German Journal on Artificial Intelligence (KI) 28(2):93-99

Künstliche Intelligenz manuscript No.
(will be inserted by the editor)

Rough Terrain 3D Mapping and Navigation using a Continuously Rotating 2D Laser Scanner

Mark Schadler · Jörg Stückler · Sven Behnke

Received: September 19th, 2013 / Accepted: February 27th, 2014

Abstract Mapping, real-time localization, and path planning are prerequisites for autonomous robot navigation. These functions also facilitate situation awareness of remote operators. In this paper, we propose methods for efficient 3D mapping and real-time 6D pose tracking of autonomous robots using a continuously rotating 2D laser scanner. We have developed our approach in the context of the DLR SpaceBot Cup robotics challenge. Multi-resolution surfel representations allow for compact maps and efficient registration of local maps. Real-time pose tracking is performed by a particle filter observing individual laser scan lines. Terrain drivability is assessed within a global environment map and used for planning feasible paths. Our approach is evaluated using challenging real environments.

1 Introduction

In order to accomplish navigation and exploration tasks using a mobile robot, a number of problems must be solved—including environment mapping, localization, drivability assessment, and path planning. These challenges become increasingly difficult in unconstrained 3D environments, such as those encountered during planetary exploration missions. Additionally, sensing hardware should be accurate, efficient, affordable, and obey auxiliary constraints including weight and space limitations.

All authors are with
Autonomous Intelligent Systems,
Computer Science Institute VI,
University of Bonn, Germany
E-mail: schadler@iai.uni-bonn.de

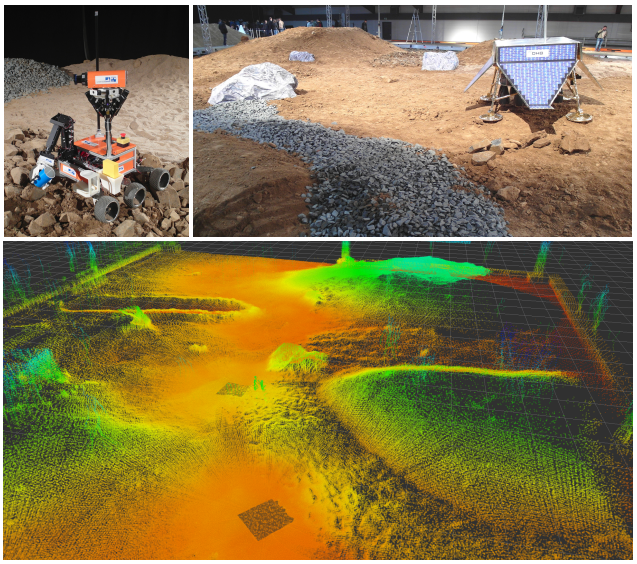

Fig. 1 Top left: Our robot Explorer for the DLR SpaceBot Cup. Top right: SpaceBot Cup arena. Bottom: 3D scans of the arena aligned with our approach.

To fulfill sensing hardware requirements, we use a continuously rotating laser scanner for environment perception. Local multi-resolution surfel maps allow for a compact 3D environment representation and facilitate efficient registration. We acquire local maps of the environment in a stop-and-go scanning scheme (Fig. 1). The 3D scans are registered to each other and aligned

into a global map using graph optimization of the 6D scan poses.

The 6D robot pose is continuously tracked during motion from the 2D scan lines of the continuously rotating laser scanner. We efficiently determine the observation likelihoods of individual scan lines in an allocentric multi-resolution surfel map, which is used for Monte Carlo localization.

Drivability assessment is performed using a 2.5D surfel map created during exploration. In this map, surfaces characteristics including incline, smoothness, and neighbor continuity can be easily extracted and evaluated for drivability—resulting in a navigation graph. Using heuristic search, we determine cost-optimal paths within the navigation graph.

Our evaluation demonstrates that accurate 3D maps can be built and the pose of the robot is tracked in real-time. Our approach has been used on our entry NimbRo Centauro at the DLR SpaceBot Cup.

2 Related Work

Most research on mapping and localization using laser scanners in 3D environments focuses on the 2D subproblem [5, 8]. Even more recent works have remained in two dimensions [4, 20]. Due to increased availability of depth sensors, research on mapping and localization in 3D has recently boomed [22, 23].

