2013 IEEE/RSJ International Conference on
Intelligent Robots and Systems (IROS)
November 3-7, 2013. Tokyo, Japan

Fast 3D Mapping in Highly Dynamic Environments using Normal Distributions Transform Occupancy Maps

Jari Saarinen, Todor Stoyanov, Henrik Andreasson and Achim J. Lilienthal

Abstract— **Autonomous vehicles operating in real-world industrial environments have to overcome numerous challenges, chief among which is the creation and maintenance of consistent 3D world models. This paper focuses on a particularly important challenge: mapping in dynamic environments. We introduce several improvements to the recently proposed Normal Distributions Transform Occupancy Map (NDT-OM) aimed for efficient mapping in dynamic environments. A careful consistency analysis is given based on convergence and similarity metrics specifically designed for evaluation of NDT maps in dynamic environments. We show that in the context of mapping with known poses the proposed method results in improved consistency and in superior runtime performance, when compared against 3D occupancy grids at the same size and resolution. Additionally, we demonstrate that NDT-OM features real-time performance in a highly dynamic 3D mapping and tracking scenario with centimeter accuracy over a 1.5km trajectory.**

I. Introduction

Industrial applications of robotics are rapidly shifting from pre-programmed manipulators to highly mobile, intelligent robotic service vehicles. In order to operate efficiently in complex, large scale and realistic industrial environments, autonomous systems have to overcome numerous challenges, chief among which is the creation of a map or world model. As the availability of high-density 3D range sensors increases, so does the need for efficient mapping algorithms that produce accurate large-scale 3D models. One important, yet largely overlooked aspect of mapping in real world environments is the ability to build and maintain a model in the presence of dynamic entities. In order to meet the requirements for operating in realistic industrial scenarios, we expect from a 3D representation that it allows for accurate maps, that are compact in order to keep use of computational resources low and that it enables 3D mapping which can adapt to changes in the environment. The accuracy of a map can be measured as its capability to represent the sensor measurements [14], however, it should also enable high performance in the tasks that are carried out on the map.

In this paper we build upon the Normal Distributions Transform (NDT), which was originally developed in the context of 2D laser scan registration [3]. The central idea is to represent the observed range points as a set of Gaussian probability distributions. NDT has later been extended to 3D scan registration [5], [13]. In a recent paper [10] we introduced the Normal Distributions Transform Occupancy Maps (NDT-OM), which is a 3D mapping approach based on the NDT representation. NDT-OM extends the NDT representation with occupancy values, which also model free space. In addition, the NDT-OM is recursively updated. NDT-OM was shown in [10] to be feasible for mapping large-scale environments, even in the presence of dynamic changes. However, [10] does not account for two important aspects of the map update procedure. First, NDT cells often contain an estimated structure smaller than the cell volume. Thus, a sensor ray passing through an NDT cell should not necessarily be considered as evidence of emptiness. The second aspect concerns the update of an NDT cell covariance estimate in dynamic environments. In our previous work [10], measurements from dynamic objects are fused into the model, resulting in covariance estimates that represents a combination of both static and dynamic entities. In this paper, we address these shortcomings by: 1) introducing a probabilistic occupancy update approach that accounts for the shape of the estimated covariance in each cell during raycasting and 2) proposing a modified recursive covariance adaptation approach that allows faster adaptation to dynamic changes in the environment. In addition, we introduce a cluster-based raycasting approach for NDT-OM, which enables update rates faster than those of an efficient occupancy grid implementation of the same size and resolution. We will also demonstrate that the resolution has little effect on the localization accuracy obtained with a NDT-OM.

The evaluation was done in a highly dynamic test scenario. The effect of parameters are analysed by introducing metrics for evaluating the maps in dynamic environments. We demonstrate that the proposed NDT-OM maps are consistent and accurate enough to enable a 3D mapping and tracking system to achieve real-time pose estimates with centimeter accuracy over 1.5km trajectory.

The rest of this paper is organized as follows. Sec. 2 reviews prior work on 3D representations with special emphasis on the Normal Distributions Transform. Sec. 3 introduces the proposed improvements for mapping with NDT-OM, Sec. 4 describes the test setup, Sec. 5 analyzes the results and Sec. 6 concludes the paper.

II. Related Work

Several approaches for 3D spatial modeling have been proposed and successfully used in robotic mapping systems. Elevation grid maps — a 2.5D parametrization of space, obtained by associating a height value to cells organized in a 2D grid, have been used for outdoor robot navigation since the early years of robotics research (e.g. Bares et al.

All autohors are with Center of Applied Autonomous Sensor Systems (AASS), Örebro University, Sweden. Email: {first name}.{last name}@oru.se. This work has partly been supported by the European Commission under contract number FP7-ICT-600877 (SPENCER).

[1]). Gaussian Processes (GP) is an extension to the elevation grid approach. The available range sensor data can be used to learn the hyper-parameters of a GP, which can then be used to perform regression for any point in 2D space and obtain an interpolated height value, resulting in a continuous spatial model (Lang et al. [4]). Triebel et al. [16] propose the Multi-Level Surface (MLS) map as an extension to elevation grids which allows for multiple height values to be stored per cell. Triangle meshes are another method for spatial representation popular in the computer graphics community (Wiemann et al. [17]).

Occupancy grid mapping is one of the predominant modeling techniques for 2D [8] and 3D [7], [18] environments in robotics applications. The environment is represented by partitioning space into a regular grid and updating the probability of occupancy for each cell. An occupancy grid models both free and occupied space and adapts to dynamic changes in the environment through sequential updates making it only 3D representation, among those mentioned above, that is suitable for dynamic environments.

