with these points is denoted as $\mathbf{C_{PP}}$ and is depicted as the dotted ellipsoid in the figure. The weighted centroid is denoted as $\mathbf{p}_c$; the weights being taken to be inversely proportional to the trace of the covariance of the individual points $\mathbf{p}_j$. The normal of the plane turns out to be the eigenvector of $\mathbf{C_{PP}}$ corresponding to its least eigenvalue. Apart from the planar patch's $\hat{\mathbf{n}}$ and d parameters, the extraction procedure also gives [8] their 4×4 covariance matrix $\mathbf{C}$. Thus, the plane $^k\mathcal{P}$ is an ordered set of triplets given by

$$^k\mathcal{P} \triangleq \left\{ {}^k\mathcal{P}_i \langle {}^k\hat{\mathbf{n}}_i, \, {}^k d_i, \, {}^k\mathbf{C}_i \rangle, \, i = 1 \dots \mathsf{N}_k \right\}. \quad (2)$$

If the robot moves from $\mathcal{F}_\ell$ to $\mathcal{F}_r$, and observes the coordinates of the same physical point as $^\ell\mathbf{p}$ and $^r\mathbf{p}$ respectively, these coordinates are related by [9]

$$^\ell\mathbf{p} = {}^\ell_r\mathbf{R} \, {}^r\mathbf{p} + {}^\ell_r\mathbf{t}, \quad (3)$$

where, the translation $^\ell_r\mathbf{t} \triangleq \overrightarrow{\mathcal{O}_\ell\mathcal{O}_r}$, resolved in $\mathcal{F}_\ell$.

$^\ell\mathcal{P}_i \leftrightarrow {}^r\mathcal{P}_j$ means that the ith patch in the scan taken at $\mathcal{F}_\ell$ corresponds to the jth patch in $\mathcal{F}_r$. If we assume that the planes in the two frames have been renumbered so that planes with the same index physically correspond, then substituting (3) in (1) and comparing coefficients gives

$$^\ell\hat{\mathbf{n}}_i = {}^\ell_r\mathbf{R} \, {}^r\hat{\mathbf{n}}_i, \qquad {}^\ell\hat{\mathbf{n}}_i \cdot {}^\ell_r\mathbf{t} = {}^\ell d_i - {}^r d_i. \quad (4a)$$

The registration problem now consists of estimating $^\ell_r\mathbf{R}$ and $^\ell_r\mathbf{t}$ by solving the above in a least-squares sense. This procedure also yields the covariance of the registration solution. To find the actual correspondences between the planar-patches is the task of the MUMC algorithm [1].

II. ROBUSTNESS IN THE CASE OF A COMPLETE ABSENCE OF AN INITIAL MOTION ESTIMATE

Algorithms like ICP which match scans by iterative attraction to a local minimum of the registration-cost-function rely on a good initial guess of inter-scan pose change— e.g. that provided by vehicle odometry– for finding the right registration. Matching two scans in the absence of any initial guess is a formidable task, especially if the inter-scan pose change is considerable compared to the field of view (FOV) and the maximum range of the 3D range-sensor involved. In case the sensor has a large FOV, the main culprit for matching failures is occlusion.

MUMC searches the global correspondence-space of *large* planar-patches for the consensus which maximizes geometrical consistency and hence minimizes pose-uncertainty. The globality of the search implies that MUMC does not necessarily need to have an initial guess for the pose difference. However, if such a guess is available along with its uncertainty, it can still be utilized by MUMC to do χ^2-tests and prune the global search space— hence speeding-up its execution. This paper focuses on making the performance of MUMC more robust in the case of large, totally unknown inter-scan movements, assuming a relatively big FOV of the sensor. In other words, we focus on improving and evaluating the robustness of MUMC primarily w.r.t. occlusion.

A. Improving Robustness: Unmatched Planes Overlap Metric μ_u

The central dilemma facing a scan-matching algorithm is to weigh the size of scan-overlap against the quality of overlap. If only the size is given priority, poorly overlapping scans cannot be matched or are wrongly matched. In case only the quality of match is considered, a very small area may be accurately matched without there being a global agreement. MUMC proposes to solve

this issue by employing the uncertainty-volume (determinant of the covariance matrix) of the computed pose-registration as a metric to minimize: if *consistent* patch-correspondences are added to the matched set, this metric reduces; this avoids the temptation to greedily collect all patch-correspondences which merely satisfy some threshold.

