

Registration of Non-Uniform Density 3D Point Clouds using Approximate Surface Reconstruction

Dirk Holz and Sven Behnke

Autonomous Intelligent Systems Group, University of Bonn, Germany

Abstract

3D laser scanners composed of a rotating 2D laser range scanner exhibit different point densities within and between individual scan lines. Such non-uniform point densities influence neighbor searches which in turn may negatively affect feature estimation and scan registration. To reliably register such scans, we extend a state-of-the-art registration algorithm to include topological information from approximate surface reconstructions. We show that our approach outperforms related approaches in both refining a good initial pose estimate and registering badly aligned point clouds if no such estimate is available. In an example application, we demonstrate local 3D mapping with a micro aerial vehicle by registering sequences of non-uniform density point clouds acquired in-flight with a continuously rotating lightweight 3D scanner.

1 Introduction

3D scanners provide robots with the ability to extract spatial information about their surroundings, detect obstacles and avoid collisions, build 3D maps, and localize. In the course of a larger project on mapping inaccessible areas with autonomous micro aerial vehicles (MAVs), we have developed a light-weight 3D scanner [13] specifically suited for the application on MAVs. It consists of a Hokuyo 2D laser scanner, a rotary actuator and a slip ring to allow continuous rotation. Just as with other rotated scanners, the acquired point clouds (aggregated over one full or half rotation) show the particular characteristic of having non-uniform point densities: usually a high density within each scan line and a certain angle between scan lines which depends on the rotation speed of the scanner (see **Figure 1**). In our setup, we aggregate individual scans of the continuously rotating laser scanner and form 3D point clouds over one half rotation (covering an omnidirectional field of view). To compensate for the MAV's motion during aggregation, we use visual odometry [18] and inertial sensors as rough estimates, and transform the scans into a common coordinate frame.

Since we use the laser scanner not only for mapping and localization but also for collision avoidance, we rotate the scanner fast resulting in a particularly low angular resolution of roughly 9° . This reduces the point density in the aggregated point clouds but increases the frequency with which we perceive (omnidirectionally) the surroundings of the MAV (2 Hz). The resulting non-uniform point densities affect classic neighborhood searches in 3D and cause problems in local feature estimation and registration. To compensate for the non-uniform point densities, we extend the state-of-the-art registration algorithm Generalized-ICP [19] to include topological surface information instead of the 3D neighborhood of points.

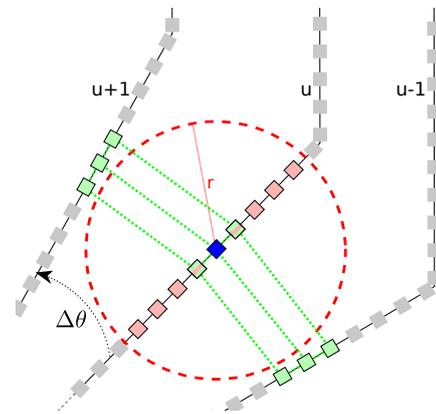


Figure 1: Classic neighbor searches in non-uniform density point clouds may only find points in the same scan line (red), whereas a topological neighborhood (green) can better reflect the underlying surface

The remainder of this paper is organized as follows: after giving a brief overview on related work in Section 2, we present our approach in Section 3. In experiments, we demonstrate the superior performance of our approach and discuss the achieved results in Section 4.

2 Related Work

2.1 Laser Scanners for MAVs

For mobile ground robots, 3D laser scanning sensors are widely used due to their accurate distance measurements even in bad lighting conditions and their large field-of-view. For instance, autonomous cars often perceive obstacles by means of a rotating laser scanner with a 360° horizontal field-of-view, allowing for the detection of obstacles in every direction [12].

Up to now, such 3D laser scanners are rarely used on

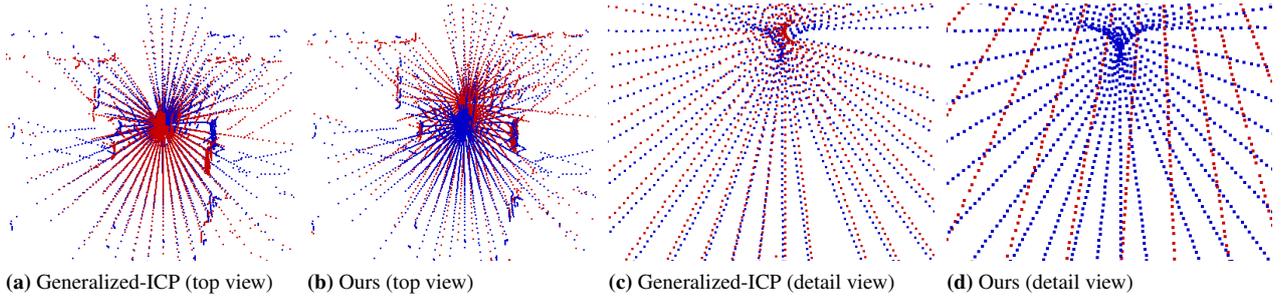


Figure 2: Registering two point clouds (red and blue points) in an example scenario: the original Generalized-ICP implementation suffers from the degradation of the underlying covariance matrices and incorrectly aligns the two point cloud origins and the individual scan lines (a,c). Our approach correctly registers the two point clouds (b,d).

