Comparative Evaluation of Range Sensor Accuracy in Indoor Environments

Todor Stoyanov, Athanasia Louloudi, Henrik Andreasson and Achim J. Lilienthal
Center of Applied Autonomous Sensor Systems (AASS), Örebro University, Sweden

***Abstract*— 3D range sensing is one of the important topics in robotics, as it is often a component in vital autonomous subsystems like collision avoidance, mapping and semantic perception. The development of affordable, high frame rate and precise 3D range sensors is thus of considerable interest. Recent advances in sensing technology have produced several novel sensors that attempt to meet these requirements. This work is concerned with the development of a holistic method for accuracy evaluation of the measurements produced by such devices. A method for comparison of range sensor output to a set of reference distance measurements is proposed. The approach is then used to compare the behavior of three integrated range sensing devices, to that of a standard actuated laser range sensor. Test cases in an uncontrolled indoor environment are performed in order to evaluate the sensors' performance in a challenging, realistic application scenario.**

I. Introduction

In recent years, a multitude of range sensing devices have become available at more affordable costs. Notably, 2D laser range sensors have demonstrated reliable performance and therefore have been widely used for both research and industrial applications. As the complexity of application scenarios considered in mobile robotics increases, so does the use of 3D range sensors. Although precise commercial 3D laser sensors are available, their prohibitively high cost has limited their use. Actuated laser range finders (aLRF) have thus been the most widely used 3D range sensors in the mobile robotics community. Usually aLRF sensors utilize commercially available 2D laser sensors of high accuracy and well known precision, resulting in a reliable measurement system. Nevertheless, aLRF sensors have several disadvantages — the lack of commercial availability of an integrated system, slow refresh rates on the order of 0.1Hz, a high weight and a high number of moving parts.

Several competing sensor technologies that attempt to solve the problems of aLRF systems have been proposed. Recently, time-of-flight (ToF) and structured light cameras have become more available and more widely used. ToF sensors operate by emitting modulated infrared light and measuring the phase shift of the reflected signal. Typically, ToF cameras can deliver dense range measurements at high frame rates of up to 50Hz. Structured light cameras can produce similar measurements, using a projected pattern that is observed by a CCD camera with a known baseline distance. Although the ToF and structured light sensor technologies have a lot of potential for use in mobile robotics, both are affected by several error sources. It is thus very important to evaluate the accuracy of these emerging sensors, compared to that of the actuated LRF.

Over the course of development of ToF sensors, several works have evaluated their accuracy, in the context of sensor calibration. Kahlmann [1] proposes several calibration routines for the SwissRanger SR-2 and SR-3000 cameras. In order to evaluate their effect on the range measurement quality, he proposes to scan a flat wall and compare offsets from the expected distance. Later, several precisely engineered optical calibration setups, as well as a calibration track line are used. Linder et. al. also use standard optical calibration setups (checkerboard images), in order to calibrate PMD ToF cameras [2]. Fuchs et. al. [3] evaluate distance accuracy using a modified checkerboard pattern and a pair of ToF cameras mounted on an industrial manipulator. Chiabrando et. al. [4] perform a comprehensive evaluation of the SR-4000 ToF camera, also using a pre-fixed optical calibration pattern and a tilted calibration board. While these set-ups constitute an important test case for ToF cameras, they do not capture the complexity of typical uncontrolled environments encountered by mobile robots. May et. al. [5] consider several real-world environments and measure the distance to well known objects in these set-ups, resulting in a more complete accuracy evaluation. Several works have also implicitly evaluated the accuracy of ToF ranging systems by considering their utility in the context of mapping [6] [7], obstacle avoidance [8] and object modeling [9].

Prior work dealing with the accuracy of ToF systems has uncovered complex error sources. Features of the environment, such as dark textures, sharp edges, foreground objects and distant bright objects, all introduce measurement errors in the obtained ranges. Although investigations into structured light camera accuracy are lacking, similar complexity of the error sources may be expected. Thus, the evaluation of both ToF and structured light cameras in strictly controlled environments and engineered test scenarios may not properly reflect on their performance in a real world setting. It is therefore important to develop a procedure for a holistic evaluation of novel range sensing devices for the purposes of mobile robotics applications.

