

Towards Autonomous Wheelchair Systems in Urban Environments

Chao Gao, Michael Sands, and John R. Spletzer

Abstract In this paper, we explore the use of synthesized landmark maps for absolute localization of a smart wheelchair system outdoors. In this paradigm, three-dimensional map data are acquired by an automobile equipped with high precision inertial/GPS systems, in conjunction with light detection and ranging (LIDAR) systems, whose range measurements are subsequently registered to a global coordinate frame. The resulting map data are then synthesized *a priori* to identify robust, salient features for use as landmarks in localization. By leveraging such maps with landmark meta-data, robots possessing far lower cost sensor suites gain many of the benefits obtained from the higher fidelity sensors, but without the cost. We show that by using such a map-based localization approach, a smart wheelchair system outfitted only with a 2-D LIDAR and encoders was able to maintain accurate, global pose estimates outdoors over almost 1 km paths.

1 Introduction

We are interested in developing smart wheelchair systems capable of autonomous navigation in unstructured, outdoor environments. Our primary work to date in this area has been with the Automated Transport and Retrieval System (ATRS) [1]. ATRS enables independent mobility for drivers in wheelchairs by automating the stowing and retrieval of the driver's wheelchair system. While ATRS has been commercialized, and its smart-chair system does in fact navigate autonomously, its autonomy is limited to an area immediately adjacent to the host vehicle. We would like to build on these results to support a greater range of smart-chair applications. Key to this objective is a robust means for outdoor localization.

Localization is a fundamental enabling technology for mobile robotics, and as a result a very active research area. Although the problem in structured, indoor environments might be considered solved, robust localization outdoors is still an open research problem. While the community has made significant strides recently in terms of vehicle autonomy outdoors [2], much of this has been achieved through

C. Gao, M. Sands, and J. Spletzer
Lehigh University, Bethlehem, PA, USA, e-mail: chg205, mjs207, josa@lehigh.edu

sensor suites using tightly coupled inertial/GPS navigation systems costing up to \$100K or more. Such a solution is impractical in terms of both size and cost for applications such as ours. In the absence of reliable GPS measurements, fall-back strategies are similar to those used indoors and involve extracting strong features from RADAR, LIDAR, vision sensors, *etc.*, and tracking their relative position and uncertainty estimates over time [3, 4, 5]. However, such approaches are more fragile when used outdoors due to the absence of continuous structure and the much larger problem scale.

Urban environments represent an interesting and important middle-ground as over 80% of the U.S. population resides in cities and suburbs [6]. The availability of GPS measurements in these areas for pose estimation can typically be assumed, but multi-path errors from buildings, trees, *etc.*, can significantly compromise its accuracy. Fortunately, these same structures can be used as landmark features to yield accurate relative position estimates. In this paper, we investigate a paradigm where a smart wheelchair system relying upon lower cost sensors localizes with the assistance of three-dimensional maps generated by a vehicle equipped with a high fidelity sensor suite. These maps are synthesized *a priori* to identify robust, salient features which can be used as landmarks for robot localization. By leveraging these maps, the wheelchair gains many benefits obtained with the higher-fidelity sensors but without the cost.

2 Related Work

Our work relates to research efforts in three-dimensional mapping as well as robot localization and mapping. Generating three-dimensional maps of urban areas has been investigated by several groups, so there is significant previous work that can be leveraged [7, 8, 9, 10, 11, 12]. Unlike most of this work, our focus is generating and processing three-dimensional maps with respect to a global frame, which will be reusable and readily extended by any user. With the explosion of data services, we expect the availability of such maps to be commonplace in the future [13]. Our motivation is that by leveraging such maps, lower fidelity sensors could be employed. This has been demonstrated routinely indoors (*e.g.*, MCL with sonar vs. SLAM with LIDAR), and we believe the analogy will hold outdoors. Many features in urban environments are viable candidates for landmarks. For example, corner features are often used in EKF localization and mapping approaches as their position can be reduced to a single point [14]. Building facades and walls might also be used [8]. Indeed even signage can be detected and recognized [15]. While we are ultimately interested in integrating aspects of each of these features within our synthesized map representation, in this paper we limit our focus to pole features as landmarks. In this context, pole features would correspond to street lamps, trees, parking meters, street signs, *etc.* Such features are prevalent in most urban areas.