The Normal Distribution Transform (NDT) representation was first introduced by Biber and Strasser [3] for 2D scan matching. NDT is a grid based surface representation. The basic idea is to first accumulate sensor measurements into grid cells and then use them to compute a sample mean and a covariance for each cell. Thus, the NDT map is a set of normal distributions that describe the probability of a point being measured at a particular physical location. The NDT representation was extended to 3D (again for scan matching purposes) by Magnusson and Duckett [6] (later published in [5]). Additionally, 3D-NDT maps have been used in autonomous navigation context [12].

Stoyanov et al. in [14] compared the accuracy of 3D-NDT representation to grid-based [18] and triangle mesh based [17] representations using extensive tests conducted in different environments. An important result of [14] was that 3D-NDT typically provides the best accuracy even at low resolutions.

In [19] Yguel and Aycard proposed a method for updating the distributions in each cell using error-refinement (ER). They also proposed to use an occupancy update method based on whether a cell was observed empty or occupied. However, their approach only tracks the occupancy probability of cells containing a Gaussian component, thus disregarding the modeling of free space. In addition, the methodology proposed in [19] disregards an important property of maps with Gaussian components — namely the fact that not all of the space inside a cell containing a Gaussian component is necessarily occupied.

In our prior work [10], we introduced the Normal Distributions Transform Occupancy Maps (NDT-OM) — a 3D mapping approach using the NDT representation. NDT-OM introduced recursive covariance updates and models the occupancy of each cell using the log-likelihood approach [9], similarly to an occupancy grid. In this paper we propose three improvements to NDT-OM mapping: 1) An improved sensor model for occupancy updates; 2) a covariance adaptation scheme and; 3) an efficient raycasting approach. In addition we propose a map quality metrics for NDT-OM and evaluate the performance of NDT-OM in highly dynamic 3D mapping and tracking scenario.

III. Normal Distributions Transform Occupancy Mapping In Dynamic Environments

A. NDT-OM updating in dynamic environments

A key procedure in occupancy grid mapping [8] is to perform raycasting, i.e., tracing the ray that connects the sensor origin and sensor beam endpoint. All cells along the beam length can then be updated with evidence for being empty, while the final endpoint cell is updated with a high probability of occupancy. This procedure can be efficiently implemented using log-likelihood formulation [9]. The previously introduced NDT-OM formulation [10], applies a similar approach. Formally, a cell c_i in NDT-OM is represented with parameters $c_i = \{\mu_i, P_i, N_i, p(m_i|z_{1:t})\}$, where μ_i and P_i, are the parameters of the estimated Gaussian component, N_i is the number of points used for estimation of normal distribution parameters so far, and $p(m_i|z_{1:t})$ is the probability of the cell being occupied. The minimum amount of parameters required to define an NDT-OM cell is 11 (mean, upper diagonal of covariance, number of points and occupancy probability). In this paper we are using this representation, but proposing improvements especially considering mapping in dynamic environments.

Fig. 1 illustrates the NDT-OM update procedure. The sensor is situated in the bottom left corner, the red dots are the measurements returned by the sensor, red lines illustrate the trace of the rays and the yellow ellipses represent occupied cells with estimated normal distributions. In this paper we will focus on cases C and D.

In NDT mapping the cell often contain an estimated structure smaller than the cell volume. As a consequence, a sensor ray can pass through an occupied cell without hitting an obstacle, see Case C in Fig. 1. Using the standard beam tracing procedure as in occupancy grid mapping in this case will cause an occupied cell to be wrongly updated with evidence of being empty. To account this, in Sec. III-C we will formulate an approach that evaluates the consistency of a measurement with the map by computing the likelihood for the ray to pass through a cell occupied with a Gaussian component. This likelihood is then used to set the evidence for the occupancy update.

Another NDT update related issue in dynamic environments is that the distribution within the cell can change over time (Case D in Fig. 1). This results for the covariance estimate containing both static and dynamic points and the result does not correspond to actual state of the cell. In order to speed up the convergence of the covariance estimate, we propose a covariance adaptation parameter in Sec. III-B.

Finally, an occupancy update step is a computationally intensive operation mainly due to the raycasting. High density sensors, such as Velodyne HDL-32 used in this paper, produces hundreds of thousands of points every second.

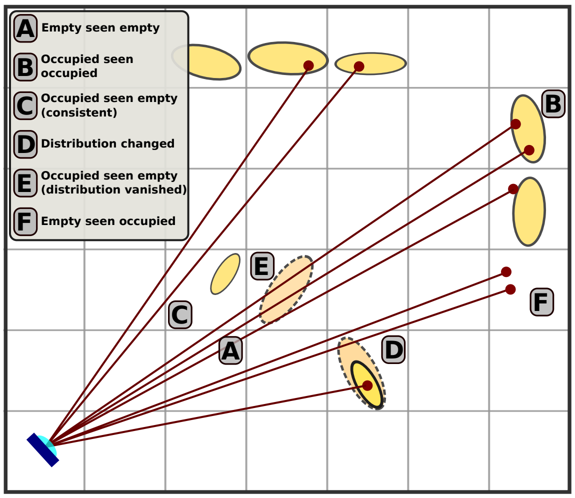

Fig. 1. Illustration of different cases of NDT occupancy mapping. Cases A, B, E and F correspond to standard occupancy mapping. Case C requires that the consistency between the map and the measurement is checked in order to update the cell correctly. Case D requires that there is a mechanism for adaptation of the existing normal distribution.