In the original formulation of MUMC [1], the agreement of *unmatched* planes was not explicitly considered. In case of complete absence of initial guesses and the presence of occlusion, robustness can be improved by also evaluating planes in the two scans for which no correspondences were found. We introduce here a metric μ_u to compute the extent of overlap of unmatched planes. It is computed after [1, step 11 of Algorithm 2], when a set of potential plane correspondences Γ has been found, and the least-squares registration $^\ell_r\mathbf{R}$, $^\ell_r\mathbf{t}$ it implies, is computed. Now the set Γ needs to be evaluated w.r.t. the unmatched planes.

The basic idea is that the overlap can be measured as a χ^2 distance in terms of the weighted scatter matrix and the weighted centroid of the patches as depicted in Fig. 1. Assume that we want to evaluate whether the previously unmatched planes $^\ell\mathcal{P}_i$ and $^r\mathcal{P}_j$ overlap. The matched set Γ and its associated registration $^\ell_r\mathbf{R}$, $^\ell_r\mathbf{t}$ are considered fixed and certain.

$$^\ell\mathbf{q}_j \triangleq {}^\ell_r\mathbf{R} \, {}^r\mathbf{p}_{c,j} + {}^\ell_r\mathbf{t}, \quad (5)$$

$$\boldsymbol{\Sigma} \triangleq {}^\ell\mathbf{C}_{\mathbf{PP},i} + {}^\ell_r\mathbf{R} \, {}^r\mathbf{C}_{\mathbf{PP},j} \, {}^\ell_r\mathbf{R}^\mathsf{T}, \quad (6)$$

$$\chi_v^2 = ({}^\ell\mathbf{q}_j - {}^\ell\mathbf{p}_{c,i})^\mathsf{T} \boldsymbol{\Sigma}^{-1} ({}^\ell\mathbf{q}_j - {}^\ell\mathbf{p}_{c,i}). \quad (7)$$

For each unmatched plane $^r\mathcal{P}_j$, we find the $^\ell\mathcal{P}_i$ having the minimum χ_v^2 to it, such that, additionally, $^r\mathcal{P}_j$ and $^\ell\mathcal{P}_i$ are also translationally consistent [1, Sec. III-A2] with $^\ell_r\mathbf{t}$ and rotationally consistent, i.e. $^\ell\hat{\mathbf{n}}_i \cdot ({}^\ell_r\mathbf{R} \, {}^r\hat{\mathbf{n}}_j) \approx 1$. To be considered feasible, this minimum χ_v^2 should also be less than $\chi^2_{3,t\%}$, which is the χ^2 value for 3 d.o.f. at the significance level of $t\%$. We selected $t = 1\%$. If a specific index i in the ℓ-set is found to pair with more than one index in the r-set, the pairing with the lesser value of the minimum χ_v^2 is taken and the other is rejected. We now have an additional set of surmised correspondences denoted Γ_s from among the unmatched set of planes in addition to the previously fixed set Γ of matched planes. Hence,

$$\mu_u = \frac{\#\Gamma_s + \#\Gamma}{\mathsf{N}_r}, \qquad \Gamma \leftarrow \Gamma \cup \Gamma_s. \quad (8)$$

where, $\#$ denotes the set's size and N_r is the number of r-patches being matched. In [1, step 12 of Algorithm 2], we now consider the appended set Γ any further only if $\mu_u \geq \bar{\mu}_u$, i.e. the overlap is at least as big as a given threshold.