The use of poles features as landmarks has been investigated by other researchers. This includes the work of [16, 17], among others. The primary focus of these efforts was SLAM with a ground vehicle (*i.e.*, an automobile) where “cylinder” features were segmented using vision and/or LIDAR systems, and tracked over time. This

technique enabled mobile localization and mapping in outdoor, unstructured environments over relatively long distances (*e.g.*, 100s of meters). We propose to build upon these efforts by first building large-scale three-dimensional maps, synthesizing these maps to identify strong landmark features, introducing a refinement stage to improve map consistency, and then leveraging these maps with an ultimate goal of improving localization performance outdoors.

3 Map Generation

Data Acquisition. Our vehicle for data acquisition was “Little Ben,” which previously had served as the Ben Franklin Racing Team’s entry in the DARPA Urban Challenge [18]. Vehicle pose was provided by an Oxford Technical Solutions RT-3050, which uses a Kalman filter based algorithm to fuse inertial measurements, GPS updates with differential corrections, and odometry information from the host vehicle. It provides 6-DoF pose updates at 100 Hz with a stated accuracy of 0.5 meters circular error probable (CEP). Range and bearing measurements from a pair of roof mounted, vertically scanning Sick LMS291-S14 LIDAR systems were then registered to the current vehicle pose to generate high-resolution range maps. The two LIDARs are highlighted (circled red) in Fig. 1 (Left). We used two LIDARs to improve map reconstruction by reducing scene occlusion and for redundant measurements to reduce noise effects. During data acquisition, Ben was driven at 8-12 km/hr. This corresponded to a LIDAR scan spacing of $\approx 4\text{-}6$ cm, which allowed even thin pole features (*e.g.*, street signs, parking meters) to be captured reliably.



Fig. 1 Development Platforms. (Left) Vehicle used for data acquisition. Ben integrates an OXTS RT-3050 and a pair of vertically scanning Sick LMS291-S14 LIDARs (circled red). (Center-Right) Our smart-chair platform integrates LIDAR, vision, GPS, and odometry sensors. Its computer interface and on-board power distribution enable a range of sensors and accessories to be quickly integrated for prototyping purposes.

LIDAR Calibration. Ultimately, we need to register the acquired range scans to a common world frame W . This registration requires knowledge of the extrinsic parameters (rotation R and translation T) of both the vehicle frame V , and the front and back LIDAR frames (F, B) versus time with respect to W . The vehicle parameters $R_V^W(t), T_V^W(t)$ are estimated directly by the OXTS RT-3050 at 100 Hz. As the LIDARs are related to the vehicle frame by a rigid transformation, we need only recover the extrinsic parameters of the LIDARs with respect to the vehicle frame. For this work, we developed a novel approach to simplify the calibration process. Noting that points in the world frame and LIDAR frames were related by rigid transforma-

tions, we could recover the LIDAR extrinsic parameters with a sufficient number of point correspondences between the front and back LIDARs. To facilitate correspondence tracking over time, 6 poles on a $\approx 50 \times 100$ meter calibration loop were used as landmarks. On each pole, two retro-reflective targets were placed 1 meter apart for a total of 12 target points $X_T = [x_1, y_1, z_1, \dots, x_{12}, y_{12}, z_{12}]^T$ in W . These targets were automatically segmented from the environment by thresholding the LIDAR remission measurements. A sample landmark pole is shown in Fig. 2 (left).

Point correspondences were obtained by driving multiple cycles around our calibration loop. Our calibration process then consisted of two stages. The first was to remove deterministic error between successive calibration loops caused by GPS jumps. This step was accomplished by using the first loop to generate reference landmark positions $X_1 = [x_{1,1}, y_{1,1}, z_{1,1}, \dots, x_{12,1}, y_{12,1}, z_{12,1}]^T$, where $x_{i,j}$ denotes the estimated x-position of the i^{th} target in W during the j^{th} calibration loop. The deterministic shift $S_j = [s_{xj}, s_{yj}, s_{zj}]^T$ of the j^{th} loop was estimated by

$$S_j^* = \underset{S}{\operatorname{argmin}} \|X_1 - X_j - S_{j \times 12}\|^2 \quad (1)$$

where $S_{j \times 12} \in \mathbb{R}^{36}$ is the estimated deterministic error S_j replicated for each target point. This minimization problem was solved using a least-squares approach. S_j was then treated as a bias, and the value of X_j for each loop was adjusted accordingly.

