Performance Characterization of LIDAR Based Localization for Docking a Smart Wheelchair System

C. Gao[1], I. Hoffman[2], T. Miller[1], T. Panzarella[2] and J. Spletzer[1]
[1] Lehigh University, Bethlehem, PA, USA,
spletzer@cse.lehigh.edu
[2] Freedom Sciences, LLC, Philadelphia, PA, USA,
tpanzarella@freedomsciences.com

***Abstract*— The Automated Transport and Retrieval System (ATRS) [1] represents an alternative to van conversions for automobile drivers with lower body disabilities. The core of ATRS is a "smart" wheelchair that navigates between the driver's position and a powered lift at the rear of the vehicle. From an automation perspective, autonomously docking the wheelchair onto the lift platform was the most significant technical challenge due to its geometry constraints. This work investigates the use of a LIDAR-based localization scheme to meet these requirements. Testing was conducted at the component and system level. Results indicate the proposed approach is both accurate and robust under benign and degraded visibility conditions. We expect ATRS to enter the market in late 2007.**

I. INTRODUCTION

A van conversion currently represents the sole personal transportation solution for an individual in a wheelchair. While enabling independent mobility, van conversions represent a costly and unsafe transportation solution. To eliminate these shortcomings, we have developed a technology-based alternative - ATRS. A primary benefit of ATRS is the separation of the operator and his/her wheelchair during both vehicle operations and entry/exit. This eliminates the potential for injuries or deaths caused by improper securement (as the operator is seated in a crash-tested seat system), or lift malfunctions [2]. Furthermore, by eliminating the drastic and permanent vehicle modifications associated with van conversions, ATRS will cost significantly less.

In describing the ATRS operational procedures, we refer to Figure 1. When the operator returns to his/her automobile, a keyless entry is used to both unlock the vehicle and to deploy the traversing driver's seat. The operator then positions the wheelchair, and performs a seat-to-seat transfer (pose A). After this, the wheelchair is deployed to the rear of the vehicle (pose B). In our proof-of-concept system, this side traversal was completely autonomous [3]. In the current system – referred to colloquially as "ATRS-Lite" – the wheelchair is remotely controlled by the vehicle operator via a joystick located at the user interface (UI). Once the chair enters the LIDAR's field-of-view at the rear of the vehicle (pose C), it is automatically tracked. The UI then cues the operator to place the wheelchair into "docking" mode. This enables the van-side computer to transmit control inputs in real-time to the chair over a dedicated RF link. This allows autonomous reliable docking (locking in place) onto the lift platform (pose D).

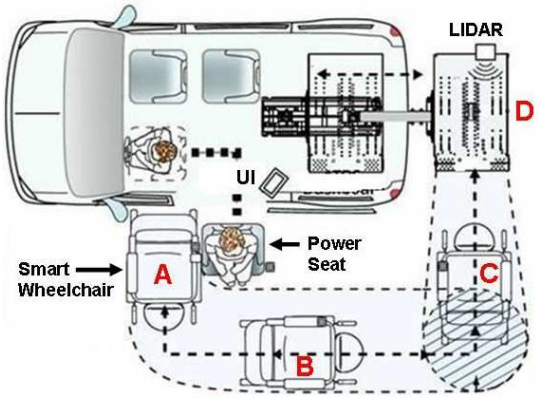

Fig. 1: ATRS concept diagram illustrating the process for autonomously stowing the smart wheelchair.

II. LOCALIZATION APPROACH

A SICK LMS291 LIDAR is the primary sensor used for estimating the wheelchair pose with respect to the lift platform. Its integration into the ATRS wheelchair lift platform is shown in Figure 2 (left). The LMS291 measures the line-of-sight range to objects in the environment over a 90° field of view, with a discretization of 0.5°. Measurements from each LIDAR scan can be written as a tuple of the form $z_m = [r, \alpha, \gamma]_m$, $m = 0...180$, where r_m and γ_m denote the measured range to and reflectivity of the m^{th} feature at a bearing of $\alpha_m = m/2 - 45°$ with respect to the LIDAR sensor frame. To simplify the feature segmentation process, two cylindrical retro-reflectors fiducials were permanently affixed to the front of the wheelchair, as shown in Figure 2 (center). The returned reflectivity signal from the retro-reflectors allowed them to easily be isolated in the LIDAR scan, and as a consequence greatly simplified the feature segmentation task. This is illustrated in the sample LIDAR scan at Figure 1 (right-bottom).