Using multi-resolution maps to maintain high performance and low memory consumption has been investigated by several groups. Hornung et al. [11], for example, implement a multi-resolution map based on octrees (OctoMap). Ryde et al. [24] use voxel lists for efficient look-up. Both of these approaches consider mapping in 3D with a voxel being the smallest map element. Similar to our approach, the 3D-NDT [22] represents point clouds as Gaussian distributions in voxels at multiple resolutions. Our multi-resolution surfel maps (MRSMap [28]) adapt the maximum resolution with distance to the sensor to incorporate measurement characteristics. Our registration method matches 3D scans on all resolutions concurrently, utilizing the finest common resolution available between both maps, which also makes registration efficient. By exploiting the neighborhood of scan lines due to the continuous rotating 3D scan acquisition, map aggregation can also be made very efficient.

Some approaches integrate scan lines of a continuously rotating laser scanner into 3D maps while the robot is moving [1, 2, 6, 21, 26]. While continuous motion is mandatory for many applications, in planetary exploration the stop-and-scan scheme is beneficial. It simplifies mapping and avoids inaccuracies that could arise with approaches performing mapping under motion.

Some works focus on localization. Khoshelham [14] proposes using solely planar objects for localization in 3D within indoor environments. Kuemmerle et al. [17] apply Monte Carlo localization in multi-level surface maps [30], which represent occupied height intervals on a 2D grid. Klaess et al. [15] model the environment in surfel maps in a fixed resolution, similar to the 3D-NDT [22]. They then localize in these maps using a tilting 2D laser by matching line elements extracted from the 2D scan lines in a particle filter framework, assuming motion of the robot in the horizontal plane. Our approach does not need to make this assumption.

Path planning for driving in planar 2D indoor environments is a well-studied topic in robotics. Hornung et al. [10] consider the variable footprint of a mobile manipulation robot for 2D path planning. They derive a multi-layered 2D representation for planning in an OctoMap. Klaess et al. [15] map the environment in 3D surfel grids using a 3D laser scanner and derive a 2D navigation map that considers the height of the robot. For navigation on non-flat terrain, several approaches generate 2D cost maps from sensor readings to essentially treat path planning in 2D [7, 9, 19]. Frequently, the terrain is modeled in elevation grid maps, on which planning can be performed [13, 18]. Recently, Stoyanov [27] proposed a wavefront-propagation path planner in 3D-NDT maps. We use efficient search-based planners and propose a robust approach to traverse between multiple 3D scans.

3 Multi-Resolution Surfel Maps

We use multi-resolution surfel maps (MRSMaps [28]) to efficiently represent environments. Octrees are the natural data structure for multiple-resolution information storage in 3D. Within octree voxels, both the surface shape parameters and surface reflectance distribution are stored as a surface element (surfel). Surfels approximate points within a voxel, which are considered normally distributed, by a sample mean and covariance.

These are stored through all resolutions in the octree. Thus, a non-leaf node maintains the statistical properties of all descendants allowing for quick sampling of the map at any resolution. In contrast to the error model of RGB-D sensors, the maximum resolution at a measured point is now determined in linear dependency of the distance of the point from the sensor. This implicitly captures decreasing sampling density with distance from the sensor caused by the constant angular resolution of our 3D laser-range finder.

The registration of local multiresolution surfel maps, which are created from 3D scans, is implemented in two main steps: surfel association and pose optimization.

Surfels are associated between maps from the finest resolution to coarser resolutions until associations have been determined for the entire map. We search for associations in a local query volume whose size is inversely proportional to the resolution of the surfel. Surfels that have been associated in a previous iteration are re-associated with the best matching surfel found within the direct voxel-grid neighbors of the previous association. Associations are made between surfels having the closest Euclidean distance in position and shape-texture descriptors [28] within the query volume.

Given a target map m_m and a 3D laser scan of an environment, we model pose optimization as finding the pose x that maximizes the likelihood $p(z|x, m_m)$ of observing the laser measurements z at the pose x in the target map m_m. Poses $x = (q, t)^T$ are represented by a translational part $t \in \mathbb{R}^3$ and unit quaternion q. After creating a map m_s from the 3D scan measurements z, we determine the observation likelihood between the source and target map given a pose x

$$p(m_s|x, m_m) = \prod_{(i,j)\in A} p(s_{s,i}|s_{m,j}), \tag{1}$$

where A is the set of surfel associations and $s_{u,v} = (\mu_{u,v}, \Sigma_{u,v})$ is the surfel v in map u. As we model surfels as normal distributions, we can easily calculate the observation likelihood of two associated surfels.