The second stage was to remove the influences of random error in the calibration process. In doing so, we needed to estimate the extrinsic parameters for each LIDAR ($R_B^V, T_B^V, R_F^V, T_F^V$), as well as the positions of our 12 targets (X_T) in W . We note that for the same world point \mathbf{x}_W ,

$$R_{VF}^W(R_F^V \mathbf{x}_F + T_F^V) + T_{VF}^W = R_{VB}^W(R_B^V \mathbf{x}_B + T_B^V) + T_{VB}^W \quad (2)$$

where the VF and VB subscripts are used to denote the vehicle transformation to the world frame corresponding to the different vehicle poses when \mathbf{x}_W was observed by the front and back LIDARs, respectively. Thus, we can solve for both the LIDAR extrinsics as well as the target positions with a minimum of 16 point correspondences between the front and back LIDARs. Since the vehicle pose varies with each calibration loop, 12 unique correspondences can be obtained from each loop cycle. As a result, a large number of correspondences can be acquired very quickly. We solved for (2) using a non-linear minimization solver. However, one final enhancement was added first to remove measurement outliers. For this, we employed a ‘‘constrained’’ RANSAC approach [19], where we instantiated each model hypothesis with a small number (2-4) of correspondences at random from *each target*. This enhancement ensured that the error residuals were balanced across the entire calibration loop.

Fig. 2 shows representative results of merging front and back LIDAR data using the measured extrinsics (center) and those estimated from the calibration process (right). Improvements in data fusion and scene reconstruction from the calibration process are clearly visible. The mean absolute error (MAE) for the error residuals between the two LIDAR reprojections was 12.67 cm. As the performance of the pose system is 50 cm CEP, these results were considered satisfactory.




Fig. 2 (Left) One of six landmark poles used during the LIDAR calibration process. The retro-reflective targets could be automatically segmented to track correspondences across multiple calibration loops. Reconstruction results before and after our calibration phase. Improvements in data fusion with the front (red) and back (blue) LIDARs from the calibration phase are clearly visible.

Landmark Synthesis. Segmenting pole features was accomplished using a two-step clustering approach: (1) recursively cluster points within each scan, and (2) merge clusters in successive scans where appropriate. In both steps, a Euclidean distance threshold was used as the clustering criterion. For a cluster to be accepted as a pole feature, validation gates were placed on cluster size. Furthermore, only strongly vertical clusters were accepted by examining the covariance matrix C of the associated feature points' positions. Specifically, the eigenvector associated with the largest eigenvalue λ_{max} of C should be close to $[0, 0, 1]^T$. Only after clearing these validation gates was the cluster accepted as a landmark in the synthesized map.

4 Wheelchair Localization

The smart-chair used in this work is based upon an Invacare M91 Pronto power wheelchair with Mk5 electronics. It integrates high resolution optical encoders, a Hokuyo UTM-30LX LIDAR system, a 1024x768 Point Grey digital video camera, and a Garmin 18 WAAS enabled GPS system. For this work, the UTM-30LX was the wheelchair's sole exteroceptive sensor. When compared to the ubiquitous Sick LMS2xx LIDARs, it is extremely compact. In our current integration, the LIDAR and camera system are mounted on the opposite arm as the manual joystick controller as shown in Fig. 1 (center). The configuration is comparable in size to the joystick controller box.

Landmark Detection. In our localization paradigm, the wheelchair employs LIDAR and odometry sensors in conjunction with the synthesized landmark map. Implicit in this approach is the assumption that the landmarks can be reliably segmented. However, unlike the landmark synthesis phase, the wheelchair LIDAR must rely entirely upon two-dimensional scan data. To compensate for this, our landmark detection strategy used two approaches dependent upon pole feature geometry. The first step in either approach was clustering registered point returns from the wheelchair LIDAR scan in Euclidean space. Cluster diameters were then used to discriminate between larger diameter pole features (trees, telephone poles, street lamps, *etc.*) and narrower ones (parking meters, traffic sign posts, *etc.*).

Larger diameter clusters with 5 or more points were fit as circle features. Several additional validation gates followed based upon circle geometry and residual fitting

error before a feature was accepted as a landmark candidate. We found empirically that larger diameter landmarks could be reliably detected at ranges $\approx 75 \times$ the feature radius, meaning that a landmark with a radius of 10 cm would typically be detected at a range of about 7.5 meters. Circle fitting was not appropriate for smaller diameter clusters. As such, these features were tracked over time. If they were persistent and no other clusters were detected within a given distance threshold, they were accepted as landmark candidates. Using such an approach, smaller diameter features could reliably be detected at ranges ≤ 5 meters. Candidate landmarks were passed to the data association module for additional processing.