To test the stability of using a reflectivity threshold for segmenting the wheelchair fiducials, we measured the maximum range at which the retro-reflector's return signal would maximize the LIDAR's 8-bit reflectivity measurement ($\gamma = 255$). The return signal was still maximized at ranges in excess of 14 meters, easily encompassing our operational tracking envelope of 4 meters. This allowed two simple binary filters on reflectivity and range for segmenting the

Fig. 2: (Left) Wheelchair lift platform showing LIDAR integration. (Center) Smart wheelchair with retro-reflectors fixed to the front. (Right) Range (top) and reflectivity (bottom) data from a single LIDAR scan. The retro-reflector reflectivity values (two peaks) greatly simplify the feature segmentation process.

target features:

1) $\gamma_m > \gamma_{min}, \quad m = 0,1,...,180$
2) $r_m < r_{max}, \quad m = 0,1,...,180$

During all testing, $\gamma_{min} = 250$, $r_{max} = 400$ cm (in hindsight, γ_{min} could have been set to 255 based upon the test results for even stronger filtering). After this filtering phase, valid features were clustered in Euclidean space as "candidate fiducials" based upon the known fiducial size ($d = 5.2$ cm). Every hypothetical fiducial pair was then evaluated against the known fiducial baseline (44 cm). Any pair within ± 2 cm tolerance of the known baseline was considered a valid chair pose. If and only if 1 valid chair pose was obtained, the wheelchair segmentation was considered to be successful. Otherwise, the user would be notified via the UI to take corrective action (*i.e.*, reposition the chair) and the segmentation process repeated. This approach has proven to be extremely reliable in real-world conditions. With the position of both fiducials known, estimating the chair position and orientation (assuming a ground plane constraint) was straight-forward.

Fig. 3: Test fixture used in fiducial testing. The turntable simulated wheelchair motion, while allowing the estimated wheelchair positions to be projected to a single point.

III. COMPONENT LEVEL TESTING

To characterize localization system performance, a modified turntable (shown in Figure 3) was used. This allowed simulated wheelchair position estimates to be projected to a single point (the center of rotation) so that "ground truth" localization accuracy could be well characterized. Also, the rotation closely modeled the maximum expected wheelchair fiducial velocities (0.4 m/s and 0.9 rad/s). For these tests, the fiducials were mounted on an arm and spaced 44 cm apart (the same distance as on the wheelchair used in development).

A. Baseline Localization Accuracy

The turntable used for testing was initially placed at a nominal position of (x,y)=(1.0m,0.0m) with respect to the LIDAR coordinate frame. The turntable was then actuated, and the fiducials tracked by the SICK for a minimum of 5,000 scans. The x-distance to the turntable was then increased by 0.5 meters until 4 meters, and the process repeated.

A sample plot of the estimated x-y coordinates is shown at Figure 4 (left), while summary data for all positions are at Figure 4 (right). Assuming an unbiased estimator, *i.e.*, the mean position of the distribution is the actual center of rotation, the mean absolute error at all ranges was less than 7.5 mm; at ranges less than 2.5 m, it was < 4 mm. This is an order of magnitude smaller than the docking tolerance ($\approx 4-5$ cm) which highlights the efficacy of the localization approach.

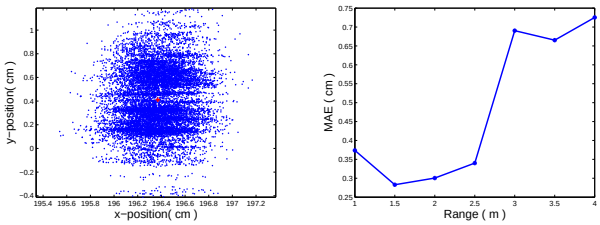

Fig. 4: (Left) Position estimates of the turntable at 2 m. Over 95% dots distributed within 1.2×1.2 cm^2 area. The mean absolute error at all ranges (Right) was less than 7.5 mm.

B. Degraded Fiducial Testing

The purpose of these subtests was to simulate the effects of wear and tear on the retro-reflector fiducials and estimate its impact on localization performance. During these subtests, the left retro-reflector fiducial was covered on the front side side by a 2 cm black stripe. The rear sides of both fiducials were left unchanged. This test setup is illustrated in Figure 5. As a result, when the front side of the turntable arm assembly was visible to the LIDAR, we collected data for

the "damaged" configuration. The rear side corresponded to our undamaged baseline. This allowed a simultaneous comparison without any other changes in the test configuration.

Fig. 5: (Left) Simulated damaged fiducials on the arm's front side. (Right) Undamaged fiducials on the rear side. This allowed a simultaneous comparison of damaged fiducial position accuracy vs. an undamaged baseline system.