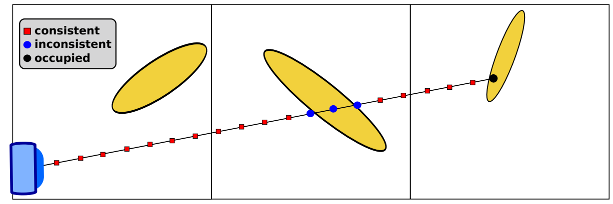

Fig. 2. NDT-OM update step. The sensor ray travels through three cells containing Gaussian components. The first cell is consistent according to map and measurement, while the second one is inconsistent.

Raycasting has to be performed for each measurement inspecting many cells for each measurement. In Sec. III-D we utilize the recursive formulation of NDT-OM, which allows us to first compute a temporary NDT map from the measurement. We then represent the laser measurement with a vector of Gaussian components obtained from the temporary map and use the mean values to update the occupancy of NDT-OM. This leads to an efficient NDT-OM update, since computing an NDT map is fast and the number of Gaussians is substantially less than number of raw sensor measurements.

B. Adaptive Recursive updating of NDT cell

The basic Recursive Covariance Update (RCU) procedure introduced in [10] computes an exact sample mean and covariance, given two sets of measurements $\{z_i\}_{i=1}^{m}$ and $\{z_i\}_{i=m+1}^{m+n}$, $z_i \in \mathbb{R}^3$. The combined estimate of the sample mean given in [10] can be written as

$$\mu_{1,m+n} = \frac{1}{m+n} T_{1,m+n}, \tag{1}$$

where $T_{1,m+n} = T_{1,m} + T_{m+1,m+n}$, $T_{1,m} = \sum_{i=1}^{m} z_i$, and $T_{m+1,m+n} = \sum_{i=m+1}^{m+n} z_i$. Similarly the combined covariance estimate is

$$P_{1,m+n} = \frac{1}{m+n-1} S_{1,m+n}, \tag{2}$$

where $S_{1,m+n} = S_{1,m} + S_{m+1,m+n} + \frac{m}{n(m+n)}(\frac{n}{m}T_{1,m} - T_{m+1,m+n})(\frac{n}{m}T_{1,m} - T_{m+1,m+n})'$, $S_{1,m} = \sum_{i=1}^{m}(z_i - \frac{1}{m}T_{1,m})(z_i - \frac{1}{m}T_{1,m})'$, and $S_{m+1,m+n} = \sum_{i=m+1}^{m+n}(z_i - \frac{1}{n}T_{m+1,m+n})(z_i - \frac{1}{n}T_{m+1,m+n})'$.

Equations (1) and (2) can be used to update exact sample mean and covariance over a sequence of observations. With respect to long-term mapping in dynamic environment we observe that 1) $T_{1,m}$ and $S_{1,m}$ grow unbounded with the number of measurements added, which could lead to numerical instabilities and overflows; and 2) the resulting distribution will carry information from all available measurements, hence, the presence of measurements from dynamic objects will result in incorrect distributions.

NDT-OM in the case of Eqs.1 and 2, with basic RCU, stores the number of points used for estimation as $N_p = n+m$. In order to speed up the convergence of the covariance estimates to correct distribution we introduce a covariance adaptation parameter M, which regulates the amount of points considered used for the estimate, that is

$$N_p = \begin{cases} n+m, & n+m < M \\ M, & n+m \geq M \end{cases}. \tag{3}$$

This approach maintains the current mean and covariance estimates unchanged, however, M sets the weight for the T, and S. Setting a small M means that the adaptation is faster, but on the other hand the history is lost quickly and on the contrary, setting M to a large value means that the adaptation is slow, but the history is better preserved.

C. New sensor model for occupancy updating with NDT-OM

In the case of an occupancy grid [8] the common practice is to update the occupancy by following the trace of a measurement from the sensor pose towards the sensor reading. All the cells along the ray are updated with low occupancy probability while the last cell is updated with high one. NDT maps have typically large cell size and often the object in the cell does not occupy the full volume of the cell. An example of the situation is illustrated in Fig. 2, where a sensor ray passes through three cells. The leftmost cell should not be updated empty, since the measurement is consistent with the map, while the center cell is clearly inconsistent. Thus, in order to consistently update the occupancy of NDT-OM the evidence should not be constant, but a function of disagreement between the observation and the map. Below, we will derive a sensor model that is based on a measure of inconsistency between the observation and the map. The derivation assumes that each cell along the sensor ray is visited once and that there is only one normal distribution per cell.

A line passing through the sensor origin $x_s \in \mathbb{R}^3$ and a measurement $z_i \in \mathbb{R}^3$ can be defined as:

$$\bar{x}(t) = l \cdot t + l_0, \tag{4}$$

where $l = \frac{z_i - x_s}{\|z_i - x_s\|} = (l_x, l_y, l_z)'$ is the direction of the line, l_0 is some point on the line and $t \in \mathbb{R}$ is a parameter. Given a normal distribution $N(\mu_i, P_i)$, the likelihood along the line is given by the function:

$$p_i(\bar{x}(t)) = exp(-\frac{1}{2}(\bar{x}(t) - \mu_i)' P_i^{-1}(\bar{x}(t) - \mu_i)). \qquad (5)$$