Data Association. There are inherent limitations in using two-dimensional LIDAR measurements to segment three-dimensional landmarks. As a result, the landmark detection process may erroneously validate other environmental features as pole features. The impact of these false positives on localization performance was mitigated through a data association phase.

Several sources of uncertainty exist in the synthesized landmark locations within our map. These sources include uncertainty introduced by the pose system, errors in LIDAR extrinsic calibration, LIDAR range errors, and noise associated with the landmark synthesis process itself to name but a few. Uncertainty in landmark position was modeled by associating a covariance matrix Σ_l with the position of each landmark while Σ_w and Σ_o denoted covariance matrices associated with the uncertainty in wheelchair pose and LIDAR range and bearing observations, respectively. With these so defined, we used the Mahalanobis distance D between the predicted and observed sensor measurements as our quality metric for data association, defined as

$$D = \sqrt{\mathbf{z}^T (\Sigma_o + H_w \Sigma_w H_w^T + H_l \Sigma_l H_l^T)^{-1} \mathbf{z}} \quad (3)$$

where, H_w and H_l are the Jacobians of the observation model with respect to wheelchair pose and landmark location, respectively,

$$H_w = \begin{bmatrix} \frac{x_w - x_l}{z_r} & \frac{y_w - y_l}{z_r} & 0 \\ \frac{y_l - y_w}{z_r^2} & \frac{x_w - x_l}{z_r^2} & -1 \end{bmatrix}, \quad H_l = \begin{bmatrix} -\frac{x_w - x_l}{z_r} & -\frac{y_w - y_l}{z_r} \\ -\frac{y_l - y_w}{z_r^2} & -\frac{x_w - x_l}{z_r^2} \end{bmatrix} \quad (4)$$

and $\mathbf{z} = \mathbf{z}_l - \mathbf{z}_o$ is the difference between the predicted and actual range and bearing measurements for the wheelchair LIDAR. A threshold on D served to filter out potential false positives observed during the landmark detection phase. For the case of closely located landmarks where multiple possible detection-landmark associations might be possible, the association minimizing the total assignment cost was used.

Localization Approach. Extended Kalman Filters (EKF) have been one of the most popular techniques for state estimation in mobile robotics [3, 20, 14], and we took a similar approach for estimating the wheelchair pose $\mathbf{x} = [x_w, y_w, \theta_w]^T$. In the prediction step, linear and angular velocities (v, ω) were estimated from the encoders using a differential drive model for the wheelchair. For the correction step, the observation functions were based upon LIDAR estimates for the range z_r and bearing z_α to the segmented landmark at position (x_l, y_l)

$$z_r = \sqrt{(x_w - x_l)^2 + (y_w - y_l)^2}, \quad z_\alpha = \arctan\left(\frac{y_w - y_l}{x_w - x_l}\right) - \theta_w \quad (5)$$

and the Kalman gain was calculated as $K = PH_w^T(H_wPH_w^T + H_l\Sigma_lH_l^T + \Sigma_o)^{-1}$ where H_w and H_l were as defined in (4). The process then followed a traditional EKF implementation with updates of 2-5 Hz dependent upon vehicle velocity.

Landmark Position Refinement. A shortcoming with relying heavily upon GPS for map generation is that changes in satellite geometry/visibility can lead to “jumps” in vehicle pose. These discontinuities affect map consistency. One approach to address this would be to integrate additional sensing onto the data acquisition platform and run SLAM in parallel with the data acquisition phase [12]. We propose an alternate refinement stage where SLAM is actually run by the map client – in our case the wheelchair – during an initial route traversal akin to a learning phase. This is something we envision would be done by the wheelchair user’s care provider prior to enabling completely autonomous operations. An advantage of this approach is that the landmark refinement would be tuned to the actual sensor geometries employed by the client vehicle. For our implementation, we extended our EKF localization using a SLAM approach as in [3, 20] to further refine the landmark positions. The landmark locations were then updated with the SLAM-refined landmark positions and covariance Σ_L estimates. While this did not improve the global map accuracy, it significantly improved the map consistency.

5 Experimental Results

To investigate the viability of the proposed approach, our first experiments were conducted in the parking lots around Lehigh’s Stabler Arena. Admittedly, this area was *not* representative of urban environments. However, it served as a low-traffic proving ground with sufficient pole features to first validate the concept. Fig. 3 (left) shows the raw registered range data acquired by driving Ben through the area. These data were then synthesized as outlined above, and the resulting map with embedded landmarks is shown at Fig. 3 (right). Validating the fidelity of this reconstruction is difficult due to the lack of absolute ground truth. However, we measured the distance between 25 pairs of landmarks using a Bosch DLE50 laser distance measure and compared these to the distances of corresponding synthesized landmark pairs. The mean absolute difference between the sets was 7.2 cm. We also reviewed the reliability of the landmark synthesis approach. All 71 pole features present in the area surveyed were positively detected and synthesized into the map.