To determine the map pose x of the observation, we optimize the logarithm of the observation likelihood from Eq. (1) in two stages. We calculate an initial transform using Levenberg-Marquardt (LM) optimization and use Newton's method for refinement.

4 Simultaneous Localization and Mapping

Mapping: Graph optimization is used to globally optimize the tracked pose from 3D laser scans. For each input 3D scan, a key view $v_i \in \mathcal{V}$ (pose visualized by a coordinate frame) is extracted along the sensor view trajectory and globally aligned to a reference key view. This alignment implies a geometric constraint between the key views and is thus maintained as an edge $e_{i,j} \in \mathcal{E}$ in a key-view graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$. As an additional step, all key views deemed close are registered against each other to add edges to the constraint graph. The key views are optimized in a probabilistic pose graph using the g2o framework [16]. The optimized poses of all key views are used to create an allocentric MRSMap.

Monte Carlo Localization using scan lines: The 6D pose of the laser scanner in the allocentric map is tracked with a particle filter. This pose estimate is used to initialize the registration transform for the next key view. Compared to other filtering methods that allow for non-linear motion and measurement models such as an Extended Kalman Filter (EKF), the particle filter allows for simple integration of a scan line-to-map measurement model and also has the ability to solve the global localization problem for future applications. The general idea of a particle filter is not discussed here; see [3] for a detailed introduction.

To compute a suitable state estimate, we model the *state transition* with a time-discrete linear dynamics model. The state estimation problem is posed as the estimation of the full 6-DoF configuration of the laser scanner. Odometry is also given as input to the propagation model with assumed Gaussian noise in translational and rotational parts. For our rover-type robot, we project the propagated pose onto the drivable surface. Using an inertial measurement unit (IMU), the state propagation model also constrains the pitch and roll orientations of the robot.

The *observation model* of a scan line is similar to the model for 3D scan registration. Instead of registering single scan lines to a target surfel map, we use a particle filter and determine the likelihood of a scan line given a particle pose x. The observation model measures the alignment of the scan line $Z = \{z_i\}^n \in \mathbb{R}^3$ with n point measurements to the current target map m_t, given a

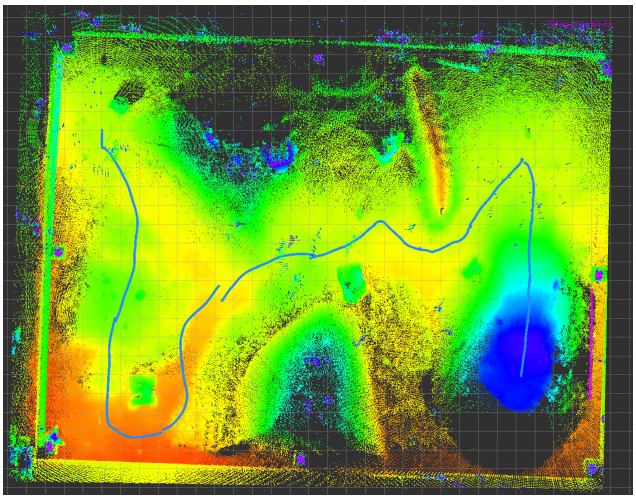

Fig. 2 Aligned 3D scans of the DLR SpaceBot Cup arena (color codes height). Localization estimate in blue.

particle pose x:

$$p(Z|x, m_t) = \prod_{i=1}^{n} p(z_i|x, A(z_i, x, m_t)),$$
$$A(z_i, x, m_t) = \underset{s_t}{\operatorname{argmin}}\ d_{\text{plane}}(s_t, T(x)z_i), \tag{2}$$

where $A(z_i, x, m_t)$ associates the transformed measurement point $z_{\text{transformed}} = T(x)z_i$ to the surfel within the target map having the smallest distance d_{plane} between the surfel plane and the target point. The potential surfels to associate are found using a volumetric query around the transformed measurement with a volume size proportional to the distance of the measurement to the sensor. If no association can be found within the region, the observation likelihood of the point is given a default no-association likelihood corresponding to the sensor models false/random measurement probability. The observation likelihood of a measurement to a surfel is given by the distance of the measurement end-point from the surfel plane under the normal distribution of the surfel.

Since we use the motion model as proposal distribution in the particle filter, the importance weight for each particle is given by its observation likelihood.