In order to find an inconsistency measure between the measurement and a map, we search for the maximum of Eq. 5 for each cell containing a Gaussian component

$$L_i = \underset{t}{argmax}(p(\bar{x_i}(t))) \qquad (6)$$

The maximum of Eq. 6 can be straight forwardly found by solving $\frac{dp_i(\bar{x}(t))}{dt} = 0$. Additionally, given a solution point $\bar{x_i}(\hat{t})$, where $\hat{t}$ is the solution for Eq. 6, we evaluate the likelihood of this point being an end point

$$L_i^e = exp(-\frac{1}{2}\left\|\bar{x_i}(\hat{t}) - z_i\right\|^2 / \sigma_s^2), \qquad (7)$$

where σ_s^2 is sensor noise. Finally, we set the evidence for updating the occupancy value of a cell m_i that is observed empty given a measurement z_i and map m is

$$p(m_i = 1 | z_i, m) = \begin{cases} 0.5 - \eta L_i(1 - L_i^e) & , cell\, occupied \\ \beta & , otherwise \end{cases}, \qquad (8)$$

where η is a scaling factor and β is a constant. Parameter η in Eq. (8) is used since both L_i and L_i^e are likelihoods. In order to compute the evidence of occupancy, the likelihoods need to be scaled. The scaling parameter η is selected so that $0 < \eta L_i(1 - L_i^e) < 0.5$ and it can be used as a parameter to determine the rate of adaptation to new information. When an empty cell is observed empty the evidence is set to a constant β, similar to the standard occupancy update. However, since NDT-OM also considers structures smaller than a single cell, β should depend on the ratio between the volume of the cell and the volume covered by previous observations and the current measurement. However, we do not have an efficient way to keep track of this ratio. The value of β determines the level of confidence about one observation; setting the value too low can cause that a cell that is first observed empty, but has an object will be interpreted as empty. We therefore heuristically use a constant value close to 0.5 ($\beta = 0.45$) in all our tests. A cell that is considered occupied, i.e, the cell where z_i falls into, is treated in the same way as in standard occupancy grid mapping.

Finally, in the occupancy update step we apply occupancy clamping [20] in order to prevent the map from becoming overconfident. Occupancy clamping simply limits the log-odd values of the map between given thresholds.

D. Raycasting using cluster means

The most time consuming operation in occupancy mapping is raycasting. The standard approach is to perform raycasting for all measurements updating all cells along the ray, which means that the map update step depends on the number of measurements and the distance values returned by the sensor. In this paper we use the mean values obtained from a temporary NDT map for raycasting to substantially reduce the number of raycasting operations. The basic principle is to first create a local NDT map using the most recent observations and then perform raycasting using only the means of Gaussian distributions in the local map. After the occupancy update, the means and the covariances of the local map are fused directly into the global map using RCU. The process is illustrated in Fig. 3, where the red ellipse represents a cluster obtained from the local map.

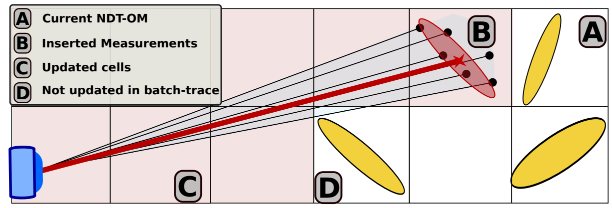

Fig. 3. Cluster mean raycasting illustrated. The red ellipse is computed in the local map from a set of recent measurements. Then the mean value (star) is used to represent all the measurement points that generated the mean value. Instead of tracing the six black lines, associated with the original measurements, we only perform one raycasting operation for the thick red line, associated with the mean value.

During this process, a local map is created at the resolution of the global map with aligned grids. This ensures that the resulting covariance and mean estimates are exactly the same as given by RCU when using directly the original scan. While raycasting using the mean value, we weight the observation with the number of points in the cluster, as if the mean value was observed by all the measurements in the cluster (the red line in Fig. 3 is used to represent all the black lines). This maintains the update of occupancy exact for the occupied cell, however, the raycasting is not exact. The mean value differs from the original measured points and thus there is a chance that a slightly different set of cells are selected for occupancy update. However, as demonstrated in the results section of this paper, the discrepancy is minimal and does not introduce inconsistencies in practice, while resulting in dramatically improved update rates.

IV. Experiment Setup

The improvements proposed in this paper are mostly related to the performance of 3D mapping with NDT-OM in dynamic environments. Accordingly, we present experiments in a dynamic environment. The test sequence is recorded at a basement of Örebro University (referred as "basement" from now on). The main area of activity is approximately 25m x 25m and covers two rooms connected by a corridor (see Fig. 4c). The test area is not particularly large-scale, but it was selected because of two reasons: 1) The area is sufficiently small so that a substantial fraction of the environment can easily be changed during the experiments and 2) the area is equipped with an industrial Automatically Guided Vehicle (AGV) laser reflector system that provides a ground truth pose estimate with approximately one centimeter accuracy[1].

The test vehicle — a training platform for AGV operators (visible in Fig. 4a and b) was controlled by the AGV on-board control system to repeat a given trajectory (see Fig. 4c) throughout the experiment. The data collected during the experiment includes odometry, ground truth pose from the navigation system and Velodyne HDL-32 sensor data. The Velodyne HDL-32 produces approximately 700,000 points per second and the data was used in our evaluation without pre-filtering (except for a height cutoff for measurements higher than 2.1m). All performance evaluations were done using a single core implementation with Intel(R) Core(TM) i7-3770K CPU @ 3.50GHz with 16GB of memory.

The experiment started with an empty, static basement (Fig. 4a). During the experiment four persons spread boxes around the basement so that gradually the layout of the area changed substantially (Fig. 4b). At the end of the experiment the boxes were collected, such that the final state of the basement was identical to its initial configuration. The total number of boxes used was forty 0.4m x 0.4m x 0.6m boxes and one hundred 0.4m x 0.3m x 0.3m boxes. The experiment lasted 36 minutes and the total trajectory length was 1552m. Additionally a 300m trajectory was collected in the static basement to create a ground truth map.