This work extends a recently proposed method for spatial representation accuracy evaluation [10] to perform a comparison of range sensor measurements. Three integrated 3D range sensors — the SR-4000 and Fotonic B70 ToF cameras and the Microsoft Kinect structured light camera are compared to a standard aLRF sensor. This article proceeds with a description of the accuracy evaluation approach. Section III describes the test setup and evaluation environments considered. Section IV presents an analysis of the obtained results, followed by a summary of the major contributions.

II. Accuracy Evaluation

In order to objectively compare the performance of different 3D range sensors, a well-defined framework is needed. It is also important to define well-formed comparison criteria and obtain meaningful, easy to interpret statistical results. The main objective of this work is to propose a method to compare the accuracy of range measurements in an uncontrolled environment, closely related to the application scenario in which the sensors are to be used.

The information returned by 3D range sensors consists of discrete samples from a real world environment. Due to differences in sensing methodology, position and orientation of the devices, the discrete points obtained by each sensor do not correspond to precisely the same physical locations. Thus, a simple, per-point distance error would not be an accurate measure of the discrepancies between different sensors. Rather, in order to compare a test point set P_t to a reference measurement P_r, a continuous model $\mathcal{M}(P_r)$ should be used. The model $\mathcal{M}$ should be capable of evaluating the occupancy of any arbitrary point in the field of view of the sensor. Several alternative probabilistic modeling techniques could be used for constructing $\mathcal{M}$. Building on previous results in spatial representation accuracy evaluation [10], the Three-Dimensional Normal Distributions Transform (3D-NDT) is chosen for constructing the reference model $\mathcal{M}(P_r)$.

A. The Three-Dimensional Normal Distributions Transform (3D-NDT)

The Normal Distributions Transform was originally developed in the context of 2D laser scan registration [11]. The central idea is to represent the observed range points as a set of Gaussian probability distributions. In order to better account for outliers in the data, a mixture of a Gaussian distribution and a uniform distribution may also be used [12]. NDT has been extended to three dimensions [13] and applied to the domains of 3D scan registration [14] and loop detection [15], as well as change detection [16] and path planning [17]. One of the key advantages of the 3D-NDT is the fact that it forms a piecewise smooth spatial representation, resulting in the existence of analytic derivatives. Consequently, standard optimization methods can be employed to produce state-of-the-art registration of 3D point clouds using the 3D-NDT representation [14].

Assuming that a set of n point samples $P = \{(x_i, y_i, z_i)\}$ has been drawn from a Gaussian distribution $\mathcal{N}(\boldsymbol{\mu}, \Sigma)$, the maximum likelihood estimates of the covariance and mean can be obtained from the observations:

$$\boldsymbol{\mu} = \frac{1}{n} \sum_{i=0}^{i=n} \begin{bmatrix} x_i \\ y_i \\ z_i \end{bmatrix} \tag{1}$$

$$M = \begin{bmatrix} x_0 - \mu_x & .. & x_n - \mu_x \\ y_0 - \mu_y & .. & y_n - \mu_y \\ z_0 - \mu_z & .. & z_n - \mu_z \end{bmatrix} \tag{2}$$

$$\Sigma = \frac{1}{n-1} M M^T \tag{3}$$

The 3D-NDT partitions space into a set of disjoint voxels and fits a Gaussian distribution to the points in each cell. In order to obtain a well-fitting estimate, it is important to choose a discretisation strategy that would permit for each distribution to be well supported by the observations in P. The most typical implementation of 3D-NDT utilizes a regular grid, though the use of an OcTree data structure of irregular cell size is also possible [15]. As shown previously, the spatial indexing used has an effect on the representation accuracy of the 3D-NDT [10], as well as on the speed and accuracy of the registration algorithms that utilize it [15].

B. Comparing Range Sensor Accuracy using the 3D-NDT

Prior work in evaluating the discriminative capabilities of the 3D-NDT [10] has demonstrated that it is a robust spatial representation with better accuracy than the standard occupancy grid map. Thus, the 3D-NDT is used in this work to represent the ground truth models of the environment. An important feature of the 3D-NDT is that due to its probabilistic foundations, it can be used to evaluate the occupancy likelihood of any point in space. Thus, given a 3D-NDT model $\mathcal{M}_{ndt}$ and a query point x, a probability of occupancy can be estimated. The PDFs in each cell of $\mathcal{M}_{ndt}$ can be treated as surface generation processes, thus the probability of observing $\boldsymbol{x}$, generated from $\mathcal{N}(\boldsymbol{\mu}, \Sigma)$ is:

$$p(\boldsymbol{x}|\mathcal{N}(\boldsymbol{\mu}, \Sigma)) = \frac{1}{(2\pi)^{3/2}\sqrt{|\Sigma|}} e^{-\frac{(\boldsymbol{x}-\boldsymbol{\mu})^T \Sigma^{-1} (\boldsymbol{x}-\boldsymbol{\mu})}{2}} \tag{4}$$