III. EVALUATION OF SCANS FOR MATCHABILITY GIVEN THE GROUND-TRUTH REGISTRATION

A. The Spatial Surface-Area Overlap Metric μ_s

The spatial overlapping surface area mutually visible from the two scans to be matched is one of the factors affecting matching success of MUMC. We propose to evaluate it using an Octree to discretize the space and represent the point-clouds of the scans. It uses the original point-cloud directly and not the extracted planes. First, the two scans are aligned using the given ground-truth. The volume covered by the points is recursively split into cubic octets until the size of the voxels falls below a threshold Δ_t. Each voxel that does not contain any points is deleted. From each remaining voxel, the center point is taken for the evaluation. A voxel in one scan is considered to correspond to the closest voxel in the other scan if their mutual distance is within a threshold $\bar{d}_s$. The number

Fig. 2. Scan 0 as range-image of size $1441 \times 401 = 577841$ with gray-scale based on distance. The horizontal field of view (FOV) covers full $360°$. The vertical field of view (FOV) of the sensor is $-40°$ to $+60°$. The façade of the St. Peter's Cathedral of Bremen is visible in the image's center.

of found corresponding voxels is the spatial surface-area overlap metric μ_s.

B. The Range-Image Overlap Metric μ_r

The spatial surface-area overlap metric μ_s is an absolute area measure, but it does not necessarily correspond to the number of pixels on the range-images mutually visible from the two scans. This is especially true for open spaces where the surfaces are far away from the sensor. The overlapping 2D area measured by the number of common points in the range-images is another important metric which determines the matching success of MUMC. It can conveniently normalized to the unit-interval by scaling with the total number of points in the range-image N_R.

This normalized metric is called μ_r and is computed as follows: transform each point ${}^r\mathbf{p}_j$ on the r-range-image using the known ground-truth transform to get the vector ${}^\ell\mathbf{q}_j$. From its beam direction, one can deduce the pixel (m, n) in the ℓ-range-image which corresponds to ${}^\ell\mathbf{q}_j$. Consider all pixels in the ℓ-range-image within a δ_n neighborhood of (m, n). Each such pixel corresponds to a 3D point ${}^\ell\mathbf{p}_i$. We consider the correspondence ${}^\ell\mathbf{p}_i$ to ${}^r\mathbf{p}_j$ resolved if ${}^\ell\mathbf{p}_i$ has the minimum-distance d_{ij} to ${}^\ell\mathbf{q}_j$ within the said neighborhood, and additionally, $d_{ij} \leq \Delta_d$. We take $\Delta_d = \delta_n \tan(\Delta_\phi)\|{}^\ell\mathbf{q}_j\|$, where Δ_ϕ is the angular resolution (maximum among horizontal and vertical) of the range-image. From the number of such resolved pixel correspondences N_c, we can readily compute $\mu_r = N_c/N_R$.

IV. EXPERIMENT

We now describe an experimental dataset using which we test MUMC for robustness and evaluate the predictive power of the overlap-metrics derived in Sec. III. A RIEGL VZ-400 (http://www.riegl.com) 3D laser range finder was used to collect 13 samples shown in Figs. 3 and 4. It is an out-door urban scenario viz. the city-center of Bremen at late night. The number of markers put in the scene were about 40. The dataset can be downloaded from http://kos.informatik.uni-osnabrueck.de/3Dscans.