The data set includes dynamic obstacles (people moving boxes) as well as semi-static structural changes to the environment. Thus, the conditions are extremely challenging for mapping and localization algorithms, built upon a static environment assumption. The vehicle was continuously in operation and collecting data throughout the experiment.[2]

V. Tests and Results

In the following subsections we analyze the performance of NDT-OM using the dataset introduced in Sec. IV. A comparison of accuracy between 3D-NDT representation to grid-based [18] and triangle mesh based [17] representations was presented in [14]. Additionally in [10] we evaluated the covariance update against the one presented in [19]. We therefore restrict the subsequent comparison between the NDT-OM as proposed in [10] to the improved version presented in this work. For convenience in this section we refer to the improved approach as NDT-OMFG (standing for NDT-OM Fast and Generalized). In addition, a mapping performance comparison against [18] is provided in context of mapping with known poses.

A. Mapping with known pose

In this subsection, we evaluate the consistency of the produced models, when mapping a dynamic environment given known ground-truth robot poses (obtained from the reflector localization system). First, we build ground truth maps for both NDT-OM and NDT-OMFG using the dataset

[1]Automatically Guided Vehicle (AGV) system from Kollmorgen using Vehicle Master Controller (VMC 500), and Laser Way reflector based positioning system.

[2]A video showing the experiment is available at http://youtu.be/O7q1v960ZgE

a)
b)
c)
d)
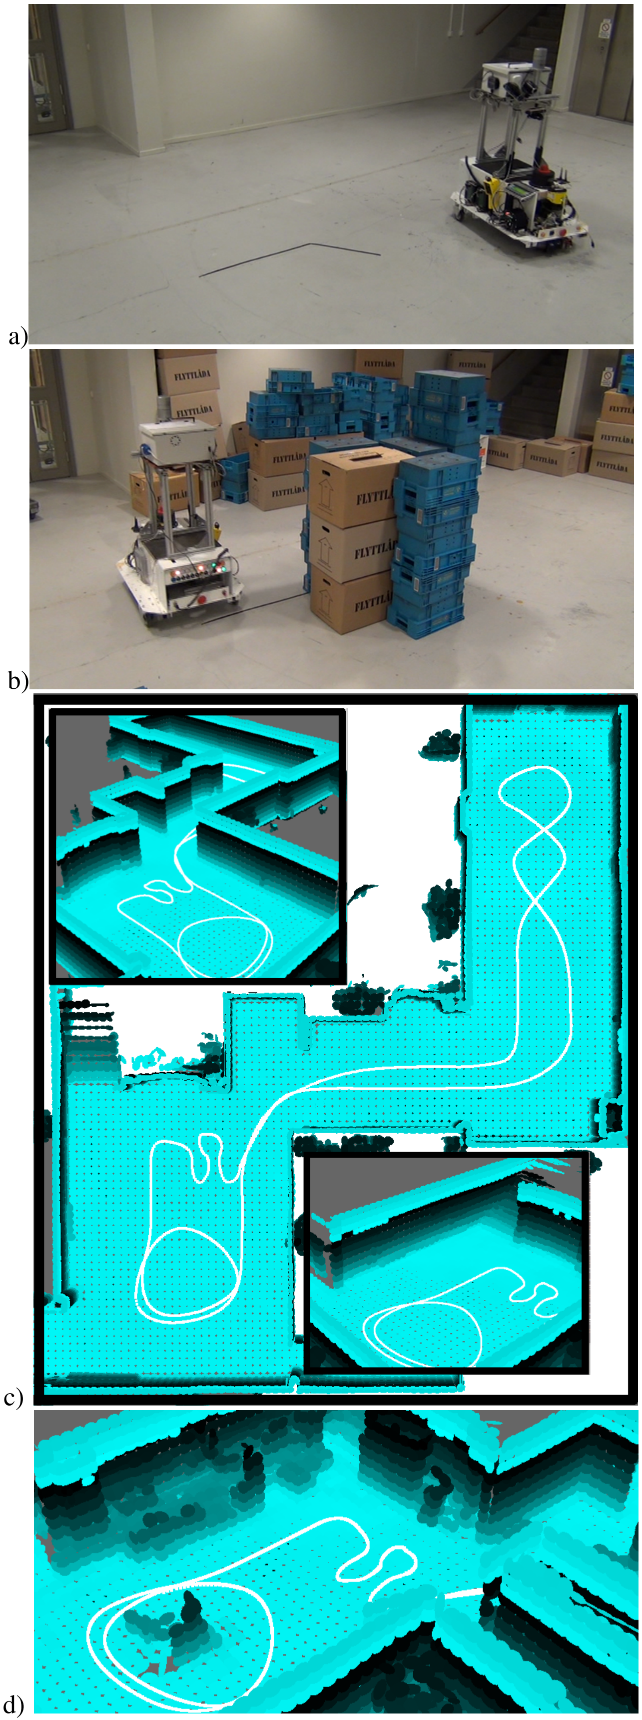

Fig. 4. Experimental setup: a) a view of the basement at the beginning (and end) of the experiment; b) during the experiment, c) 3D map of the test area with vehicle trajectory and d) state of the map during the experiment.