As the PDFs in the model $\mathcal{M}_{ndt}$ are largely disjoint, the probability $p(\boldsymbol{x}|\mathcal{M}_{ndt})$ can be evaluated by only using the closest 3D-NDT component and directly applying Equation 4. In order to declare the point in space as occupied, the probability $p(\boldsymbol{x}|\mathcal{M}_{ndt})$ can be tested against a prefixed threshold. Previous work has empirically determined that a conservative threshold value of 0.1 results in a very accurate model. For the purposes of the current evaluation, an OcTree version of the 3D-NDT is used, as it offers stable performance over varying discretisation sizes.

The proposed evaluation procedure assumes a reference 3D sensor, which produces a ground truth point cloud P_r. First, the 3D-NDT model of the reference point cloud $\mathcal{M}_{ndt}(P_r)$ is constructed. The point set output of a test sensor P_t is then considered as a set of positive examples pos_t, produced by a binary classifier associated with the sensor. Next, a set of corresponding negatively classified points neg_t can be generated, using P_t. For each point in P_t, a random offset of between 0.1 and 1 meters along the beam direction is subtracted. This method of generating free space samples is consistent with the sensor output, as each point in neg_t is located between the sensor origin and a corresponding reported obstacle point from pos_t. Next, all points in pos_t and neg_t are tested for occupancy, using the ground truth model $\mathcal{M}_{ndt}$. The samples in pos_t correctly labeled positive represent the true positives. The negatively labeled samples from pos_t are false negatives. Similarly, the positively labeled points from neg_t constitute

false positives, while the remaining points are true negatives. The true positive rate (tpr) can then be calculated as the number of true positives over all samples labeled positive by the classifier (true positives and false negatives). The false positive rate (fpr) is defined as the number of false positives over all negatively labeled points (false positives and true negatives). Thus, the point set P_t can be represented by a corresponding point $p = (tpr, fpr)$ on a Receiver Operating Characteristic (ROC) plot — a standard method of evaluating and comparing binary classifiers in machine learning.

III. Evaluation Methodology

Using the evaluation procedure proposed in Section II-B, any number of sensors can be compared against a ground truth reference scan. In this work, four different 3D range sensors were mounted on a mobile platform. The sensor setup and sample output point sets are shown in Figure 1. In order to use the evaluation methodology previously presented, the output point clouds have to be placed in a common reference system.

One possible method for extrinsic sensor calibration would be to consult the data sheets and then precisely measure the relative displacements and orientations. The main disadvantages of this approach are that it lacks robustness to changing the sensor configuration and that small errors in the orientation can be made easily, resulting in large discrepancies of the obtained measurements. A common strategy from photogammetry is the use of a calibration pattern. Due to the large amounts of noise present and low resolution of some of the sensors, this also is not a feasible alternative. Thus, the sensors were registered together directly using the obtained range measurements and the 3D-NDT registration algorithm [15]. The advantages of this approach are that using several test scans rich in features minimizes the effect of sensor noise, resulting in a more accurate calibration.

Several procedures for intrinsic parameter calibration of ToF sensors have been proposed [1]–[3]. Prior work on the SR-4000 sensor [4] has however concluded that most of these procedures have been performed by the manufacturer and are not necessary. Notably, the evaluation scenario with known ground-truth information (discussed in Section IV-A) could be further used to improve the calibration parameters and internal noise models of the evaluated sensors. While such an investigation could be of interest in correcting the internal sensor errors, it would not offer insights into environment-induced discrepancies and will not be further pursued in the current work. Thus, an out-of-the-box testing approach was used in this work, with no further sensor-specific filtering or noise removal techniques employed. Factory settings were used for all sensors, with automatic parameter estimation enabled whenever possible. The Fotonic camera was used in 25Hz mode, with four samples used to estimate each pixel depth and two alternating modulation frequencies. Finally, in order to provide for a more objective evaluation, only the overlapping parts of the sensors fields of view were used.