Over $10,000$ scans were made to ensure a minimum set size of $5,000$ samples for both the damaged and undamaged retro-reflector configurations. These trials were conducted for $x = \{1.0, 2.0, 3.0, 4.0\}$ meters, $y = 0.0$ m. A sample plot for nominal x-y coordinates of $(1.0\text{m}, 0.0\text{m})$ is shown at Figure 6 (left). As expected, the variance of the position estimates from the damaged fiducials was greater than that of the baseline system. Summary data for all trials are shown at Figure 6 (right). For even the damaged fiducial case, mean absolute errors were < 7 mm for all trials. We also observed a typical 3 mm y-position bias between the mean position of the two distribution sets. This can also be seen in Figure 6 (left). Again, this also was not unexpected, as only one fiducial was partially covered. Even with this bias, the y-position error was still less than 1 cm, which again is far below our docking tolerances.

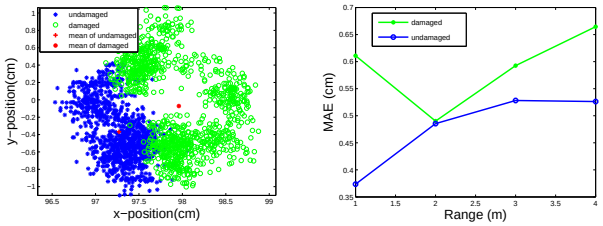

Fig. 6: (Left) Position estimation of turntable for damaged/undamaged fiducials at 1 meter. The variance from damaged fiducials (lighter, in the right of the plot) is greater than the variance from undamaged fiducials. There is also ≈ 3 mm y-position bias. (Right) Mean absolute error comparison between the damaged and undamaged fiducials in all trials. The mean absolute error was still < 7 mm in both sets.

C. Rain Testing

Testing was performed to determine if rainfall would influence the LIDAR measurement data and compromise our localization algorithm. In order to protect the turntable assembly electronics, these tests were conducted using two static fiducials. As the LMS291 integrates rain protection in the near-field (ranges < 2 meters), rain testing focused on standoff distances of $3-4$ meters from the LIDAR. Trials were conducted under five different conditions: no-rain (base-line), light, medium, heavy, and very heavy rain (*i.e.*, a thunderstorm). The positions of the fiducials and the LIDAR were fixed throughout the entire testing process to ensure fair comparison of the localization accuracy in different environments.

Figure 7 (left) summarizes the mean absolute error for the four rain scenarios at 3 meters distance. We see that light to medium rain had little to no effect on the feature segmentation and localization accuracy; mean absolute error was less than 5 mm. However, heavy rain increased the "random" error to ≈ 2 cm. It also introduced outliers in the range measurements, where a significantly shorter than actual range was obtained. Without proper filtering, these could introduce a dramatic error to the wheelchair pose estimate, as illustrated in Figure 7 (right). Fortunately, the outlier rate for all trials was very small ($\approx 1-2\%$), and can easily be handled using traditional techniques (in our case, median filtering on the range measurements). The higher random error also motivates the need for temporal filtering, and we have now integrated an Extended Kalman Filter to estimate the wheelchair pose for this purpose [1].

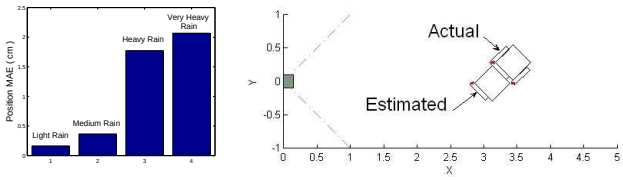

Fig. 7: (Left) Mean absolute error for light to very heavy rain rates. (Right) Severe wheelchair pose estimate error from a range measurement outlier caused by heavy rain.

IV. CONCLUSIONS

In this paper, we have outlined our LIDAR based localization approach for docking a smart wheelchair system. We have also outlined the test procedures for characterizing the subsystem performance. During the past year, this system has been integrated into the beta ATRS system and tested across a range of conditions. This included three days of continuous demonstrations at the World Congress Exposition on Disabilities (WCD 2006) in November 2006. Conference participants were also given the opportunity to test the system. Over 300 cycles of docking and undocking were conducted during this time without a single failure - testifying to the efficacy of the approach. We expect ATRS to enter the commercial market in late 2007.

REFERENCES

[1] C. Gao, I. Hoffman, T. Panzarella, and J. Spletzer, "Automated Transport and Retrieval System (ATRS): A Technology Solution to Automobility for Wheelchair Users", *in 6th Conference on Field and Service Robotics (FSR '07)*, Chamonix, France, July 2007.

[2] U.S. DOT NHTSA, "Research Note: Wheelchair User Injuries and Deaths Associated with Motor Vehicle Related Incidents", Sep 1997.

[3] H. Sermeno-Villalta and J. Spletzer, "Vision-based Control of a Smart Wheelchair for the Automated Transport and Retrieval System", *in Proceedings of the 2006 IEEE International Conference on Robotics and Automation (ICRA '06)*, Orlando, USA, May 2006.