

# Towards 3D Mapping in Large Urban Environments

Andrew Howard, Denis F. Wolf and Gaurav S. Sukhatme  
Robotics Research Lab., Computer Science Department  
University of Southern California  
Los Angeles, U.S.A.  
Email: ahoward@usc.edu, denis@usc.edu, gaurav@usc.edu

**Abstract**—This paper describes work-in-progress aimed at generating dense 3D maps of urban environments using laser range data acquired from a moving platform. These maps display both fine-scale detail (resolving features only a few centimeters across) and large-scale consistency (typical maps are approximately 0.5 km on a side). In this paper, we sketch a basic 3D mapping algorithm (paying particular attention to practical engineering details) and present preliminary results acquired on the USC University Park Campus using a Segway RMP vehicle.

## I. INTRODUCTION

Laser-based mapping in indoor environments is a well studied problem for which a number practical solutions have now been demonstrated [1]. In this paper, we seek to adapt and apply some of these solutions to the less-well-studied problem of laser-based mapping in urban environments. Our aim is to generate dense 3D maps that capture features as small as a few centimeters across, over environments that are of the order of one square kilometer in area.

Compared with indoor mapping, the problem of mapping in urban environments has some distinct characteristics. First and foremost, coarse localization (with an uncertainty of few meters) is often available from GPS; thus, the data-association problems that bedevil indoor mapping algorithms (such as determining whether or not a robot has returned to a previously visited location) can be solved directly. On the other hand, while robots operating indoors can expect to maintain a constant attitude, robots operating in urban environments will necessarily experience some pitching and rolling. The urban mapping problem must, therefore, be treated as a three-dimensional problem from the outset.

In the work that follows, we make two key assumptions: (1) the robot's altitude is constant, and (2) the environment is at least partially structured (i.e., contains built objects). Note that we *do not* assume a completely static environment: urban environments are difficult to control, and inevitably contain moving objects such as pedestrians and vehicles. The basic mapping algorithm has four steps. (1) Fine-scale localization: odometry is combined with IMU and laser range-finder data to produce an incremental pose estimate for the robot. This estimate is very accurate over short distances, but exhibits unbounded drift over the long term. (2) Coarse-scale localization: GPS is used

to determine an approximate robot pose; unlike the fine-scale localization, this estimate is accurate to only a few meters, but does not suffer from cumulative drift. As an alternative to GPS, we also introduce a modified Monte-Carlo Localization algorithm that can, in principle, localize the robot using a rough map generated from satellite or aerial imagery. (3) Coarse-to-fine localization: the key challenge for urban mapping lies in mixing fine and coarse localization estimates in a manner that preserves both local continuity and global consistency (i.e., maps must display fine-scale structure, and loops must be closed). For this step, we match features occurring across multiple scans and optimize the entire robot trajectory to satisfy both local and global constraints. (4) Map generation: using the pose estimates generated by step 3, data from the scanning laser range-finders is projected into a 3D Cartesian space, forming an extended point cloud. Environmental features manifest themselves as collections of points, while open space appears as the absence of points.

In this paper, we present a high-level sketch of these four steps, paying particular attention to practical engineering issues (the underlying theory is derived almost entirely from existing work in indoor mapping [2]–[4]). We also present early results obtained for the USC University Park Campus using a Segway RMP as the mapping robot.

## II. RELATED WORK

An excellent survey paper on robot mapping has been written by Thrun [1]. This paper identifies a relatively small set of probabilistic techniques underlying most recent mapping approaches: these include Kalman Filters [5], [6], Expectation Maximization [7], incremental maximum likelihood [2], and various hybrid methods [8]. Recent work on FastSLAM algorithms (which approximate the full posterior distribution over maps using a particle filter) should be added to this list [9].

The approach described in this paper makes use of both incremental maximum likelihood estimation (for fine-scale localization) and Lu-and-Milios-style global map alignment [3]. While these methods have known limitations (they do not maintain a posterior distribution over possible maps, for example), they are sufficient for the task of urban mapping. We also make use of local map patches to enforce global map consistency; this is similar in concept, if not in detail, to the approach taken in the *Atlas* framework [10].




Fig. 1. The mapping robot: a Segway RMP equipped with scanning laser range finders (in both vertical and horizontal planes), a Garmin GPS unit and a MicroStrain IMU.