lightweight MAVs—due to payload limitations. Instead, two-dimensional laser range finders are used [1, 5, 6, 9, 20, 22]. A statically mounted 2D laser range finder restricts the field-of-view to the two-dimensional measurement plane of the sensor, however. This poses a problem especially for reliably perceiving obstacles surrounding the MAV. When moving however, and in combination with accurate pose estimation, these sensors can very well be used to build 3D maps of the measured surfaces. Still, perceived information about environmental structures is constrained to lie on the 2D measurement planes of the moved scanner. In contrast, we use a continuously rotating laser range finder that does not only allow for capturing 3D measurements without moving, but also provides omnidirectional sensing at comparably high frame rates (2 Hz in our setup). A similar sensor is described by Cover et al. [4] and Scherer et al. [17]. Their MAV is used to autonomously explore rivers using visual localization and laser-based 3D obstacle perception. In contrast to their work, we use the 3D laser scanner for both omnidirectional obstacle perception and mapping the environment in 3D.

2.2 3D Scan Registration

The fundamental problem in 3D mapping and localization is registration in order to align the acquired 3D point clouds and estimate the poses (positions and orientations) where the clouds have been acquired. Over the past two decades, many different registration algorithms have been proposed. Prominent examples for estimating the motion of mobile ground robots using 3D scan registration are the works of Segal et al. [19], Nuechter et al. [14], and Magnusson et al. [10].

3D laser scanners built out of an actuated 2D laser range finder are usually (especially on ground robots) rotated comparably slower than ours to gain a higher and more uniform density of points. Most of the approaches to register such scans are derived from the Iterative Closest Points (ICP) algorithm [3]. It consists of several processing steps iteratively repeated until convergence (or until other termination criteria are met): it searches for corresponding (closest) points in the target point cloud for all points in

the source point cloud to be registered, and minimizes the distances between found correspondences w.r.t. some error metric. In addition, source and target point cloud can be subsampled, found correspondences can be rejected, and residual correspondences can be weighted. All of these steps can be modified and extended to improve the overall registration performance or to adapt the algorithm to specific data (see [15] for an overview). Whereas modifications in all steps can considerably improve registration performance, the error metric and its minimization have the strongest effect on the convergence behavior.

A particularly robust registration algorithm is Generalized-ICP [19]. It unifies the ICP formulation for various error metrics such as point-to-point, point-to-plane, and plane-to-plane. The effect of using this generalized error metric is that corresponding points in two 3D point clouds are not directly dragged onto another, but onto the underlying surfaces. For our non-uniform density point clouds, however, Generalized-ICP tends to fail since the local neighborhoods of points do not adequately represent the underlying surface (see an example in **Figure 2**). They either miss neighbors in previous and following scans or have a considerably higher density in the same scan. We adapt the Generalized-ICP approach to use extracted topological information from approximate surface reconstruction in the acquired 3D point clouds. For sparse data, as the non-uniform density point clouds acquired by our MAV, our extension can compensate for this effect and successfully register aggregated point clouds.

3 Approach

Generalized-ICP [19] is a particularly robust registration algorithm. Instead of minimizing the Euclidean distances $d_i^{(\mathbf{T})} = \mathbf{b}_i - \mathbf{T}\mathbf{a}_i$ between corresponding points \mathbf{a}_i and \mathbf{b}_i as in the original Iterative Closest Point (ICP) algorithm [2]:

$$\mathbf{T}^* = \arg \min_{\mathbf{T}} \sum_i \|d_i^{(\mathbf{T})}\|^2, \quad (1)$$

to find the transformation \mathbf{T}^* aligning two point clouds A and B , it minimizes a more general error metric:

$$\mathbf{T}^* = \arg \min_{\mathbf{T}} \sum_i d_i^{(\mathbf{T})T} (C_i^B + \mathbf{T} C_i^A \mathbf{T}^T)^{-1} d_i^{(\mathbf{T})}. \quad (2)$$

The effect is that corresponding points are not directly dragged onto another, but the underlying surfaces represented by the covariance matrices C_i^A and C_i^B are aligned. The covariance matrices are computed so that they express the expected uncertainty along the local surface normal at the point. Consequently, the convergence of Generalized-ICP degrades with inaccurate estimates of the covariances with classic neighborhood searches as illustrated in Figure 1.