The RIEGL VZ-400 has a maximum range of about half a kilometer. Each complete sample of 22.5 million points, with a vertical and horizontal angular sampling resolution of $0.04°$, took about 3 minutes to acquire. For planar-patches extraction, this data was sub-sampled for the same overall FOV but at an angular resolution of $0.25°$ to a range-image as shown in Fig. 2. The RIEGL sensor also returns intensity of reflection for each beam and thus certain custom-made reflector markers can be distinguished, if put in the scene as shown in Figs. 4(b) and 4(c). A RIEGL proprietary algorithm is used to detect these markers across scenes and to compute the ground-truth based on them. This algorithm sometimes requires human-intervention to distinguish the markers in the image.
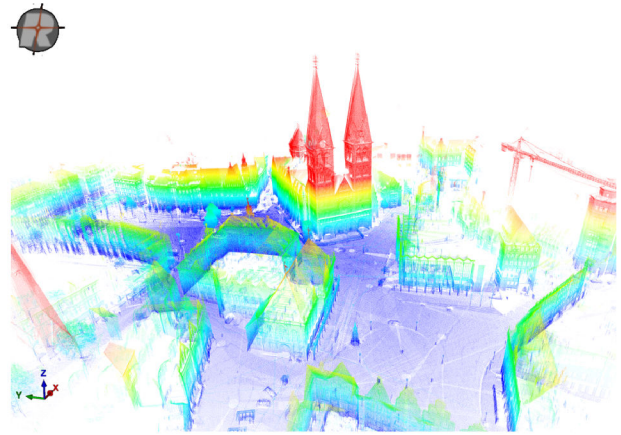
(a) The ground-truth points-based map computed using reflective markers.
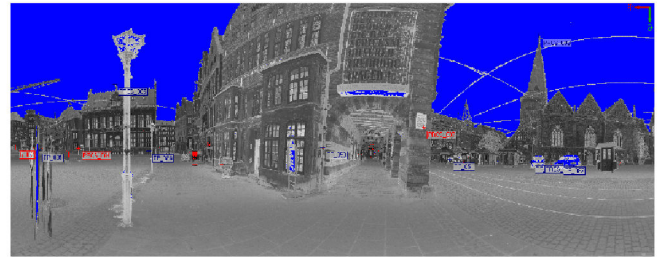
(b) Marker locations shown in one of the scenes.

(c) Close-up of a marker.

Fig. 4. The overall ground-truth computed using reflective markers.

A. Registration Comparison Results