Once all sensor data has been reliably registered in a common reference frame, the accuracy of the range measurements can be evaluated. As ground truth geometrical models of uncontrolled indoor environments are generally not available, this work uses as a reference measurement set the actuated laser points. In order to demonstrate the feasibility of this assumption, an initial test of the evaluation procedure was performed in a controlled environment. Figure 2 shows the environment used to perform an evaluation with a known ground truth. Three standardized boxes of known dimensions were placed in front of a white wall. In order to increase the complexity of the scene and test for robustness to texture, a black and white calibration board was also placed on the wall. The sensor setup was placed at three meters from the wall and used to collect data from all sensors. A ground truth model of the environment was created by hand using 3D modeling software and sampled uniformly to obtain a reference point cloud P_r. Next, a more extensive real world scenario data set was collected. Point clouds from an indoor laboratory environment were gathered at fifty distinct positions. The resulting ROC plots obtained are discussed in the next section.

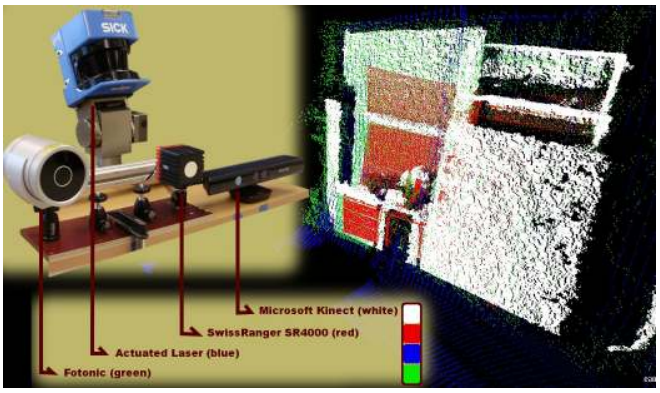

Fig. 1. Sensor setup used for the presented evaluation. An actuated SICK LMS-200, a SwissRanger SR-4000, a Fotonic B70 and a Microsoft Kinect camera were all mounted on a mobile platform. The resulting point sets were registered in a common reference system and used for evaluation.

IV. Results

A. Evaluation with Known Ground-Truth Information

Ten point clouds from each of the four sensors were collected and compared to the reference point cloud P_r, resulting in the ROC plot in Figure 3(a). Evidently, the point sets obtained by the actuated laser are spread with the largest variance of all four sensors. This behavior is unexpected, as a large spread indicates a big difference between individual point clouds. The variance can be explained by some inaccuracies in the laser sensor system — namely an inconsistency in the timing of messages originating from the laser and from the tilting actuator. The problem is further complicated by the large amount of data coming from the three high frame rate 3D sensors, resulting in high computational load of the data collection computer.

The three integrated range sensing devices all exhibit highly repeatable measurements with low variance in the scan scores. The Microsoft Kinect sensor has the best performance in this test scenario, with results very close

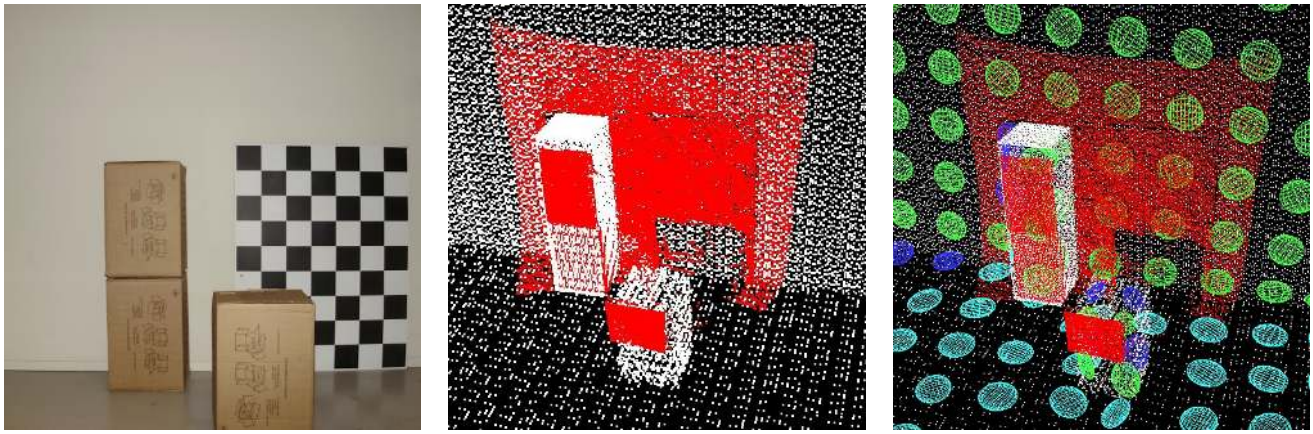

(a) (b) (c)