3.1 Approximate Surface Reconstruction

In order to get a better estimate of the underlying covariances, we perform an approximate surface reconstruction as done in our previous work [7] in the context of range image segmentation. We traverse an organized point cloud S once and build a simple quad mesh by connecting every point $\mathbf{p} = S(u, v)$ (v -th point in the u -th scan line) to its neighbors $S(u, v + 1)$, $S(u + 1, v + 1)$, and $S(u + 1, v)$ in the same and the subsequent scan line (see Figure 1). In the original approximate surface reconstruction [7], we only add a new quad to the mesh if $S(u, v)$ and its three neighbors are valid measurements, and if all connecting edges between the points are not occluded. The first check accounts for possibly missing or invalid measurements in the organized structure. For the latter occlusion checks, we examine if one of the connecting edges falls into a common line of sight with the viewpoint $\mathbf{v} = \mathbf{0}$ from where the measurements were taken. If so, one of the underlying surfaces occludes the other and the edge is not *valid*:

$$\text{valid} = (|\cos \phi_{i,j}| \leq \cos \epsilon_\phi) \wedge (d_{i,j} \leq \epsilon_d^2), \quad (3)$$

$$\text{with } \phi_{i,j} = \frac{(\mathbf{p}_i - \mathbf{v}) \cdot (\mathbf{p}_i - \mathbf{p}_j)}{\|\mathbf{p}_i - \mathbf{v}\| \|\mathbf{p}_i - \mathbf{p}_j\|}, \quad (4)$$

$$\text{and } d_{i,j} = \|\mathbf{p}_i - \mathbf{p}_j\|^2, \quad (5)$$

where ϵ_ϕ and ϵ_d denote maximum angular and length tolerances, respectively.

In addition to the measured angle $\phi_{i,j}$, checking the length of edges plays an important role in obtaining accurate approximations of the underlying surface. For segmenting range images [7], we have learned empirical noise models for this distance threshold ϵ_d . In this paper, we adapt ϵ_d to reflect the expected distances between points on the same physical surface by exploiting the (measurable or known in advance) angles in between individual scans and points therein. In particular, we neglect all edges (and quads) being longer than the expected quad diagonal (times 1.5 to account for noise):

$$\epsilon_d = 1.5 \sqrt{2} \tan \theta. \quad (6)$$

¹Data sets are available at: http://www.ais.uni-bonn.de/mav_registration

3.2 Approximate Covariance Estimates

For estimating the covariance matrix of a point, we directly extract its local neighborhood from the topology in the mesh instead of searching for neighbors. Depending on the desired smoothing level (usually controlled with the search radius), we can extend the neighborhood to include neighbors of neighbors in the topology and ring neighborhoods farther away from the point.

Instead of computing the empirical covariances as in [19], we approximate them using the local surface normals. We compute the normal \mathbf{n}_i for point \mathbf{p}_i directly on the mesh as the weighted average of the plane normals of the N_T faces surrounding \mathbf{p}_i :

$$\mathbf{n}_i = \frac{\sum_{j=0}^{N_T} (\mathbf{p}_{j,a} - \mathbf{p}_{j,b}) \times (\mathbf{p}_{j,a} - \mathbf{p}_{j,c})}{\left\| \sum_{j=0}^{N_T} (\mathbf{p}_{j,a} - \mathbf{p}_{j,b}) \times (\mathbf{p}_{j,a} - \mathbf{p}_{j,c}) \right\|}, \quad (7)$$

with face vertices $\mathbf{p}_{j,a}$, $\mathbf{p}_{j,b}$ and $\mathbf{p}_{j,c}$. We then compute C_i^A and C_i^B as in [19]:

$$C_i^A = \mathbf{R}_{\mathbf{n}_i}^A \begin{pmatrix} \epsilon & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \mathbf{R}_{\mathbf{n}_i}^{A T}, \quad C_i^B = \mathbf{R}_{\mathbf{n}_i}^B \begin{pmatrix} \epsilon & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \mathbf{R}_{\mathbf{n}_i}^{B T} \quad (8)$$

with rotation matrices $\mathbf{R}_{\mathbf{n}_i}^A$ and $\mathbf{R}_{\mathbf{n}_i}^B$ so that ϵ reflects the uncertainty along the approximated normals \mathbf{n}_i^A and \mathbf{n}_i^B . Since the used Hokuyo laser range finder is quite accurate, we choose ϵ to be small (1 cm).

4 Experiments and Results

Registration problems considerably vary depending on the availability and quality of pose estimates. Assuming an optimal (ground truth) pose estimate, the point clouds are already aligned and a correct registration result is equal to the initial estimate. That is, any transformation applied by registration is considered an error in translation and rotation. For our non-uniform density point clouds, standard registration algorithms tend to drag individual scan lines onto another instead of aligning the environmental structures thus diverging from the optimal solution.