Recently, a number of authors have treated the specific problem of 3D mapping in large urban environments using scanning laser range finders. In the mobile robotics community, Montemerlo and Thrun [11] have recently demonstrated impressive results using a conventional 2D scanner mounted on a heavy-duty pan-tilt platform (producing a de-facto 3D scanner). In the digital modelling community, multiple fixed scanners have been mounted to moving vehicles to produce similar results [12].

### III. THE SEGWAY RMP

The work described in this paper is facilitated in large part by the use of a Segway RMP as the sensor platform (Figure 1). The RMP is a two-wheeled, dynamically stabilized vehicle based on the Segway HT. As a mapping platform, this vehicle has a number of advantages: it is fast, has good endurance, can support large payloads, and admits a high vantage point for sensors. This last is particularly important for the mapping task, as it greatly reduces ‘ground-clutter’ in the sensor data.

For our current 3D mapping experiments, the vehicle is configured with one laser in the horizontal plane and one or more lasers in the vertical plane. The horizontal laser is used principally for fine-scale localization, while the vertical laser(s) are used to generate dense 3D maps. The robot is also equipped with GPS and IMU sensors.

### IV. FINE-SCALE LOCALIZATION

For 3D mapping, one must determine the robot’s 6-DOF pose; i.e., its position (latitude, longitude and altitude) and orientation (roll, pitch and heading). Since roll and pitch can be measured directly with respect the gravity vector, the pose estimation problem is effectively reduced to a 4-DOF problem (latitude, longitude, altitude and heading). Furthermore, in this paper, we make the assumption that altitude is fixed, and consider only the reduced problem of estimating the robot’s latitude, longitude and heading.

For fine-scale localization, we make combine three forms of sensor data: odometry (measured by the RMP), roll

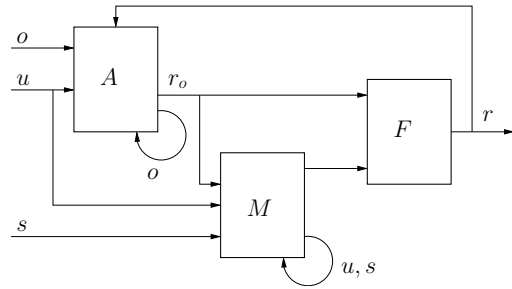


Fig. 2. Block-diagram of the fine-scale localization algorithm. The inputs are odometry data  $o$  (latitude, longitude and heading); IMU data  $u$  (roll and pitch) and laser scan data  $s$  (a set of range and bearing pairs). The output is the updated robot pose estimate  $r$  (latitude, longitude, altitude, and roll, pitch, heading). The first block  $A$  generates an updated pose estimate  $r_o$  based on the previous pose estimate  $r$  and the change in odometric pose; the second block  $M$  computes a corrected pose estimate  $r_s$  by comparing successive laser scans; the final filter block  $F$  combines these two estimates to produce the updated pose estimate  $r$ .

and pitch data (as measured by the RMP or a supplemental IMU), and laser range data. The RMP provides an odometric pose estimate (latitude, longitude and heading) based on the distance travelled by each of the two wheels; this estimate has an arbitrary offset and is subject to quite rapid drift. Figure 3, for example, shows the odometric pose estimates generated over a 2 km tour of the USC campus; in reality, the robot starts and ends at the same location. To reduce the odometric drift rate, we make use of laser range data, matching successive scans to induce a corrected pose estimate. This process is complicated by the fact that the RMP pitches quite dramatically during acceleration and deceleration; one must therefore take into account the pitch and roll information generated by the IMU, and perform a *three-dimensional* scan match. That is, the range and bearing values generated by the laser are projected into a three-dimensional Cartesian space (using IMU roll and pitch data) prior to scan matching.

The basic algorithm is illustrated in block-diagram form in Figure 2. There are two key features of this algorithm that should be noted. First, only the horizontal laser scanner (see Figure 1) is used for scan matching, as the vertical scanner provides little additional information (recall that altitude is assumed to be fixed). Secondly, urban environments contain a variety of features that make scan matching non-trivial: long grass and short shrubs, for example, generate extremely ‘noisy’ range scans that resist simple point-to-point correspondence. In addition, urban environments inevitably contain moving objects (in the form of pedestrians and vehicles) which may further confuse the scan matching algorithm. Therefore, rather than attempting to match entire range scans, we first pre-process the scans, extracting features that can be matched with high confidence and are likely to correspond to fixed objects. Specifically, we extract and match straight-line features, under the assumption that these correspond to built structures.