Fig. 3 Registered raw (left) and synthesized map data (right). The relative distance differences between synthesized landmark pairs and their real-world counterparts was about 7 cm.

Using the original synthesized map and landmark positions, the wheelchair was manually driven over a route network 960 meters in length at a fast walking speed (1.6 m/s). This first loop constituted the landmark refinement stage discussed in Section 4, and the wheelchair localized using an EKF SLAM approach with the segmented landmark positions. Using SLAM, the wheelchair was able to accurately maintain its pose over the entire 960 meter loop. We then repeated this same experiment 3 separate times using map-based localization with the updated landmark position and uncertainty estimates. All other parameters for data association and localization remained fixed. Representative results are at Fig. 4 (left). The landmark positions are denoted by red circles. The wheelchair path as estimated by the map-based localization approach is denoted by the blue line. The path as estimated by the wheelchair’s own GPS is also shown for comparison purposes (green line). The initial pose estimate of the GPS was also used to initialize the pose for map localization. Using the SLAM-refined landmark positions, all 3 trials were successfully completed. To characterize the localization accuracy, the wheelchair was driven over 6 ground-truth points (shown as “+”) whose positions relative to landmarks were measured by hand. The average position errors was 20 cm, with $3 < 1\sigma$, $5 < 2\sigma$, and all $6 < 3\sigma$ based upon the covariance estimates for Σ_W and Σ_L .




Fig. 4 Localization results using refined (left) and original landmark position estimates (right). Improving the consistency of landmark positions dramatically improved localization performance.

To motivate the need for the landmark refinement phase, we also ran these same trials using map-based localization (not SLAM) with the original landmark positions. Each of these trials ended in failure. This typically occurred at a portion of the course where the inter-landmark spacing required the wheelchair to rely upon dead reckoning for over 20 meters of travel. Significant error and uncertainty in wheelchair pose accumulated during this time, resulting in an incorrect feature association. This is shown in Fig. 4 (right). However, the open-loop travel was not the sole culprit. From a subsequent analysis, we determined that a fairly significant GPS shift occurred during the data acquisition phase for building the map. As a result, a fraction of the map exhibited a shift > 1 meter with respect to the map’s initial coordinate frame. This shift significantly contributed to the data association failure, and the correct robot pose could not be recovered by the EKF. As the SLAM algorithm updated the landmark positions on the fly, it was robust to this shift error. The subsequent landmark refinement stage mitigates the impact of GPS jump.

Our final experiment involved a similar test in South Bethlehem, PA adjacent to Lehigh’s campus. This was a representative urban environment, with a significantly higher density of landmarks than seen during the Stabler testing. During this test, the wheelchair was manually driven approximately 720 meters along the sidewalk at a velocity of ≈ 1 m/s. Results from this trial using SLAM are shown at Fig. 5, where the landmarks positions (red circles), the SLAM estimated path (blue line) and GPS path (green line) are superimposed on a satellite image. While completing the loop was prevented by ongoing building construction, the end position was consistent with localization estimates. Again, results compared favorably to the wheelchair’s WAAS-corrected GPS estimate. We have not been able to repeat this trial using localization with the SLAM-enhanced map due to seasonal weather conditions, but expect to in the near future.




Fig. 5 Map-based localization of the wheelchair (blue line) vs. GPS position estimates (green line).

6 Discussion

In this paper, we investigated the acquisition, synthesis and application of three-dimensional maps by a smart wheelchair for map-based localization. Since the maps were registered to a global frame, they provide a means for absolute position estimation in urban areas even in the absence of GPS. In our experiments, our wheelchair system was able to maintain accurate pose estimates after traveling hundreds of meters using such an approach. While we are satisfied with the results to date, we do realize this is just a first step. Pole features were an obvious first choice for landmarks, and we are now beginning to synthesize additional features into the map (*e.g.*, building corners). We are also interested in vision based signage detection, as these can provide nearly-unique IDs for inferring global position. We also assume the ability to automatically segment pedestrians from the environment. Our current implementation fuses results from vision and LIDAR systems. The camera uses the Haar-like feature based classifier for face detection from OpenCV [21], while the LIDAR segments candidate clusters based upon geometry constraints. Individually, both systems have high rates of false positives. However, this can be reduced dra-

matically by only accepting tracks when both sensors report a detection. A downside is that significant false-negatives remain. We are continuing to refine this approach.