5 Navigation Planning

For navigation planning, we assume that the drivable surface can be projected onto a horizontal 2D representation, i.e., there are no overhanging structures the robot shall drive under and upon. Thus, we chose a two-dimensional grid of surfels as the data structure for drivability assessment and planning. Surface information must be stored within the grid to allow for terrain assessment and surfels are well suited to this task.

Using the allocentric localization map for navigation planning could be prone to slight misalignment errors between the 3D scans. As a basis for navigation, we fuse all 3D scans at a predefined resolution in a 2D grid. We overcome discretization effects by generating all individual MRSMaps in a common reference frame. Within a 3D scan, we merge surfels along the height direction to avoid the discretization of inclines. In regions where 3D scans overlap, multiple surfels correspond to a cell of the 2D grid. Surfels within a short distance from the scan position are used directly, because they have been created from dense high-quality measurements. Outside of this radius, we keep the surfel with higher mean height, if the height difference between the surfels is larger than a threshold. Otherwise, we choose the surfel having the highest number of measurement surface points.

5.1 Drivability Assessment

We find the drivable cells in the 2D grid using region growing. We evaluate the drivable characteristics of the current grid cell from the robot foot print in the grid map. This region is modeled by a bounding circle—an overestimate that neglects the effects of orientation on cost. The drivability of an individual surfel is determined from several features, as illustrated in Fig. 3.

Data Coverage: The number of cells containing valid surfels divided by the total number of cells within the region indicates the degree of surface knowledge. The recommended minimal region coverage depends on several factors including surface properties, safety considerations, and mapping resolution. We cannot require 100% coverage, because occlusions and missing measurements may create small gaps in the surfel grid.

Bumpiness: We define local bumpiness $b(c)$ of a cell c as the maximum difference between its mean height $\mu_z(c)$ and the mean height of all eight directly neighboring surfels $n \in \mathcal{N}(c)$:

$$b(c) = \max_{n \in \mathcal{N}(c)} \{\mu_z(c) - \mu_z(n)\}. \tag{3}$$

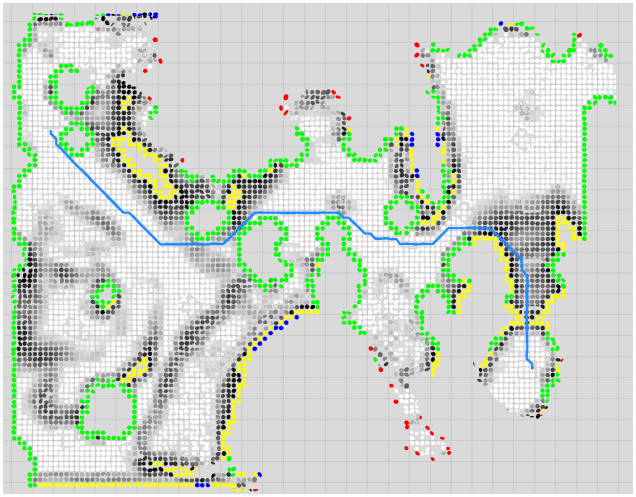

Fig. 3 Navigation planning for Fig. 2. Drivable cells are shaded by cost. Undrivable cells due to lack of data coverage (red), large incline (blue), large region bumpiness (green), and combined cost (yellow). Planned path in blue.

Bumpiness is compared to a threshold which should be set depending on the robot ground clearance, wheel size, and general planning safeness required.

Incline: Incline refers to the angle of a surface with reference to gravity, measured by the IMU. To determine the incline of a cell, we aggregate the statistical data of all surfels within the robot foot print. The normal vector of this robot-sized surfel is compared against the gravity vector for determining incline $i_r(c)$. The incline threshold should be related to the driving characteristics of the robot considering both motor capabilities and tipping resistance.

Drivability Cost: The cost of a cell $C(c)$ is dependent upon both region bumpiness and incline. The cost is calculated as $C(c) = \alpha b_r(c) + \beta i_r(c)$, where α and β are weighting parameters for bumpiness and incline, respectively. Thresholding these combined cell costs must ensure safety in case of simultaneous effects of bumpiness and incline.

5.2 Path Planning

Cost-optimal paths are efficiently planned using the A* algorithm. We use the Euclidean distance as the heuristic and consider traversability and drivability costs as edge and node costs. Fig. 3 shows an example result.