Fig. 2. Setup for Evaluation with Known Ground Truth. Figure 2(a) shows a picture of the controlled testing environment. Figure 2(b) shows the output from the SwissRanger SR-4000 camera (red), registered together with the ground truth point set (white.) Finally, Figure 2(c) shows the SwissRanger point set along with the 3D-NDT distributions constructed from the ground truth model. The normal distributions are drawn as ellipsoids at three standard deviations, with color coding indicating the dominant orientations of the distributions.

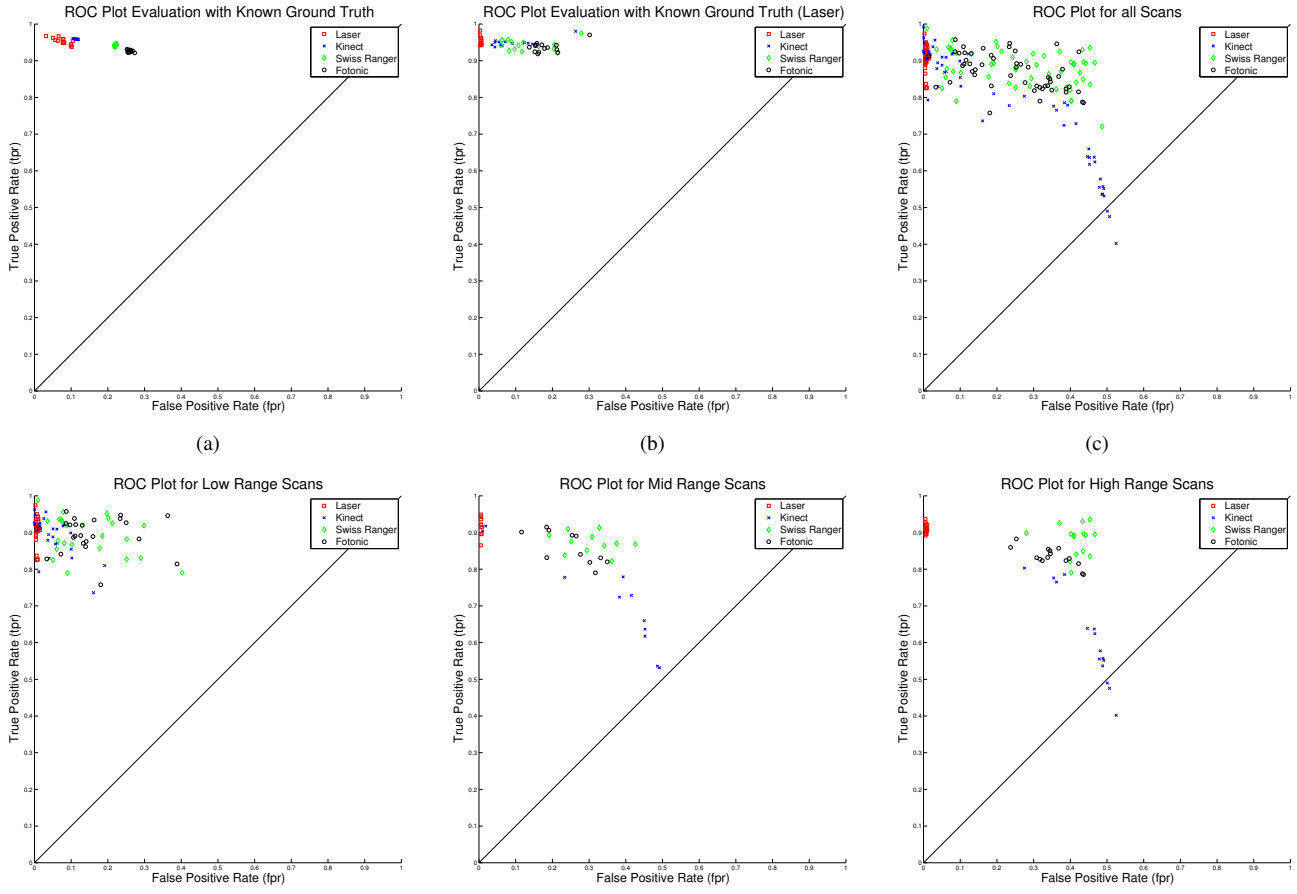

(a) (b) (c)