The results are summarized in Table I. For each scan-pair, the first row shows the registration result of the MUMC algorithm, the second row gives the 1σ uncertainty values of this result computed by MUMC; finally, the third row gives the ground-truth, as determined using reflective markers. We emphasize that no initial guess (odometry or otherwise) was provided to MUMC. The successfully matched scan sequences from scan-1 to scan-5 and from scan-6 to scan-12 are shown in Fig. 5. In these cases, the registration result of MUMC agrees quite well with the ground-truth— indeed the yaw is always correct to within half a degree. Only the z displacement shows minor discrepancy: this is attributed to the ground not being matched because of its unevenness. This resulted in the z displacement being computed in many cases using overlap. In Table I n_ℓ and n_r denote the number of planar-patches considered for matching– they are the top 12% of all patches sorted according to their evidence. The number of correspondences found is denoted by n_c.

For the pair $0 - 1$ a false registration was obtained which is shown in Fig. 6. For the pair $5 - 6$, MUMC correctly identified

(a) Scan 1   (b) Scan 2   (c) Scan 3

(d) Scan 4   (e) Scan 5   (f) Scan 6

(g) Scan 7   (h) Scan 8   (i) Scan 9
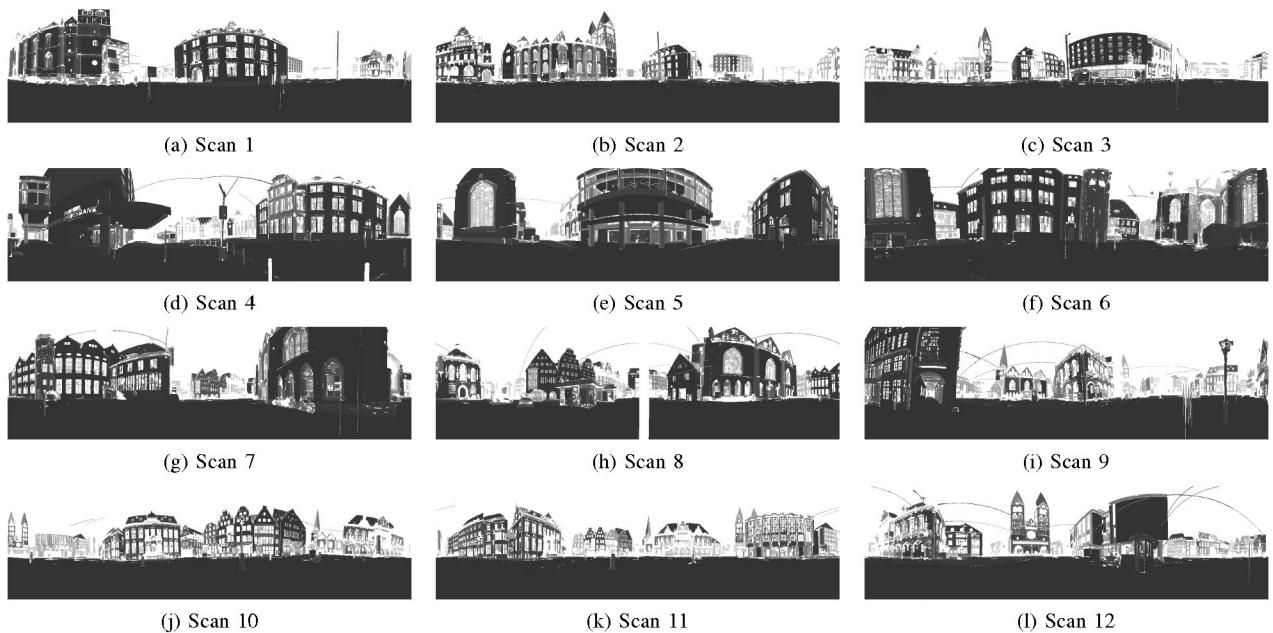
(j) Scan 10   (k) Scan 11   (l) Scan 12