Table 1 Distances between parking garage 3D scans in m.

scans	0-1	1-2	2-3	3-4	4-5	5-6	avg err
actual	4.20	4.94	6.53	4.45	4.54	8.54	
tracking	4.30	5.26	6.39	4.44	4.85	9.40	0.29
SLAM	4.18	4.90	6.37	4.40	4.49	8.43	0.08

For plan execution, waypoints are derived from it, and a local navigation plan is updated with high frequency based on the real-time egocentric terrain model obtained from the omnidirectional depth cameras [25].

6 Experiments

Experiments were performed in the arena of the DLR SpaceBot challenge, on a benchmark dataset of the Autonomous Space Robotics Lab from the University of Toronto, and a real environment similar to those encountered during search and rescue missions.

A laser-scanner was mounted on a mobile vehicle that recorded odometry and sensor data. This data was processed in real-time using an Intel Core i7-3610QM running at 2.3 GHz with 16 GB RAM (parking lot dataset) and a robot on-board Intel Core i7-4770K (max. 3.5 GHz) with 32 GB RAM. Localization was performed using 250 particles. To ensure dense point clouds from 3D laser scans, we rotated the laser at a slow speed ($\frac{1}{15}$ Hz) when creating full 3D scans. For tracking, we wished to maximize visible space per time period and rotated the laser at 1 Hz.

Parking Garage: The parking garage environment is relatively large for an enclosed space (approximately 25×60 m) and contains various structures including vehicles, girders, support beams, and windows. Within this environment, seven full 3D scans were taken at an average distance of 5.53 m between key views. Ground truth was hand-measured and consists of the relative distances between scanning poses. Using this metric, registration yielded an average error of 8 cm with a maximum error of 16 cm. Table 1 details the tracking accuracy compared to both ground truth measurements and the graph optimized registration pose estimate. Fig. 4 shows the estimated path from scan line tracking including registration poses. Fig. 5 shows the aligned point cloud generated during this experiment.

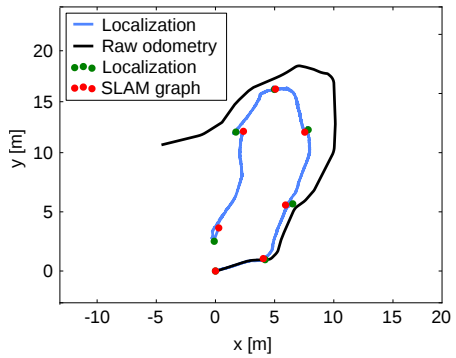

Fig. 4 Parking garage experiment. Tracked vehicle movement vs. raw odometry. Circles indicate 3D scan registration.

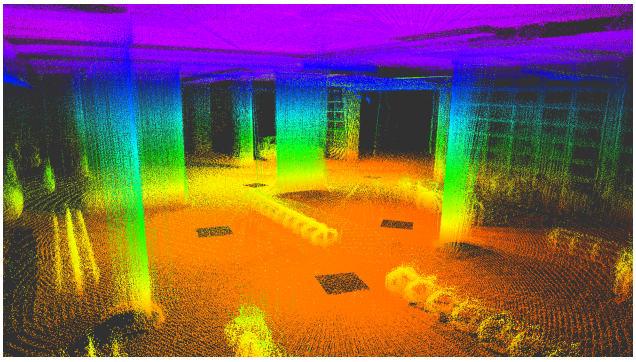

Fig. 5 Aligned scans from parking garage, colored by height.

DLR SpaceBot Cup Arena: Using data collected at the DLR SpaceBot Cup 2013, results are presented illustrating the accuracy of the mapping system and the ability to plan trajectories spanning several 3D scans. Note that during the recording, persons have walked through the arena causing spurious measurements. Here, odometry is a fused estimate of wheel odometry, IMU measurements, and visual odometry. We used fovis [12] to estimate motion from eight RGB-D cameras viewing in all directions around the robot. Shown in Fig. 1, the SpaceBot Cup arena models a rough-terrain environment similar to those potentially visited by space exploration. Within the arena, 14 3D scans were taken at an average distance of 4.33 m. To more closely resemble the open-environments in space, we have removed ceiling and wall measurements from the 3D scans before processing. Fig. 2 shows aligned scans and localization estimate and Fig. 3 illustrates navigation planning.