In order to assess the performance and reliability of our approach, and to evaluate both its divergence behavior and convergence behavior, we use two different data sets that we make publicly available¹. In a final experiment, we show that we can build 3D maps by registering non-uniform density point clouds (angular resolution of roughly 9°) acquired in-flight by a micro aerial vehicle.

4.1 Data Set and Error Metric

In order to evaluate convergence and divergence behavior of our approach for different angular resolutions, we have created a dataset of organized point clouds containing ground truth pose information. It was recorded using the

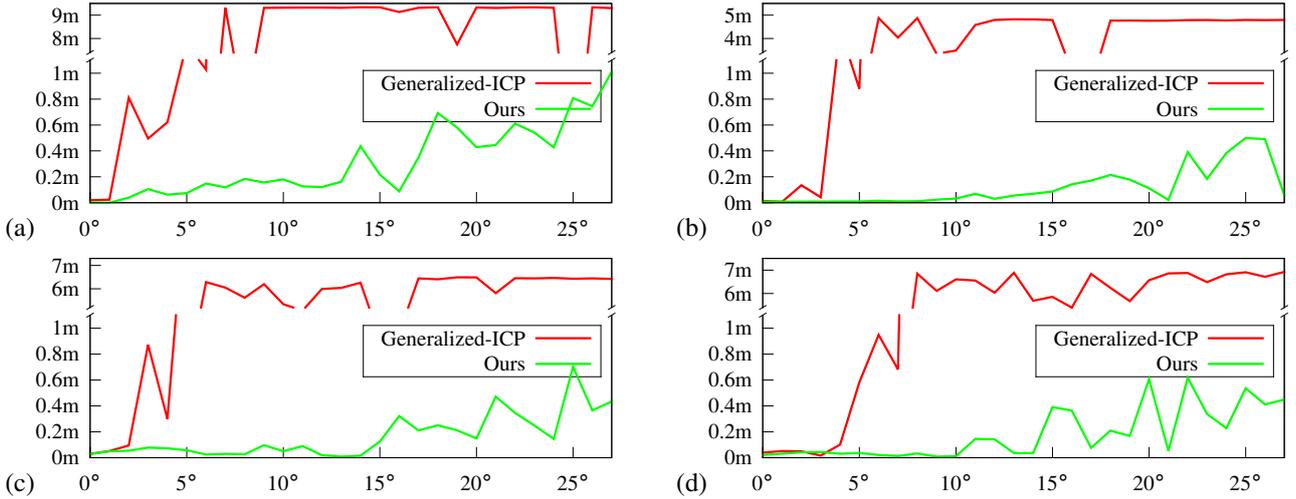


Figure 3: Evaluation of the divergence behavior on four scan pairs (a-d). Registration (translation) errors increase with an increasing angle θ between scan lines. Our approach can compensate for non-uniform densities and achieves fairly acceptable results even for larger angles between scans.

same rotating laser scanner but on a mobile ground robot standing still while acquiring 3D point clouds—thus avoiding inaccuracies in laser scan aggregation. The dataset contains point clouds from eight different poses with a total of 6890 2D laser scans acquired over multiple full rotations at each pose to obtain comparably dense point clouds. The total trajectory length between the eight poses is roughly 50 m. It was recorded by Schadler et al. [16] in the arena of the DLR SpaceBot Cup² competition for semi-autonomous exploration and mobile manipulation in rough terrain. For the dataset, we collected all 2D scan lines acquired at each of the poses, sorted them by rotation angle and re-organized the data to obtain eight full resolution ($\theta \approx 0.3^\circ$) organized point clouds. We annotated each point cloud with the ground truth pose estimate obtained from an accurate multi-resolution surfel mapping approach for dense point clouds [16]. For the experimental evaluation, we generated thinned out versions of these eight original point clouds with different angular resolutions and angles $\theta \in [1^\circ, 90^\circ]$, respectively.

For the evaluation, we measure registration success in terms of the registration error. In particular, for consecutive point clouds acquired at times i and $i + 1$, we inspect the relative deviations \mathbf{E}_i with

$$\mathbf{E}_i := (\mathbf{Q}_i^{-1}\mathbf{Q}_{i+1})^{-1}(\mathbf{P}_i^{-1}\mathbf{P}_{i+1}) \quad (9)$$

between ground truth poses \mathbf{Q} and estimated poses \mathbf{P} . As suggested in [21], we focus on the translation error e_t with

$$e_t = \|\text{trans}(\mathbf{E}_i(\Delta))\|_2, \quad (10)$$

i.e., the Euclidean distance between the estimated (relative) pose estimates.