Fig. 3. The scans as range-images with gray-scale based on distance. The first (0th) scan is shown in Fig. 2. The distribution of points is analyzed in Figs. 7, 8, 9.

scan-matching failure using the metric μ_u described in Sec. II-A. These two failures are further analyzed in terms of the overlap metrics in the next sub-section.

B. Evaluation of Various Scan Metrics

Figs. 7, 8, 9 show the various scan-wise and pair-wise metrics:

- Fig. 7 confirms that most of the points in the samples are located on planar patches.
- Fig. 8 shows the spatial surface area overlap metric μ_s which was described in Sec. III-A for $\Delta_t = 50$mm and $\bar{d}_s = 250$mm computed using the ground-truth registration. It shows that the pair 5-6 has the minimum value of μ_s and thus foretells the matching-failure of this pair. However, for the pair $0 - 1$, μ_s is low, though not unusually so.
- Fig. 9 shows the normalized metric μ_r in log scale. It was described in Sec. III-B and has been computed for $\delta_n = 8$ and $\Delta_\phi = 0.25°$ and $N_R = 577841$. This metric shows a clear global minimum at $0-1$ and hence foretells its matching failure. Other pairs having low values of this metric, however, successfully match.

We can thus conclude that the overlap metrics μ_s and μ_r are helpful in predicting scan-matching failure of MUMC. Their low values imply unreliability of the matching result.

V. CONCLUSION

The MUMC is algorithm is made more robust by adding an overlap test based on unmatched planes. The absolute accuracy of the algorithm is obtained for the first time by comparing its results to the given ground-truth in a relatively large outdoor urban scenario. The few matching failures were explained in terms of two independent overlap metrics which need to be used together to make effective predictions regarding MUMC's matching success for the scans involved.

(a) Scans 1 to 5.

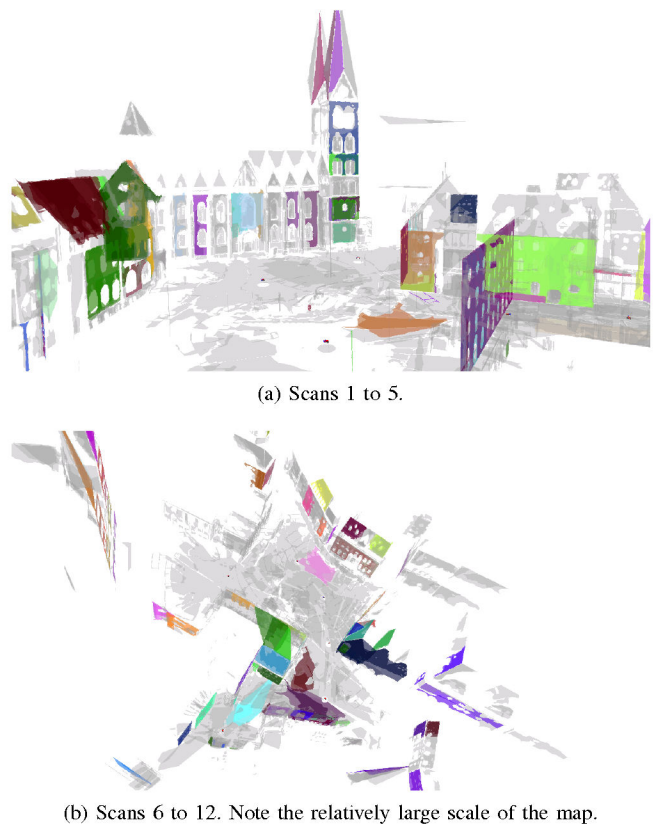(b) Scans 6 to 12. Note the relatively large scale of the map.