References

1. C. Gao, I. Hoffman, T. Miller, T. Panzarella, and J. Spletzer, "Autonomous Docking of a Smart Wheelchair for the Automated Transport and Retrieval System (ATRS)," *Journal of Field Robotics*, vol. 25, no. 4-5, pp. 203–222, 2008.
2. M. Buehler, K. Iagnemma, and S. Singh, "Special issues on the 2007 darpa urban challenge," *The Journal of Field Robotics*, vol. 25, no. 8-10, Aug-Oct 2008.
3. M. Dissanayake, P. Newman, S. Clark, H. Durrant-Whyte, and M. Csorba, "A solution to the simultaneous localization and map building (SLAM) problem," *IEEE Trans. on Robotics and Automation*, vol. 17, no. 3, pp. 229–241, June 2001.
4. P. Newman and J. Leonard, "Pure range-only subsea slam," in *IEEE International Conference on Robotics and Automation*, Taiwan, China, Sep 2003.
5. F. Masson, J. Nieto, J. Guivant, and E. Nebot, "Robust autonomous navigation and world representation in outdoor environments," *Mobile Robots Perception and Navigation*, pp. 299–320, 2007.
6. United Nations, "Population Division: World Urbanization Prospects," Feb 2008.
7. H. Zhao and R. Shibasaki, "Reconstructing textured cad model of urban environment using vehicleborne laser range scanners and line cameras," in *Second International Workshop on Computer Vision System (ICVS)*, Vancouver, Canada, 2001, p. 284295.
8. C. Frueh and A. Zakhor, "Reconstructing 3d city models by merging ground-based and air-borne views," *Lecture Notes in Computer Science*, vol. 2849, pp. 306–313, 2003.
9. A. Georgiev and P. Allen, "Localization Methods for a Mobile Robot in Urban Environments," *IEEE Trans. on Robotics*, vol. 20, no. 5, pp. 851–864, Oct 2004.
10. A. Howard, D. F. Wolf, and G. S. Sukhatme, "Towards 3d mapping in large urban environments," in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, Sendai, Japan, September 2004, pp. 419–424.
11. B. Jensen, J. Weingarten, S. Kolski, and R. Siegwart, "Laser range imaging using mobile robots: From pose estimation to 3d-models," in *1st Range Imaging Research Day*, Zürich, Switzerland, September 2005, p. 129144.
12. D. Borrmann, J. Elseberg, K. Lingemann, A. Nüchter, and J. Hertzberg, "Globally consistent 3d mapping with scan matching," *Journal of Robotics and Autonomous Systems (JRAS)*, vol. 56, no. 2, pp. 130–142, February 2008.
13. Educating Silicon, "Google Street View - Soon in 3D?," <http://www.educatingsilicon.com/>, Apr 2008.
14. M. Altermatt, A. Martinelli, N. Tomatis, and R. Siegwart, "SLAM with Corner Features Based on a Relative Map," in *IEEE/RSJ IROS*, October 2004.
15. T. Kingston and C. Laflamme, "Automated road sign detection and recognition," *Journal of the International Municipal Signal Association*, pp. 46–49, Jan/Feb 2007.
16. F. Ramos, J. Nieto, and H. Durrant-Whyte, "Recognising and Modelling Landmarks to Close Loops in Outdoor SLAM," in *IEEE ICRA*, Mar 2007.
17. J. Nieto, T. Bailey, and E. Nebot, "Recursive scan-matching SLAM," *Robotics and Autonomous Systems*, vol. 55, pp. 39–49, 2007.
18. J. Bohren, J. Derenick, T. Foote, J. Keller, A. Kushleyev, D. Lee, B. Satterfield, J. Spletzer, A. Stewart, and P. Vernaza, "Little Ben: The Ben Franklin Racing Team's Entry in the 2007 DARPA Urban Challenge," *Journal of Field Robotics*, 2008, accepted for publication.
19. M. Fischler and R. Bolles, "Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography," in *Communications of the ACM*, 1981.
20. S. Thrun, "Robotic mapping: A survey," Tech. Rep. CMU-CS-02-111, Carnegie Mellon University, 2002.
21. Intel Corporation, *OpenCV*, <http://sourceforge.net/projects/opencv/>.