Figure 3 shows a comparison between the raw odometric pose estimates and the corrected values; note that, after a

tour of about 2 km, the corrected estimate is within a few tens-of-meters of the correct value.

## V. COARSE-SCALE LOCALIZATION

While the fine-scale localization described in the previous section is relatively accurate, drift cannot be entirely eliminated. We therefore make use of a second localization algorithm to generate pose estimates that are less accurate, but have bounded error. Clearly, GPS can be used for this purpose, within certain limitations; most importantly, in urban environments, GPS is occasionally unavailable due to occlusions or multi-patch effects. Figure 4, for example, shows the GPS position estimates generated on the USC campus. Estimates are shown only for those locations in which three or more satellites were visible, in which case the stated accuracy of the estimate is approximately 5 m (non-differential GPS).

As an alternative (or supplement) to GPS, we have been investigating the use of a Monte-Carlo Localization (MCL) algorithm specially modified for use in large urban environments. Importantly, while these algorithms require a prior map of the environment, this map need not contain the level of detail we seek in the final 3D map. Indeed, the map used for MCL may be a rough 2D representation obtained from aerial or satellite imagery. This is, in effect, a bootstrapping process: a rough 2D map is used for coarse localization of the robot, and a detailed 3D map is subsequently generated by making use of these coarse estimates.

The basic Monte-Carlo Localization algorithm [13], [14] uses a particle filter to maintain robot pose estimates. Compared with other Bayesian estimation algorithms (such as Kalman filters), particle filters have the key advantage of being able to represent multi-modal distributions (the robot may be in more than one place at a time). As a consequence, particle filters are largely self-initializing: given a sufficiently large particle set, the filter will always converge to the correct robot pose.

Most implementations of the MCL have been used in indoor environments, and assume that we have access to both good maps and reliable sensor information. Thus, indoor maps typically partition the environment into regions that are either *occupied* or *free*. In our case, however, the urban environment is not well mapped, and contains a variety of features that give rise to unpredictable sensor readings (such as trees, bushes, pedestrians and vehicles). We therefore introduce a third type of region into the map: a *semi-occupied* region that may or may not generate laser returns. Thus, for example, buildings are represented as occupied regions, streets and walk-ways are represented as free regions, and parking spaces, grass, and gardens as semi-occupied regions. We also assume that semi-occupied regions cannot be occupied by the robot, further constraining the set of possible poses. Figure 5 shows a published map of the USC campus, along with the corresponding map used by the MCL algorithm (note the semi-occupied regions around most of the buildings).

Our Urban MCL implementation has two additional enhancements: it uses the adaptive sampling method described in [15] to control the size of the particle set (resulting in very significant speed-up), and laser data is pre-processed using IMU data to compensate for the pitching of the Segway RMP platform.

Figure 4 shows the pose estimates generated by the urban MCL algorithm over a tour of the campus. These results were generated by processing the raw data twice: during the first pass, the robot's initial pose was unknown, and the filter took some time to converge to a singular pose value; once this pose was established, the data was subsequently processed *backwards* to generate a complete trajectory for the robot. Note that MCL provides pose estimates in areas where GPS is unavailable, and that GPS and MCL estimates are sometimes in disagreement (in this case, visual analysis suggests that it is the GPS estimates that are misleading). The estimates are also complimentary, in the sense that MCL is most accurate in the heavily built-up sections of the campus, while GPS is most accurate in open areas. Note that while we have not yet attempted to fuse GPS and MCL estimates, these results suggest that such a fusion could be advantageous.

## VI. COARSE-TO-FINE LOCALIZATION: CLOSING LOOPS

The fine and coarse scale localization methods described in the previous sections have complementary properties. Using fine-scale localization alone, one can project data from the laser range-finders into a 3D Cartesian space, generating a point cloud representation of the environment (this is similar to the indoor mapping approach described in [16], where a horizontally mounted laser is used for fine-scale localization, while a vertically mounted laser is used to generate 3D volumetric data). Figure shows a point cloud generated using this process; the representation is highly detailed, and includes features such as trees and bushes, road signs, stairs, doorways, parked vehicles and pedestrians. Fine-scale localization is, of course, subject to drift, and maps produced in this way will eventually become inconsistent. In contrast, coarse-scale localization is not subject to drift: the robot's pose relative to the environment is known with some bounded error. Unfortunately, since this error bound can be quite high (of the order of several meters), coarse-scale localization alone cannot be used to generate detailed maps of the environment. Thus, in order to generate maps that are both locally detailed and globally consistent, one must combine both forms of localization.