Autonomous Space Robotics Lab Benchmark: Using benchmark data provided by the Autonomous Space Robotics Lab (ASRL) from the University of Toronto Institute for Aerospace Studies, a comparison against alternative mapping methods is possible [29]. The a200_met dataset consists of 25 3D laser scans including stereo visual odometry. Because 2D scans during driving are not available, we used visual odometry for localization. Our mapping system produces the sharp point cloud shown in Fig. 6.

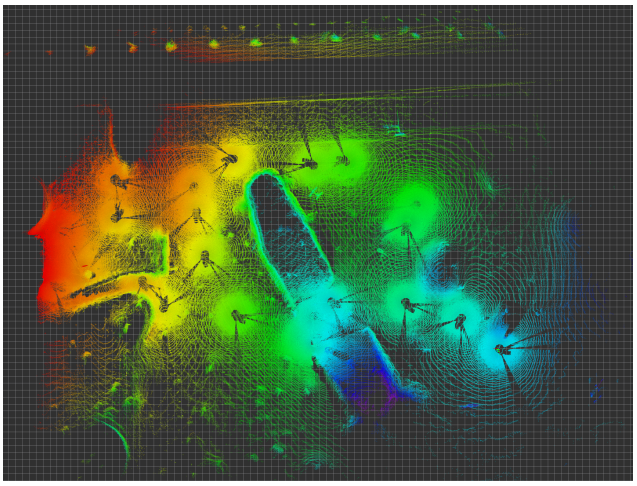

Fig. 6 Aligned scans on the ASRL a200_met dataset.

Run-Time Performance: Localization using the particle filter is limited by the laser scanner data acquisition rate. Using the Hokuyo laser scanner running at 40 Hz, a localization rate of 30–35 Hz was observed, depending on the number of valid measurements.

The SpaceBot Cup arena dataset is used for mapping performance analysis, as 3D scans are both dense and well distributed for graph optimization. With an average 3D scan containing approx. 516,000 valid points, map to map registration is performed with an average time of 1.40 s using a maximum mapping resolution of 2.5 cm. Map construction at this resolution takes an average of 0.96 s. Using a maximum resolution of 12.5 cm results in an average registration time of 0.24 s and map creation time of 0.49 s with the same sensor data.

7 Conclusion

We presented an integrated approach for robot navigation with a continuously rotating 2D laser scanner. We use local multi-resolution surfel maps to represent

3D scans which are efficiently registered and aligned in a global map using graph optimization. The motion of the robot is tracked in real time with respect to an allocentric map using a particle filter. We assess drivability based on local surface geometry and plan cost-optimal navigation trajectories. In the experiments, we demonstrated successful mapping and tracking in various environments.

Ideas for future work include integrating occupancy mapping for an improved observation likelihood model. We may also consider improved proposals in our particle filter through registration of 2D scan lines to the map.