4.2 Divergence Behavior

In order to evaluate the divergence behavior of our approach for different angular resolutions, we have chosen pairs of consecutive point clouds from the data set and registered the respective thinned out copies. In a comparative evaluation, we registered the scans of each pair using both the original Generalized-ICP algorithm and our variant with topological surface information. **Figure 3** shows the results of this comparison with decreasing angular resolution (increasing angle θ between scans).

The presented results include only four of the seven pairs of consecutive point clouds; all results are available online together with the data set. Both algorithms achieve optimal registration results for the dense point clouds with deviations from ground truth of only few centimeters. In fact, it is hard to tell whether the pose estimate used as ground truth is better or worse than the achieved alignment. For increasing angles between scan lines, the Generalized-ICP algorithm quickly starts to fail showing the aforementioned behavior of dragging individual scan lines (and the scan origins) onto another instead of aligning sensed environmental structures. In its extreme, both scan origins coincide and the maximum error in the registration results reflects the Euclidean distance between the ground truth pose estimates. Our approach achieves fairly acceptable results even for very low angular resolutions (angles between scans of $\theta > 15^\circ$). For smaller angles ($\theta \leq 10^\circ$), the resulting alignments are very accurate.

4.3 Convergence Behavior

The evaluation of the divergence behavior is particularly important for the type of data and registration problem that we address in this paper. Still, pose estimates of such high

²NimbRo Centauro at the DLR SpaceBot Cup: <http://www.ais.uni-bonn.de/nimbRo/Centauro>

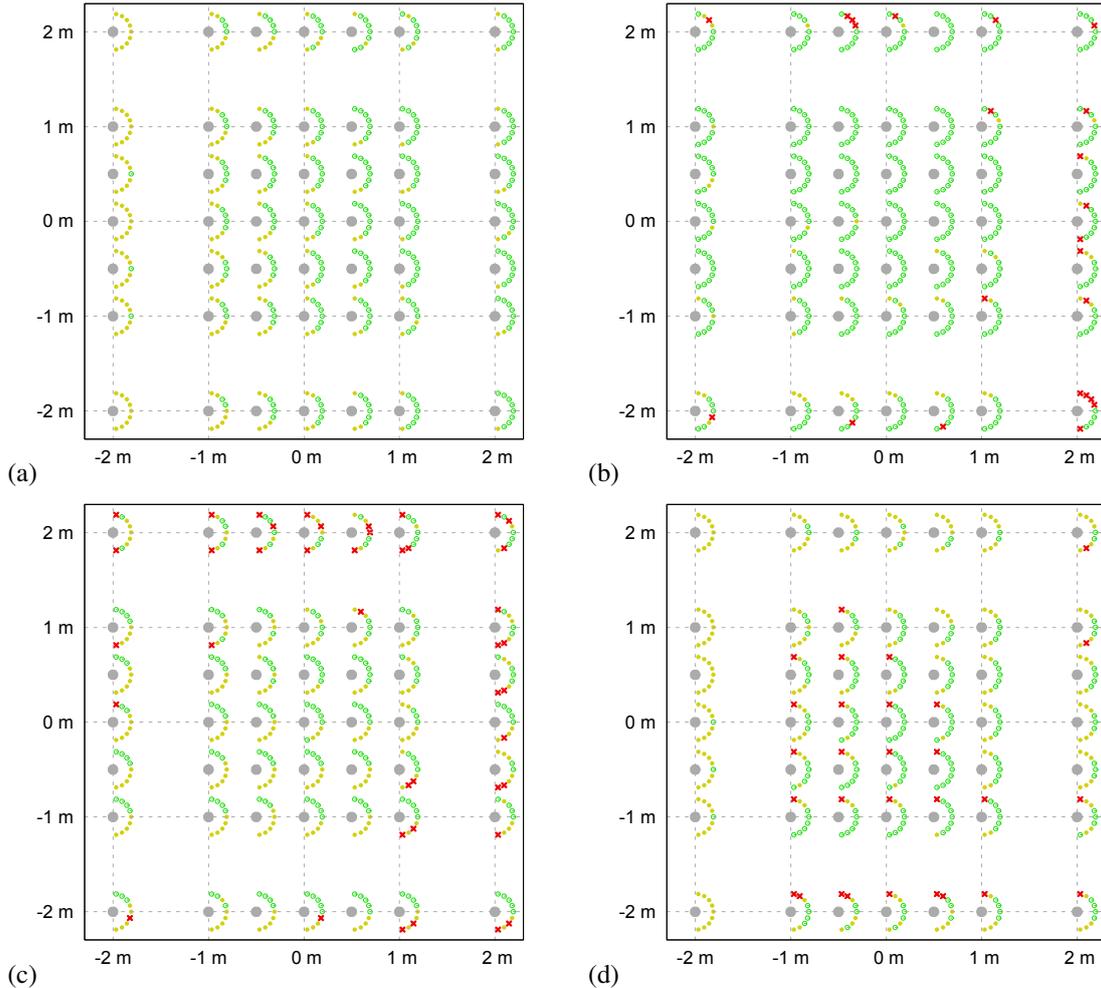


Figure 4: Evaluation of the convergence behavior on four scan pairs (a-d): registration success for different deviations from the ground truth pose estimates ($\theta = 9^\circ$). Registration success is measured w.r.t. translation error using a strict threshold (green) and a weaker threshold (yellow). Each subplot encodes seven initial orientations from -80° to 80° with 0° pointing along the x axis. Note that the original Generalized-ICP algorithm fails in all cases as can be expected from the evaluation of its divergence behavior at this angular resolution.