Fig. 5. The plane-matching sequences. Planar-patches whose correspondences have been determined using MUMC are shown in same translucent color. Unmatched planes are translucently grayed out. Note the selectivity of the MUMC algorithm in choosing the correspondences which minimize the uncertainty volume.

TABLE I

COMPARISON OF PLANE-MATCHING RESULTS. $\bar{\mu}_u$=0.23

Pair	Overlap	Rotation [deg.]			Translation [mm.]			n_ℓ	n_r	n_c	Time [s]	Succ.
		roll	pitch	yaw	t_x	t_y	t_z					
$0 \rightarrow 1$		−	−	−	−	−	−	−	−	−	−	×
Ground-truth		0.32	0.77	-157.20	-1513	41223	-189					
$1 \rightarrow 2$	μ_u=0.31	0.46	1.08	43.03	-23700	-23916	721	47	49	15	54.527	√
$\pm 1\sigma$		0.002	0.003	0.002	11	14	75					
Ground-truth		-0.08	-0.03	43.48	-23452	-24141	-640					
$2 \rightarrow 3$	μ_u=0.4	0.95	-0.51	-24.48	-9703	-36018	261	49	45	18	49.631	√
$\pm 1\sigma$		0.002	0.002	0.001	20	9	79					
Ground-truth		0.07	0.08	-24.79	-9972	-36047	-322					
$3 \rightarrow 4$	μ_u=0.24	3.20	1.09	158.52	36604	914	508	45	55	13	74.596	√
$\pm 1\sigma$		0.001	0.002	0.001	5	4	10					
Ground-truth		1.13	1.11	158.48	36575	917	1063					
$4 \rightarrow 5$	μ_u=0.24	0.32	0.73	72.70	-22163	124	151	55	54	13	125.891	√
$\pm 1\sigma$		0.003	0.007	0.007	3	1	30					
Ground-truth		0.40	0.67	72.72	-22198	123	142					
$5 \rightarrow 6$		−	−	−	−	−	−	−	−	−	−	×
Ground-truth		-1.36	0.55	-132.30	-21377	5685	-98					
$6 \rightarrow 7$	μ_u=0.35	1.46	0.59	-78.79	5315	-20390	-164	60	54	19	145.546	√
$\pm 1\sigma$		0.001	0.001	0.001	9	10	12					
Ground-truth		0.59	1.11	-78.76	5327	-20383	-189					
$7 \rightarrow 8$	μ_u=0.27	-0.80	1.31	-51.62	25266	-8089	-550	54	52	14	124.979	√
$\pm 1\sigma$		0.001	0.001	0.007	6	7	15					
Ground-truth		0.30	1.69	-51.52	25295	-8059	-447					
$8 \rightarrow 9$	μ_u=0.27	2.15	-2.28	-124.57	-8501	26675	-889	52	55	15	96.916	√
$\pm 1\sigma$		0.003	0.003	0.002	25	11	15					
Ground-truth		2.23	-2.36	-124.58	-8528	26661	-932					
$9 \rightarrow 10$	μ_u=0.39	4.29	-0.31	153.52	-11942	-22764	769	55	52	20	92.937	√
$\pm 1\sigma$		0.003	0.002	0.001	5	4	54					
Ground-truth		3.78	-0.21	153.52	-11964	-22750	753					
$10 \rightarrow 11$	μ_u=0.42	-0.09	-1.80	-78.98	-12220	20743	-796	52	52	22	80.312	√
$\pm 1\sigma$		0.004	0.004	0.003	17	27	16					
Ground-truth		0.86	-1.02	-79.05	-12226	20764	-663					
$11 \rightarrow 12$	μ_u=0.39	-0.73	0.16	-140.58	4920	-37675	159	52	51	20	70.546	√
$\pm 1\sigma$		0.002	0.002	0.001	8	7	10					
Ground-truth		0.15	-0.13	-140.58	4922	-37676	689					
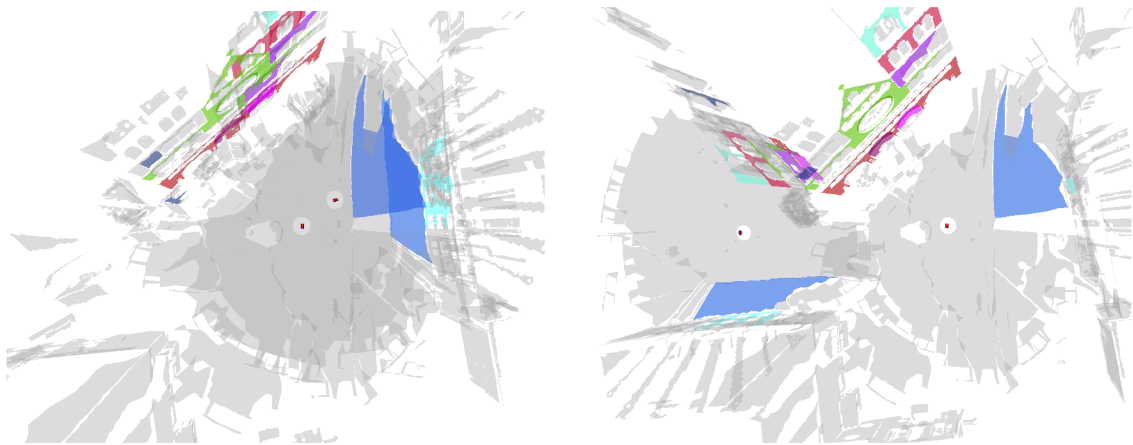
(a) The plausible-looking wrong registration. Compare with the ground-truth below. It is due to very less range-image overlap along with the recurrence of similar geometry.