(d) (e) (f)

Fig. 3. Receiver Operating Characteristic (ROC) plots for all the evaluation scenarios discussed. Figure 3(a) shows the ROC points for the test scans in the ground truth scenario, for all four sensors. Figure 3(b) shows the ROC points in the same scenario, when considering the actuated laser points as ground truth data. Figure 3(c) shows the ROC plots for all sensors in the lab environment scenario. Finally, Figures 3(d)–3(f) show partitions of the lab environment test set, according to the average point distances from the sensor. The figures show point sets at $0 - 3m$, $3 - 5m$ and $> 5m$ ranges respectively.

to those of the laser sensor. The two time-of-flight cameras exhibit similar performance, with the SwissRanger slightly outperforming the Fotonic camera.

Figure 3(b) shows results from the same test scenario, with the laser sensor point clouds used as a reference sensor, providing the point cloud P_r. Comparing with Figure 3(a) can give a good sense of the expected accuracy of the results presented later. First, due to the inaccuracies in the laser system discussed earlier, measurements are considerably more spread across the ROC plot. Notably though, the comparative performance of the evaluated sensors remains the same, with the Kinect sensor producing the best results. The average values of the sensor accuracy also remain close to their ground truth values, making a strong case for the proposed comparison strategy, especially so in the absence of reliable ground truth information.

B. Evaluation in an Uncontrolled Environment

Fifty sets of test examples collected from an indoor laboratory environment were used to generate the ROC plot shown in Figure 3(c). The ROC points associated with the measurements from the actuated laser are as expected concentrated in the top-left area of the plot. The obtained points are spread over a small area, featuring a high rate of true positives and a negligible rate of false positives. As the laser points were used to construct the ground truth 3D-NDT models, this behavior is normal and only serves as a reference for the best possible results on the data set. The points from the other three sensors are spread over a large area of the ROC plot and do not seem to offer a clear distinction between the sensors. A closer look at the collected data can however offer some insights about the sensor performance.

It is important to note that in the operational mode used (millimeter precision), the SICK LMS-200 has a maximum range of eight meters. The operational ranges of the three integrated 3D sensors, as declared by the manufacturers are between 0.8 and 3.5 meters for the Kinect, up to 5 meters for the SR-4000 and between 0.1 and 7 meters for the Fotonic B70. Thus, a separate evaluation, depending on the average distances measured in the collected point sets is necessary. Figures 3(d)–3(f) show the results obtained for partitions of the evaluation set into point clouds containing ranges of up to 3, between 3 and 5 and above 5 meters respectively. The distributions of points on the ROC plots obtained are notably more concentrated and easier to interpret.

When observing objects that are closer than three meters, the Microsoft Kinect structured light camera performs the best, delivering point clouds with a high likelihood conditioned on the laser measurements. The Fotonic and SwissRanger cameras have slightly worse performance, with a slightly higher rate of false positives. Table I summarizes the average performance of each sensor on the full data set, as well as the partitioned sets. As also testified by Figure 3(d), the performance of the Kinect sensor is very close to that of the laser, for short range environments. The two time-of-flight cameras did not show significant differences in the short range test set. In the indoor environment tested, the two cameras performed in a similar fashion over the whole test set, with the SwissRanger naturally degrading on scans containing longer distances.

Data Set	True Positive Rate (TPR)	$\sigma(TPR)$	False Positive Rate (FPR)	$\sigma(FPR)$
aLRF full	0.91	0.02	0.01	0.01
Kinect full	0.77	0.15	0.22	0.20
SR-4000 full	0.87	0.05	0.26	0.14
F B700 full	0.86	0.04	0.24	0.11
Kinect low	0.89	0.05	0.05	0.05
SR-4000 low	0.89	0.05	0.14	0.10
F B700 low	0.89	0.04	0.16	0.09
Kinect mid	0.71	0.13	0.34	0.18
SR-4000 mid	0.87	0.02	0.30	0.07
F B700 mid	0.85	0.04	0.25	0.07
Kinect high	0.61	0.12	0.44	0.07
SR-4000 high	0.87	0.04	0.41	0.04
F B700 high	0.83	0.02	0.35	0.05