quality are rare, especially on a flying micro aerial vehicle. Instead, initial pose estimates tend to be noisy and may considerably deviate from the optimal alignment. Here, the central question is if and how well a registration algorithm converges to the optimal solution.

In order to evaluate the convergence behavior of our approach, we have used the same four pairs of point clouds as in the evaluation of the divergence behavior. Instead of using all available angular resolutions, we focus on the expected angular resolution of our scanner when flying (i.e., $\theta = 9^\circ$). For each scan pair, we registered the respective point clouds using our approach while using different deviations from the ground truth pose estimates. In particular, we simulate inaccuracies in the initial pose estimates using translation errors of up to 2 m along the x and y axes (i.e., the plane the robot is moving on) and rotation errors of up to 80° about the z axis (i.e., affecting the robot’s heading estimate).

In order to measure registration success for the different initial conditions, we have chosen two thresholds for the final translation error e_t (Eq. 10): a stricter one (0.25 m) and a weaker one (1 m) similar to the evaluation of registration algorithms by Magnusson et al. [11]. The intuition behind the two thresholds is that poses within the stricter translation threshold are difficult to tell apart for a human observer; poses within the weaker threshold are inaccurate but still fairly well aligned. We consider a registration as failed if the translation error exceeds the weaker threshold.

Figure 4 reports the results of the evaluation with different initial conditions for the same four scan pairs as used in the evaluation of the divergence behavior (Section 4.2). As can be seen, our approach fails in only very few cases (with high initial rotational error and/or high translational error). In the majority of registrations (even with high initial rotational and translational errors), our approach achieves an acceptable alignment even with the strict threshold.

4.4 Registration of In-flight 3D Scans

In a final experiment, we have recorded a dataset with the continuously rotating laser scanner on the flying MAV. The MAV was flying through a parking garage of 40×20 m. Overall, the dataset contains a total of 4420 2D scan lines which are aggregated to 200 3D scans (each aggregated over one half rotation of the scanner). The overall trajectory length is 73 m (traveled in 100 s). We used two fish-eye stereo camera pairs on the MAV and visual odometry [18] to estimate the motion of the MAV and aggregate the individual 2D laser scans to 3D point clouds.

In order to provide a proof-of-concept of our approach to register non-uniform density point clouds acquired by a flying MAV using a lightweight scanner, we have selected four sequences of consecutive aggregated 3D point clouds (data sets *A-D*, with 10 point clouds each) from the original data set. Using our registration approach, we aligned all pairs of consecutive point clouds in a data set, and used the resulting pose estimates to build a 3D occupancy map [8] of the surrounding environmental structures. **Figure 5** shows detailed results (including all aligned scan pairs) for data set *A*. As can be seen, our approach accurately aligns all scans and allows for constructing locally consistent 3D maps of the environment. Whereas Generalized-ICP failed for every single scan pair, our approach yields consistent alignments for all four data sets (see **Figure 6**). We report runtimes of our approach (averaged over all 199 scan pairs in the original data set, on a single core of a 2.7 GHz Intel Core i7 notebook CPU without parallelization) in **Table 1**. As can be seen, we can register point clouds as they are acquired in real time.

Table 1: Runtimes for preprocessing and registration
Preprocessing (per point cloud)

Organization & aggregation	$\ll 1$ ms
Approx. surface reconstruction	≈ 1 ms
Normal and covariance estimation	< 1 ms

Registration (per pair of consecutive clouds)

Setting up the search tree	≈ 2 ms
GICP w. approx. covariances	60 ± 17 ms

5 Conclusion and Future Work

Point clouds acquired by 3D laser scanners that are composed of rotating 2D laser range finders have the particular characteristic of having non-uniform point densities: a high density in each laser scan and a certain angle between laser scans. With increasing angle, these non-uniform densities influence neighborhood searches which in turn negatively affect scan registration. In order to register such point clouds reliably, we have extended the state-of-the-art Generalized-ICP algorithm to include topological information from approximate surface reconstructions.

Using own data sets (that we make publicly available), we demonstrated that our approach adequately registers non-

uniform density point clouds in case of good initial estimates, and in the case of badly aligned point clouds if no such estimate is available. Furthermore, we have shown that our approach can be used to build locally consistent 3D maps of the surroundings with a light-weight 3D laser scanner on a flying micro aerial vehicle.