References

1. Anderson, S., Barfoot, T.D.: Towards relative continuous-time SLAM. In: Robotics and Automation (ICRA), IEEE Int. Conf. on (2013)
2. Bosse, M., Zlot, R.: Continuous 3D scan-matching with a spinning 2D laser. In: Proc. of the IEEE Int. Conf. on Robotics and Automation (ICRA), pp. 4312–4319 (2009)
3. Chen, Z.: Bayesian filtering: From Kalman filters to particle filters, and beyond. Statistics pp. 1–69 (2003)
4. Cole, D., Newman, P.: Using laser range data for 3D SLAM in outdoor environments. In: Robotics and Automation (ICRA), IEEE Int. Conf. on (2006)
5. Dellaert, F., Fox, D., Burgard, W., Thrun, S.: Monte Carlo localization for mobile robots. In: Robotics and Automation (ICRA), IEEE Int. Conf. on (1999)
6. Elseberg, J., Borrmann, D., Nuechter, A.: 6DOF semi-rigid SLAM for mobile scanning. In: Intelligent Robots and Systems (IROS), IEEE/RSJ Int. Conf. on (2012)
7. Ferguson, D., Likhachev, M.: Efficiently using cost maps for planning complex maneuvers. In: Proc. of the ICRA Workshop on Planning with Cost Maps (2008)
8. Fox, D., Burgard, W., Thrun, S.: Markov localization for mobile robots in dynamic environments. In: Journal of Artifical Intelligence (1999)
9. Gerkey, B.P., Konolige, K.: Planning and control in unstructured terrain. In: Proc. of the ICRA Workshop on Path Planning on Costmaps (2008)
10. Hornung, A., Phillips, M., Jones, E.G., Bennewitz, M., Likhachev, M., Chitta, S.: Navigation in three-dimensional cluttered environments for mobile manipulation. In: Robotics and Automation (ICRA), IEEE Int. Conf. on (2012)
11. Hornung, A., Wurm, K.M., Bennewitz, M., Stachniss, C., Burgard, W.: OctoMap: an efficient probabilistic 3D mapping framework based on octrees. Autonomous Robots **34**, 189–206 (2013)
12. Huang, A.S., Bachrach, A., Henry, P., Krainin, M., Maturana, D., Fox, D., Roy, N.: Visual odometry and mapping for autonomous flight using an RGB-D camera. In: Proc. of the Int. Symp. on Robotics Research (ISRR) (2011)
13. Kewlani, G., Ishigami, G., Iagnemma, K.: Stochastic mobility-based path planning in uncertain environments. In: Intelligent Robots and Systems (IROS), IEEE/RSJ Int. Conf. on, pp. 1183–1189 (2009)
14. Khoshelham, K.: Automated localization of a laser scanner in indoor environments using planar objects. In: Int. Conf. on Indoor Positioning and Navigation (2010)
15. Klaess, J., Stueckler, J., Behnke, S.: Efficient mobile robot navigation using 3D surfel grid maps. In: Proc. German Conf. on Robotics (ROBOTIK) (2012)
16. Kuemmerle, R., Grisetti, G., Strasdat, H., Konolige, K., Burgard, W.: G2o: A general framework for graph optimization. In: Robotics and Automation (ICRA), IEEE Int. Conf. on, pp. 3607–3613 (2011)
17. Kuemmerle, R., Triebel, R., Pfaff, P., Burgard, W.: Monte Carlo localization in outdoor terrains using multi-level surface maps. In: Proc. of the Int. Conf. on Field and Service Robotics (FSR) (2007)
18. Kwak, J., Pivtoraiko, M., Simmons, R.: Combining cost and reliability for rough terrain navigation. In: 9th Int. Symposium on Artificial Intelligence, Robotics and Automation in Space (iSAIRAS) (2008)
19. Lee, J., Pippin, C., Balch, T.: Cost based planning with rrt in outdoor environments. In: Intelligent Robots and Systems (IROS), IEEE/RSJ Int. Conf. on (2008)
20. Lin, K.H., Chang, C.H., Dopfer, A., Wang, C.C.: Mapping and localization in 3D environments using a 2D laser scanner and a stereo camera. Journal of Information Science and Engineering 28 pp. 131–144 (2012)
21. Maddern, W., Harrison, A., Newman, P.: Lost in translation (and rotation): Fast extrinsic calibration for 2D and 3D LIDARs. In: Robotics and Automation (ICRA), IEEE Int. Conf. on (2012)
22. Magnusson, M., Duckett, T., Lilienthal, A.: Scan registration for autonomous mining vehicles using 3D-NDT. Journal of Field Robotics **24**(10), 803–827 (2007)
23. Nuechter, A., Lingemann, K., Hertzberg, J., Surmann, H.: 6D SLAM with approximate data association. In: Int. Conf. on Advanced Robotics, pp. 242–249 (2005)
24. Ryde, J., Hu, H.: 3D mapping with multi-resolution occupied voxel lists. Autonomous Robots **28**, 169–185 (2010)
25. Schwarz, M., Behnke, S.: Local navigation in rough terrain using omnidirectional height. In: Proc. German Conf. on Robotics (ROBOTIK) (2014)
26. Stoyanov, T., Lilienthal, A.: Maximum likelihood point cloud acquisition from a mobile platform. In: Advanced Robotics (ICAR), Int. Conf. on, pp. 1–6 (2009)
27. Stoyanov, T., Magnusson, M., Andreasson, H., Lilienthal, A.J.: Path planning in 3D environments using the normal distributions transform. In: Intelligent Robots and Systems (IROS), IEEE/RSJ Int. Conf. on (2010)
28. Stückler, J., Behnke, S.: Multi-resolution surfel maps for efficient dense 3D modeling and tracking. Journal of Visual Communication and Image Representation **25**(1), 137–147 (2014)
29. Tong, C., Gingras, D., Larose, K., Barfoot, T., Dupuis, E.: The canadian planetary emulation terrain 3D mapping dataset. Int. J. of Robotics Research (IJRR) (2013)
30. Triebel, R., Pfaff, P., Burgard, W.: Multi-level surface maps for outdoor terrain mapping and loop closing. In: Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (2006)