TABLE I

AVERAGE AND STANDARD DEVIATION OF TRUE AND FALSE POSITIVE RATES FOR DIFFERENT SENSORS AND SCAN TYPES

In the evaluation performed, no sensor achieved performance comparable to the one of the laser sensor at the full distance range. Depending on the application scenario and intended use of the sensors, additional evaluation might be needed to asses the feasibility of using an integrated 3D range device, in place of an actuated laser sensor. For example, if the primary use of the sensor is collision avoidance, distances of less than three meters from the robot are usually a sufficient look-ahead interval. Thus, any of the three evaluated 3D range sensors would perform well. If a mapping application is envisioned, the uncertainty in far-away objects position would have to be explicitly taken into account, if one of the evaluated 3D range cameras is to be used. Amplitude thresholding for the time-of-flight cameras, or a distance cutoff for the structured light sensor might be necessary to ensure better performance.

In order to demonstrate the precision of the evaluated sensors at shorter ranges, an additional test was performed. The point sets, obtained by all four sensors were filtered, eliminating all obstacle points found at ranges above 3.5 meters. Thus, when only close points are considered both for model building and testing, the ROC plots in Figure 4 are obtained. The improvement in performance is evident, with the average accuracy of the Kinect sensor increasing from $(0.77, 0.22)$ TPR/FPR ($\sigma(0.20, 0.15)$) to $(0.90, 0.06)$ TPR/FPR ($\sigma(0.03, 0.01)$). A similar performance gain can be observed for the Fotonic B70 camera. The SwissRanger SR-4000 is not influenced by the relaxed conditions, due to the severe *backfolding* effects that occur in high range scans — due to the modulation frequency used bright objects at more than five meters cannot be distinguished from faint objects at very close range. This effect was not observed in the Fotonic ToF sensor, due to the use of multiple modulation frequencies for range disambiguation.

V. DISCUSSION

This work presented a comparative evaluation of three integrated 3D range cameras. A novel evaluation methodology,

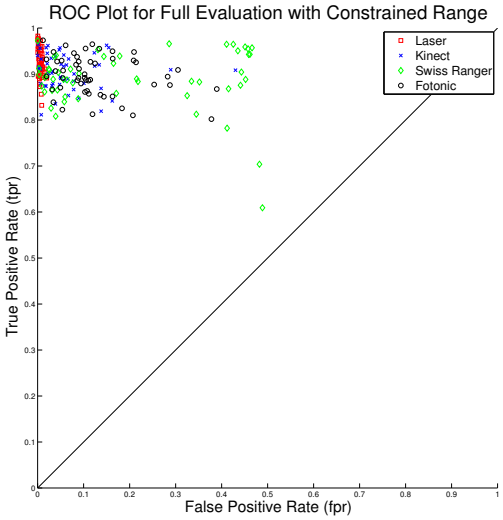

Fig. 4. ROC plot for the full indoor environment data set, considering only ranges of less than 3.5 meters. The performance of the Kinect and Fotonic cameras improves drastically.

based on previous work on spatial representation evaluation, was used to compare the outputs of the three cameras. The proposed approach offers easy to interpret statistics for the descriptive power of different sensors, compared to a set of reference range observations. In simple environments, handcrafted precise geometric models can be used to evaluate absolute sensor performance. In addition, a sensor with known performance can be used to collect reference measurements in arbitrarily complex application scenarios. A direct comparison of the ranges obtained with a different sensor can then be performed, allowing for further analysis into the usability of novel range sensors. The proposed methodology does not make any assumptions about the operational principle of the evaluated range sensors and can thus be used for generic benchmarking of newly available devices.

The accuracy of three integrated 3D range sensors — a SwissRanger SR-4000 and Fotonic B70 ToF cameras and a Microsoft Kinect structured light camera, was compared to that of an actuated laser range sensor. The three devices evaluated can deliver high frame rates and dense range measurements, but none of them features an accuracy similar to that of the laser system. A notably different conclusion can be reached when explicitly considering environments of constrained size. When all range measurements are concentrated within a sphere of radius 3.5 meters, the Microsoft Kinect sensor has an accuracy, similar to the one obtained by the actuated laser sensor. In conclusion, when considering smaller environments the Kinect and Fotonic B70 sensors can be directly used as a high data rate substitute of an aLRF system.

Future work in sensor evaluation will focus on a more in-depth analysis of the acquired data set and more extensive comparison of the expressive power of the sensors considered. One option to be explored is the direct comparison of the 3D-NDT models that can be constructed from the range measurements returned by each sensor. Other directions to be further explored are a diversification of the set of sensors evaluated and the collection of a larger test sample from different environments. Emphasis will also be placed on outdoor environments and some application specific scenarios, in particular the scenario of grasping deformable objects, stacked in a container.