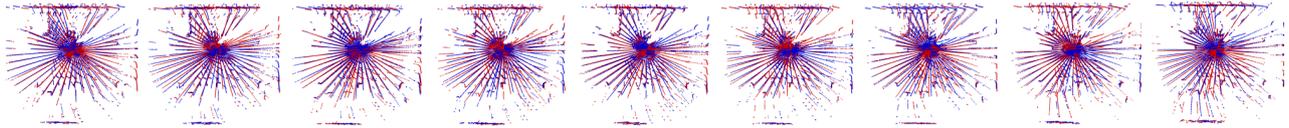
W.r.t. limitations, as a pure registration algorithm, our approach can only be used to align pairs of point clouds acquired close to each other, or to build local 3D maps over short sequences of acquired point clouds. In order to build globally consistent maps, a back-end for simultaneous localization and mapping is needed. It is a matter of future work to extend our registration approach to a fully fledged SLAM framework for mapping with micro aerial vehicles.

Acknowledgements

This work has been supported partially by grant BE 2556/7-1 of the German Research Foundation (DFG).

References

- [1] A. Bachrach, R. He, and N. Roy. Autonomous flight in unstructured and unknown indoor environments. In *Proceedings of the European Micro Aerial Vehicle Conference (EMAV)*, pages 1–8, 2009.
- [2] P. J. Besl and N. D. McKay. A Method for Registration of 3-D Shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(2):239–256, 1992.
- [3] P. J. Besl and N. D. McKay. A method for registration of 3-D shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 14(2):239–256, 1992.
- [4] H. Cover, S. Choudhury, S. Scherer, and S. Singh. Sparse tangential network (SPARTAN): Motion planning for micro aerial vehicles. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 2820–2825, 2013.
- [5] S. Grzonka, G. Grisetti, and W. Burgard. Towards a navigation system for autonomous indoor flying. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 2878–2883, 2009.
- [6] S. Grzonka, G. Grisetti, and W. Burgard. A fully autonomous indoor quadrotor. *IEEE Transactions on Robotics*, 28(1):90–100, 2012.
- [7] D. Holz and S. Behnke. Fast range image segmentation and smoothing using approximate surface reconstruction and region growing. In *Proceedings of the International Conference on Intelligent Autonomous Systems (IAS)*, 2012.
- [8] A. Hornung, K. M. Wurm, M. Bennewitz, C. Stachniss, and W. Burgard. OctoMap: An efficient probabilistic 3D mapping framework based on octrees. *Autonomous Robots*, 2013. Software available at <http://octomap.github.com>.
- [9] S. Huh, D.H. Shim, and J. Kim. Integrated navigation system using camera and gimbaled laser scanner for indoor and outdoor autonomous flight of UAVs. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3158–3163, 2013.
- [10] M. Magnusson, T. Duckett, and A. J. Lilienthal. Scan registration for autonomous mining vehicles using 3D-NDT. *Journal of Field Robotics*, 24(10):803–827, 2007.
- [11] M. Magnusson, A. Nüchter, C. Lörken, A. J. Lilienthal, and J. Hertzberg. Evaluation of 3D registration reliability and speed – a comparison of ICP and NDT. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 3907–3912, 2009.



(a) Registered scan pairs from left to right $(i, i + 1)$: (1,0), (2,1), (3,2), (4,3), (5,4), (6,5), (7,6), (8,7), (9,8).

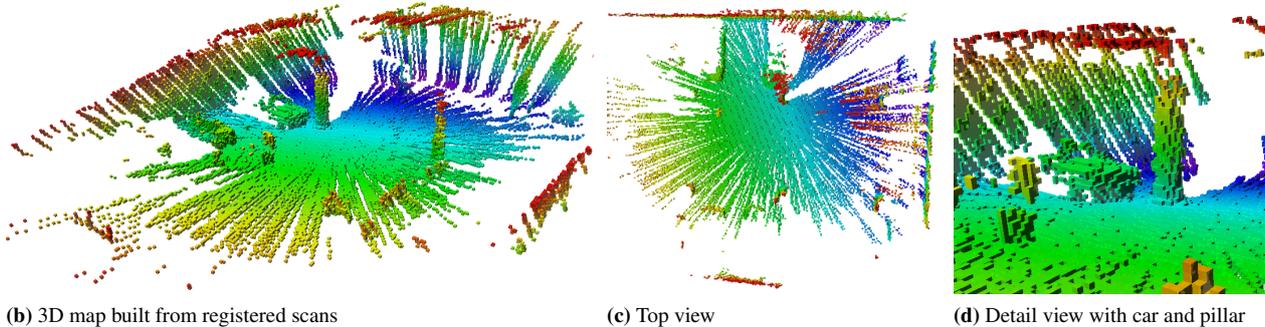


Figure 5: 3D Mapping using in-flight 3D point clouds (data set *A*): example result obtained by registering pairs of consecutive scans (a) and building a 3D occupancy map (b-d). The data set spans a trajectory of 2 m and 223 scans. For visualization, parts of walls and ceilings have been removed. Note that pairwise registration tends to drift and can only be used to aggregate short sequences of 3D point clouds. For globally consistent map building, a back-end for simultaneous localization and mapping is needed. In all map visualizations, colors encode height with the base coordinate frame rotated according to the initial orientation of the MAV (slight pitch angle when moving forward).