recorded in a static basement environment. Next, we use the dynamic data set to build a map using different occupancy clamping limits (that is the threshold for maximum occupancy/emptiness confidence a cell is allowed to reach) and for NDT-OMFG different values of the covariance adaptation constant M from Eq. 3. The dynamic environment maps produced in this manner are compared to the ground-truth map, using a consistency score defined as:

$$score = \frac{1}{\sum_i s(c_i^{gt}, c_i^{gt})} \sum_i s(c_i^{gt}, c_i), \tag{9}$$

where c_i^{gt} is an observed cell in the ground truth map, c_i is a cell from the evaluated map with same index, and

$$s(c_i^{gt}, c_i) = \begin{cases} L_2(c_i^{gt}, c_i) & , if\ c_i^{gt},\ c_i\ occupied \\ +1 & , if\ c_i^{gt},\ c_i\ free \\ -1 & , otherwise \end{cases}, \tag{10}$$

where $L_2(c_i^{gt}, c_i) = exp\left(\frac{-(\mu_1-\mu_2)^T(P_1+P_2)^{-1}(\mu_1-\mu_2)}{2}\right)$, $L_2 \in [0..1]$, measures the similarity of two Gaussian components as the $L2$-likelihood for normal distributions [13] and μ_1, P_1 are the parameters of the normal distribution in c_i^{gt} and μ_2, P_2 in c_i. The consistency measure given by Eq. 9 rewards consistently free cells as well as consistently occupied cells with similar Gaussian components. Fig. 5 summarizes the results of this test. The consistency value is computed according to Eq. 9 and the x-axis spans different occupancy limit values (values from 1 to 5 correspond to occupancy values 10, 50, 250, 1250 and 6250) scaled for better visualization. In all cases NDT-OMFG reaches higher consistency than NDT-OM. In addition, the consistency peaks at occupancy limit value of 1250 and the larger value of M results to a better consistency up to *M=1e4* (*M*=1e4 and 1e5 are nearly identical in Fig. 5). This implies that the adaptation is not needed in order for a map to converge into stationary state, which is a surprising finding. A further analysis of the effect of M is shown in Fig. 6. In this test, range scans are sequentially inserted into the map. At preselected test points (every 3 meters on the path), the consistency between the current range scan and the map is evaluated as:

$$score = \frac{1}{n^s} \sum_{i=1}^{n^s} L_2(c_i^m, c_i^s), \tag{11}$$

where c_i^s iterates over all cells in the scan model, c_i^m is the corresponding cell in the map and n^s is the number of components in the NDT models in the test scan. The averages of the obtained consistency scores for each approach are presented in Fig. 6. In this case, clearly, a smaller value of M and lower occupancy clamping threshold result in better performance.

The results for the two tests presented above are explained by the so called stability-plasticity dilemma [2]. In the first test the goal is to converge to the stationary map, which requires stability from the map. In the second test, in

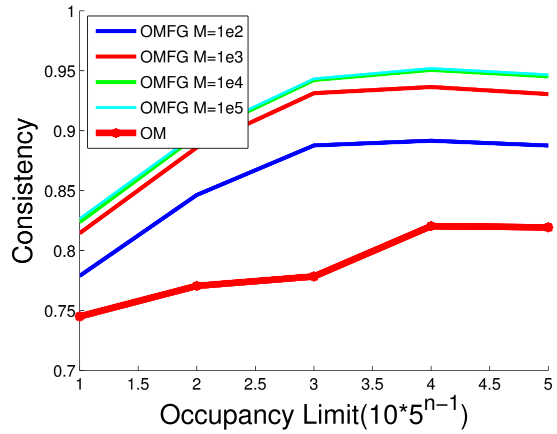

Fig. 5. Convergence to a static map using different occupancy clamping thresholds and different values of M for NDT-OMFG. The x-axis values from 1 to 5 correspond to occupancy limit values 10, 50, 250, 1250 and 6250.

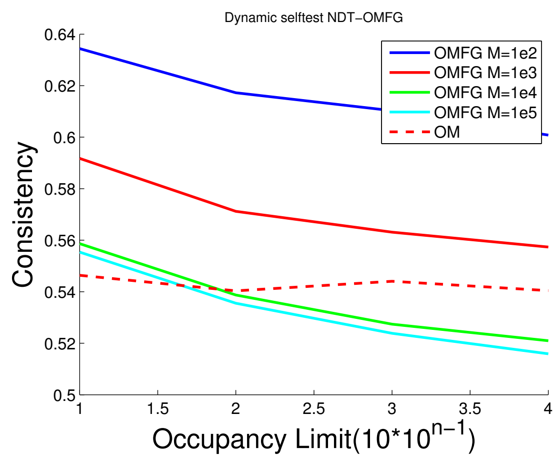

Fig. 6. Dynamic self-test. The consistency is measured as how well the measurement is explained by the map before adding the measurement into the map.

order to explain the most recent measurement in dynamic environment, the map needs a fast adaptation, which is achieved by setting a low value for both M and occupancy clamping.

As a conclusion, both the value of M and occupancy limit affect to the quality of the map. For mapping the stationary part of the environment the M and occupancy limit should be set to large values. However, in the presence of semi-dynamic changes the values should be such that the adaptation rate corresponds to the timescale of the dynamics.

Finally, Fig. 7 compares the run-time performance of different approaches. The test measures the time it takes to add one scan to the map in sequential mapping. Fig. 7 shows that NDT-OMFG is substantially (more than three times) faster than NDT-OM. For comparison, the same test was performed with a 3D occupancy grid at the same resolution as NDT maps. For testing we used a popular and efficient

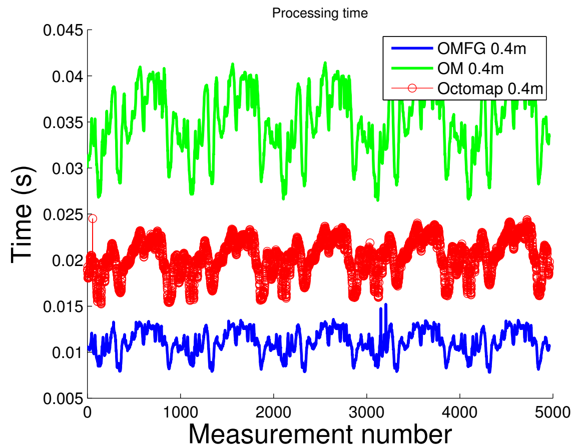

Fig. 7. Processing time comparison.