ACKNOWLEDGMENTS

This work has been partially supported by the European Union FP7 Project "RobLog — Cognitive Robot for Automation of Logistic Processes".

REFERENCES

[1] T. Kahlmann, "Range imaging metrology: Investigation, calibration and development," Ph.D. dissertation, 2007.

[2] M. Lindner, I. Schiller, A. Kolb, and R. Koch, "Time-of-Flight sensor calibration for accurate range sensing," *Computer Vision and Image Understanding*, 2010.

[3] S. Fuchs and G. Hirzinger, "Extrinsic and Depth Calibration of ToF-cameras," in *Computer Vision and Pattern Recognition, IEEE Computer Society Conference on*, 2008, pp. 1–6.

[4] F. Chiabrando, R. Chiabrando, D. Piatti, and F. Rinaudo, "Sensors for 3D Imaging: Metric Evaluation and Calibration of a CCD/CMOS Time-of-Flight Camera," *Sensors*, vol. 9, no. 12, 2009.

[5] S. May, B. Werner, H. Surmann, and K. Pervolz, "3D Time-of-Flight Cameras for Mobile Robotics," in *Proc. of IEEE/RSJ Int. Conf. on Intelligent Robots and Systems. IROS 2006*, pp. 790–795.

[6] S. May, D. Dröschel, S. Fuchs, D. Holz, and A. Nuchter, "Robust 3D-mapping with Time-of-Flight Cameras," in *Proc. of IEEE/RSJ Int. Conf. on Intelligent Robots and Systems. IROS 2009*, pp. 1673–1678.

[7] K. Pathak, A. Birk, and J. Poppinga, "Sub-pixel Depth Accuracy with a Time of Flight Sensor using Multimodal Gaussian Analysis," in *Proc. of IEEE/RSJ Int. Conf. on Intelligent Robots and Systems. IROS 2008*, pp. 3519–3524.

[8] D. Dröschel, D. Holz, J. Stückler, and S. Behnke, "Using Time-of-Flight Cameras with Active Gaze Control for 3D Collision Avoidance," in *IEEE Int. Conf. on Robotics and Automation, ICRA 2010*, pp. 4035–4040.

[9] Y. Cui, S. Schuon, D. Chan, S. Thrun, and C. Theobalt, "3D Shape Scanning with a Time-of-Flight Camera," 2010, pp. 1173–1180.

[10] T. Stoyanov, M. Magnusson, H. Almqvist, and A. J. Lilienthal, "On the Accuracy of the 3D Normal Distributions Transform as a Tool for Spatial Representation," in *Proc. of IEEE Int. Conf. on Robotics and Automation, ICRA 2011*.

[11] P. Biber and W. Straßer, "The Normal Distributions Transform: A new Approach to Laser Scan Matching," in *Proc. of IEEE/RSJ Int. Conf. on Intelligent Robots and Systems. IROS 2003*, 2003.

[12] P. Biber, S. Fleck, and W. Straßer, "A Probabilistic Framework for Robust and Accurate Matching of Point Clouds," in *26th Pattern Recognition Symposium (DAGM 04)*, 2004.

[13] M. Magnusson, T. Duckett, and A. J. Lilienthal, "3D scan registration for autonomous mining vehicles," *Journal of Field Robotics*, vol. 24, no. 10, pp. 803–827, Oct. 24 2007.

[14] 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 *Proc. IEEE Int. Conf. on Robotics and Automation, ICRA 2009*, pp. 3907–3912.

[15] M. Magnusson, "The Three-Dimensional Normal-Distributions Transform — an Efficient Representation for Registration, Surface Analysis, and Loop Detection," Ph.D. dissertation, Örebro University, Dec. 2009.

[16] H. Andreasson, M. Magnusson, and A. J. Lilienthal, "Has Something Changed Here? Autonomous Difference Detection for Security Patrol Robots," in *Proc. of IEEE/RSJ Int. Conf. on Intelligent Robots and Systems. IROS 2007*, 2007.

[17] T. Stoyanov, M. Magnusson, H. Andreasson, and A. J. Lilienthal, "Path Planning in 3D Environments Using the Normal Distributions Transform," in *Proc. of IEEE/RSJ Int. Conf. on Intelligent Robots and Systems. IROS 2010*, Taipei, Taiwan, 2010.