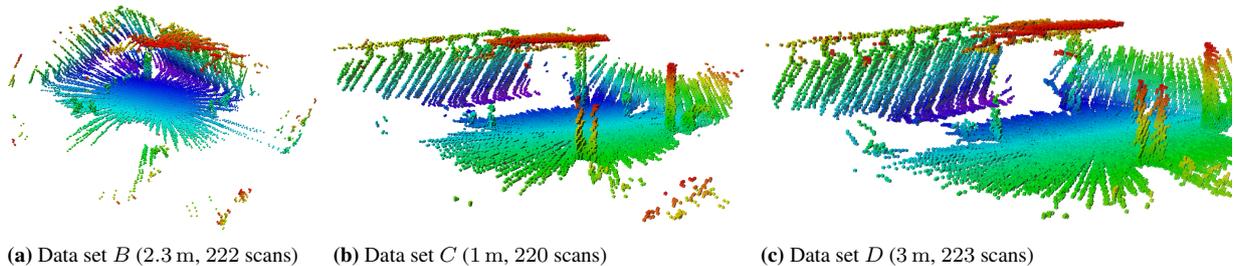


Figure 6: Visualizations of the (local) 3D maps built from the other selected data sets. All data sets are available at: http://www.ais.uni-bonn.de/mav_registration.

[12] M. Montemerlo, J. Becker, S. Bhat, H. Dahlkamp, D. Dolgov, S. Ettinger, D. Haehnel, T. Hilden, G. Hoffmann, B. Huhnke, D. Johnston, S. Klumpp, D. Langer, A. Levandowski, J. Levinson, J. Marcil, D. Orenstein, J. Paefgen, I. Penny, A. Petrovskaya, M. Pflueger, G. Stanek, D. Stavens, A. Vogt, and S. Thrun. Junior: The Stanford entry in the urban challenge. *Journal of Field Robotics*, 25(9):569–597, 2008.

[13] M. Nieuwenhuisen, D. Droschel, J. Schneider, D. Holz, T. Labe, and S. Behnke. Multimodal obstacle detection and collision avoidance for micro aerial vehicles. In *Proceedings of the European Conference on Mobile Robots (ECMR)*, pages 7–12, 2013.

[14] A. Nuechter, K. Lingemann, J. Hertzberg, and H. Surmann. 6D SLAM with approximate data association. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 242–249, 2005.

[15] S. Rusinkiewicz and M. Levoy. Efficient variants of the ICP algorithm. In *Proceedings of the International Conference on 3-D Digital Imaging and Modeling (3DIM)*, pages 145–152, 2001.

[16] Mark Schadler, Jorg Stuckler, and Sven Behnke. Rough terrain 3D mapping and navigation using a continuously rotating 2D laser scanner. *German Journal on Artificial Intelligence (KI)*, 2014.

[17] S. Scherer, J. Rehder, S. Achar, H. Cover, A. D. Chambers, S. T. Nuske, and S. Singh. River mapping from a flying robot: state estimation, river detection, and obstacle mapping. *Autonomous Robots*, 32(5):1–26, 2012.

[18] J. Schneider, T. Labe, and W. Forstner. Incremental real-time bundle adjustment for multi-camera systems with points at infinity. In *ISPRS Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences*, volume XL-1/W2, 2013.

[19] A. Segal, D. Hahnel, and S. Thrun. Generalized-ICP. In *Proceedings of Robotics: Science and Systems*, 2009.

[20] S. Shen, N. Michael, and V. Kumar. Autonomous multi-floor indoor navigation with a computationally constrained micro aerial vehicle. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 2968–2969, 2011.

[21] J. Sturm, N. Engelhard, F. Endres, W. Burgard, and D. Cremers. A benchmark for the evaluation of rgb-d slam systems. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 573–580, 2012.

[22] T. Tomic, K. Schmid, P. Lutz, A. Domel, M. Kassecker, E. Mair, I.L. Grix, F. Ruess, M. Suppa, and D. Burschka. Toward a fully autonomous UAV: Research platform for indoor and outdoor urban search and rescue. *IEEE Robotics Automation Magazine*, 19(3):46–56, 2012.