3D occupancy grid implementation, Octomap [3] [18]. Fig. 7 shows that NDT-OMFG is twice as fast as Octomap with the same resolution in our experiment.

B. Simultaneous mapping and tracking in dynamic environment

In order to demonstrate the utility of the two variations of the proposed NDT-OM mapping algorithm, it is important to evaluate how it performs in one of the most important use cases — namely as a backbone of a Simultaneous Mapping and Localization (SLAM) system. In this paper we adopt a maximum likelihood mapping principle from [15]. The map is incrementally built by adding new observation into the model at the most likely pose with respect to the map. In this paper we make use of the 3D NDT distribution-to-distribution (D2D) registration method introduced in [13]. The D2D method uses two scans, both represented as local NDT maps, and uses Newton's method to find the rigid transformation that minimizes the L_2 distance between the scans. This method has been shown to be an accurate and robust 3D registration approach [13]. In this paper, we register the NDT model of each scan directly to an incrementally built global map. The resulting pose estimate is then used to update the global map with the new scan. This naive approach does not include any loop-closing mechanism, but given the size of our test environment this is not necessary. However, in order not to confuse it into a full SLAM solution we refer to it as mapping and tracking.

Fig. 8 illustrates the Absolute Trajectory Error (ATE)[4] histograms for a) NDT-OM and b) NDT-OMFG, at a map resolution of 0.4m. The mean ATE for NDT-OM was 0.025m and for NDT-OMFG 0.024m. The ATE is computed against the ground truth pose estimate from the vehicle navigation system, which should provide us estimates within 1cm error. In both cases, the accuracy of the pose estimate is excellent and the difference between the two NDT-OM variations is small.

[3]The version of Octomap used for comparison was 1.4.3 from ros-fuerte-octomap package.

[4]ATE implementation from the Rawseed Project (http://www.rawseeds.org) was used

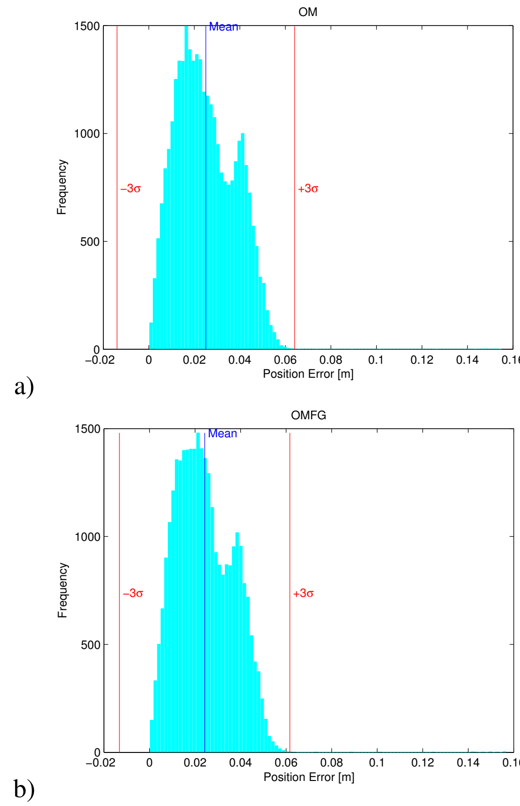

Fig. 8. Absolute trajectory error histograms for a) NDT-OM and b) NDT-OMFG in mapping and tracking trial.

Fig. 9 illustrates the performance of the NDT-OMFG-based approach at different map resolutions. The red curve shows the mean combined run-time of registration and map update, while the blue curve shows the mean ATE with one standard deviation as error bars. Fig. 9 shows that the accuracy remains at approximately 2.5cm up to a cell size of 0.5m. From 0.65m to 1.1m the accuracy is approximately 5-6cm and after that the error grows quickly. The computation time is approximately 200ms for 0.2m resolution and from 0.35m onwards it is under 100ms, which can be considered as real-time for the 10Hz sensor used in our test-case.

VI. Summary and Conclusion

This paper introduces improvements to the recently proposed Normal Distribution Transform Occupancy Map (NDT-OM) [10] approach, that allows for more efficient map updates in dynamic environments. We propose an adaptation strategy for updating the NDT-OM Gaussian components, which can be used to tune the plasticity-stability properties of the map. Furthermore, we introduce an occupancy update step that also considers the estimated shape inside each map cell. Finally, we propose a raycasting strategy for NDT-OM, based on efficient cluster mean raycasting. Together, these contributions were shown to not only improve the consistency of the NDT-OM in highly dynamic environment, but also result in more than three times faster map updates

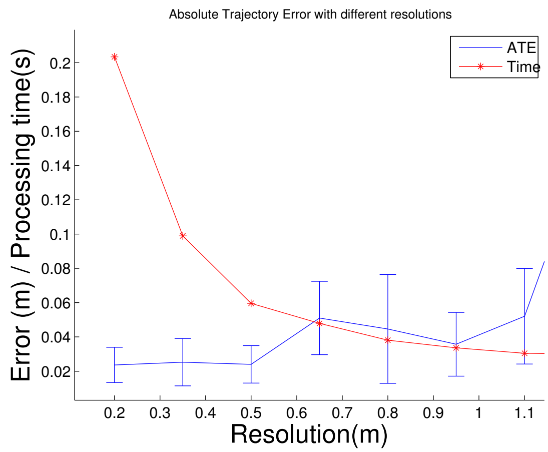

Fig. 9. Accuracy and processing time comparison with respect to the resolution.

compared to previous results. We also show that with the proposed improvements NDT-OM update rates outperform those of an efficient implementation of 3D occupancy grids even at the same resolution — at the same time NDT-OM usually works well at substantially lower resolutions than occupancy grids [11]. Finally, we demonstrate that NDT-OM can be used to obtain accurate 3D maps in real-time and despite high-level of dynamics in the environment. As a conclusion, the proposed improvements to the NDT-OM algorithm make it fully capable of producing accurate and consistent maps in environments that feature various degrees of dynamics.