(b) The ground-truth of scan-pair 0-1. The wrongly found patch-correspondences are shown in the same color.

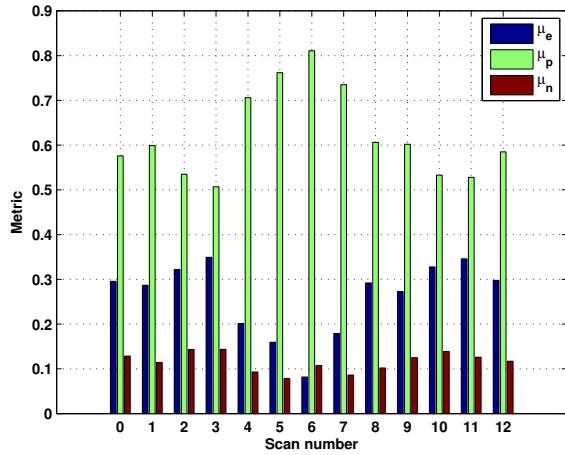Fig. 6. The wrong registration of scan-pair 0-1.

Fig. 7. Fractions of various kinds of points in the range-image. μ_e is the fraction of empty or maximum-range beams, μ_p is the fraction of points lying on planar-patches, and μ_n is the fraction of the remaining points.
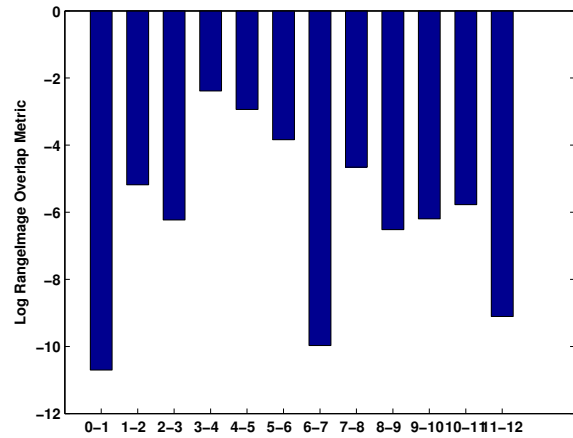

Fig. 9. Showing the range-image pairwise overlap μ_r of scans in log-scale. The chart clearly shows that scan-pair $0 - 1$ is the worst according to this metric.

[6] O. Wulf, A. Nüchter, J. Hertzberg, and B. Wagner, "Benchmarking urban six-degree-of-freedom simultaneous localization and mapping," *Journal of Field Robotics*, vol. 25, no. 3, pp. 164–180, 2008.
[7] J. Poppinga, N. Vaskevicius, A. Birk, and K. Pathak, "Fast Plane Detection and Polygonalization in noisy 3D Range Images," in *International Conference on Intelligent Robots and Systems (IROS)*. Nice, France: IEEE Press, 2008, pp. 3378 – 3383.
[8] K. Pathak, N. Vaskevicius, and A. Birk, "Uncertainty analysis for optimum plane extraction from noisy 3D range-sensor point-clouds." *Intelligent Service Robotics*, vol. 3, pp. 37–48, 2010.
[9] J. J. Craig, *Introduction to robotics – Mechanics and control*. Prentice Hall, 2005.
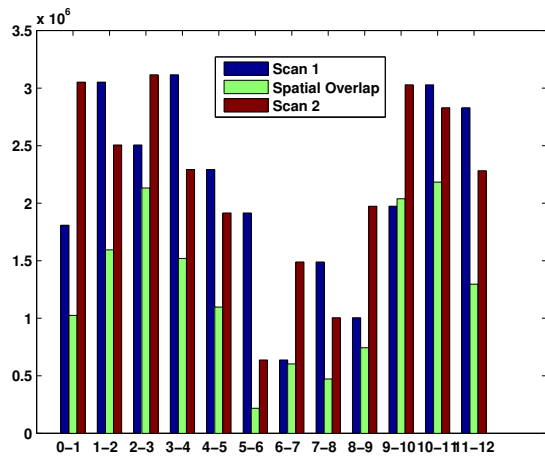
Fig. 8. Showing the absolute pairwise overlap μ_s of scans. The overlap is computed by finding corresponding voxel in scan 1 for all voxels in scan 2. Certain points in scan 2 may have the same corresponding point in scan 1, as is the case in pair 9-10. The chart clearly shows that scan-pair $5 - 6$ is the worst according to this metric.

REFERENCES

[1] K. Pathak, A. Birk, N. Vaskevicius, and J. Poppinga, "Fast Registration Based on Noisy Planes with Unknown Correspondences for 3D Mapping," *IEEE Transactions on Robotics*, vol. (in press), 2010.
[2] P. J. Besl and N. D. McKay, "A method for registration of 3-d shapes," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 14, no. 2, pp. 239–256, Feb 1992.
[3] Y. Chen and G. Medioni, "Object Modeling by Registration of Multiple Range Images," *Image and Vision Computing*, vol. 10, no. 3, pp. 145–155, 1992.
[4] K. Pathak, A. Birk, N. Vaskevicius, M. Pfingsthorn, S. Schwertfeger, and J. Poppinga, "Online 3D SLAM by Registration of Large Planar Surface Segments and Closed Form Pose-Graph Relaxation," *Journal of Field Robotics, Special Issue on 3D Mapping*, vol. 27, no. 1, pp. 52–84, 2010.
[5] S. Schuhmacher and J. Böhm, "Georeferencing of terrestrial laser scanner data for applications in architectural modeling," *International Archives on Photogrammetry and Remote Sensing (IAPRS)*, vol. XXXVI/5/W17, 2005. [Online]. Available: http://www.ifp.uni-stuttgart.de/publications/2005/schuhmacher05_venedig.pdf