All the results of this paper will be integrated as a part of future release of the oru-ros-pkg[5] open source package.

REFERENCES

[1] J. Bares, M. Hebert, T. Kanade, E. Krotkov, T. Mitchell, R. Simmons, and W. Whittaker. Ambler: an autonomous rover for planetary exploration. *Computer*, 22(6):18 –26, june 1989.

[2] P. Biber and T. Duckett. Dynamic maps for long-term operation of mobile service robots. In *Proc. of Robotics: Science and Systems (RSS)*, pages 17–24, 2005.

[3] P. Biber and W. Straßer. The normal distributions transform: A new approach to laser scan matching. In *IEEE/RSJ International Conference on Intelligent Robots and Systems(IROS 2003)*, pages 2743–2748, 2003.

[4] T. Lang, C. Plagemann, and W. Burgard. Adaptive Non-Stationary Kernel Regression for Terrain Modeling. In *Robotics: Science and Systems (RSS)*, Atlanta, Georgia, USA, June 2007.

[5] M. Magnusson, A. Lilienthal, and T. Duckett. Scan registration for autonomous mining vehicles using 3D-NDT. *Journal of Field Robotics*, 24(10):803–827, 2007.

[6] Martin Magnusson and Tom Duckett. A comparison of 3D registration algorithms for autonomous underground mining vehicles. In *Proceedings of the European Conference on Mobile Robotics (ECMR 2005)*, pages 86–91, September 2005.

[7] H. Moravec. Robot spatial perceptionby stereoscopic vision and 3d evidence grids. Technical report, CMU-RI-TR-96-34, Robotics Institute, Carnegie Mellon University, Pittsburgh, 1996.

[8] H. Moravec and A. Elfes. High resolution maps from wide angle sonar. In *Robotics and Automation. Proceedings. 1985 IEEE International Conference on*, volume 2, pages 116 – 121, mar 1985.

[9] Hans P Moravec. Sensor fusion in certainty grids for mobile robots. *AI magazine*, 9(2):61, 1988.

[5]Software available at http://code.google.com/p/oru-ros-pkg/

[10] Jari Saarinen, Henrik Andreasson, Todor Stoyanov, and Achim J. Lilienthal. Normal Distribution Transforms Occupancy Maps: Application to Large-Scale Online 3D Mapping. In *In Proceedings of IEEE International Conference on Robotics and Automation (ICRA 2013)*, 2013.

[11] Jari Saarinen, Henrik Andreasson, Todor Stoyanov, and Achim J. Lilienthal. Normal distributions transform monte-carlo localization (ndt-mcl). In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS2013)*, 2013.

[12] Todor Stoyanov, Martin Magnusson, Henrik Andreasson, and Achim J. Lilienthal. Path Planning in 3D Environments using the Normal Distributions Transform. October 18–22 2010.

[13] Todor Stoyanov, Martin Magnusson, Henrik Andreasson, and Achim J Lilienthal. Fast and accurate scan registration through minimization of the distance between compact 3d ndt representations. *The International Journal of Robotics Research*, 31(12):1377–1393, 2012.

[14] Todor Stoyanov, Martin Magnusson, and Achim J. Lilienthal. Comparative Evaluation of the Consistency of Three-Dimensional Spatial Representations used in Autonomous Robot Navigation. *Journal of Field Robotics*, 2013. to appear.

[15] Sebastian Thrun, Wolfram Burgard, and Dieter Fox. A probabilistic approach to concurrent mapping and localization for mobile robots. *Autonomous Robots*, 5(3):253–271, 1998.

[16] R. Triebel, P. Pfaff, and W. Burgard. Multi-level Surface Maps for Outdoor Terrain Mapping and Loop Closing. In *Proc. of IEEE/RSJ Int. Conf. on Intelligent Robots and Systems. IROS 2006.*, pages 2276 – 2282, 2006.

[17] T. Wiemann, A. Nüchter, K. Lingemann, S. Stiene, and J. Hertzberg. Automatic construction of polygonal maps from point cloud data. In *Proc. 8th IEEE Intl. Workshop on Safety, Security, and Rescue Robotics (SSRR-2010), Bremen, Germany*, 2010.

[18] K. M. Wurm, A. Hornung, M. Bennewitz, C. Stachniss, and W. Burgard. OctoMap: A probabilistic, flexible, and compact 3D map representation for robotic systems. In *Proc. of the ICRA 2010 Workshop on Best Practice in 3D Perception and Modeling for Mobile Manipulation*, Anchorage, AK, USA, May 2010. Software available at octomap.sf.net.

[19] M. Yguel and O. Aycard. 3D Mapping of Outdoor Environment Using Clustering Techniques. In *Tools with Artificial Intelligence (ICTAI), 2011 23rd IEEE International Conference on*, pages 403 –408, nov. 2011.

[20] Manuel Yguel, Olivier Aycard, and Christian Laugier. Update policy of dense maps: Efficient algorithms and sparse representation. In *Field and Service Robotics*, pages 23–33. Springer, 2008.