Our basic approach to this problem is as follows. First, fine-scale localization is used to generate a series of *sub-maps*, each of which corresponds to a short piece of the robot's total trajectory. Each sub-map has a *pose* and is subject to a set of *constraints* (see Figure 6). The task, then, is to determine the set of sub-map pose estimates that best satisfies the constraints (a maximum likelihood estimate). In practice, this can be done fairly easily using any of a number of non-linear optimization algorithms.




Fig. 3. Robot trajectory using odometry, and odometry plus laser scan matching. The robot executes a closed traverse of approximately 2 km.




Fig. 4. Comparison of robot trajectory estimates using GPS and MCL; note that GPS estimates are sometimes unavailable due to satellite occlusion.

Two types of constraints are present: those arising from GPS or MCL pose estimates (which constrain the global pose of each individual sub-map), and those arising from feature matching (which constrain the relative pose of pairs of sub-maps). The features in question are vertical planar regions extracted from each sub-map using a RANSAC algorithm [17]; these regions typically correspond to the exterior walls of buildings.

Note that the global constraints imposed by coarse localization (GPS or MCL) are vital to the success of

this approach; these constraints are applied first, generating a good initial fit between sub-maps. Features are subsequently matched using a simple nearest-neighbor algorithm, and this rough fit is iteratively refined.

## VII. MAPPING RESULTS

Figure 7 shows a map of the USC campus generated using the method described above. Raw data was captured by the Segway RMP over a 2 km tour of the campus, with multiple loops around both individual buildings and entire



Fig. 5. (a) Scanned map of the USC University Park campus. (b) Induced map, showing free space (white), occupied space (black), and semi-occupied space (gray). The particle filter estimate is indicated by the arrow.

*IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 419-424  
Sendai, Japan, Sep 2004*

*IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 419-424  
Sendai, Japan, Sep 2004*




Fig. 6. Feature-based fitting of sub-maps (using vertical planar features). (a) Coarse alignment using GPS or MCL; the feature correspondences are indicated arrows. (b) Final alignment using feature correspondences; note that the coarse pose estimates are now displaced somewhat.

blocks. Average speed of the robot was 1.5 m/sec. The map was generated off-line, but in real time (i.e., the time taken to generate the map is less than the time taken to tour the environment).

The final map is rendered as a point cloud containing approximately 8 million points, and can be manipulated interactively using standard VRML viewing tools. Figure 7(b) shows a detail from the map together with a corresponding still image; in addition to the buildings, palm trees and lamp-posts are clearly visible. Note that, for this tour, the lasers were mounted approximately one meter from the ground; hence only features above waist-level are visible in the map. Other configurations of the laser are possible, such that more complete maps are generated.

## VIII. CONCLUSION AND FURTHER WORK

This paper describes work-in-progress towards large-scale urban 3D mapping, and we are continuing to explore a number of extensions and alternatives to the algorithms described sections IV and VI. Nevertheless, our results to date indicate that 3D mapping in urban environments is both technically achievable and practically useful; the

level of the detail in the maps shown in 7, for example, is striking.

On the other hand, while these point cloud maps are visually compelling, they can be difficult to work with in robotic applications; much of our future work, therefore, is focused on the problem of transforming these point clouds into alternate representations (using voxels or polygons, for example). We are also considering the problem of acquiring and representing maps of *non-static* environments; thus, for example, we are attempting to fuse data acquired over several days into a single map that highlights variations in the environment.

## REFERENCES

- [1] S. Thrun, "Robotic mapping: A survey," in *Exploring Artificial Intelligence in the New Millenium*, G. Lakemeyer and B. Nebel, Eds. Morgan Kaufmann, 2002.
- [2] J. Gutmann and K. Konolige, "Incremental mapping of large cyclic environments," in *Proceedings of the IEEE International Symposium on Computational Intelligence in Robotics and Automation (CIRA)*, 2000.
- [3] F. Lu and E. Milios, "Globally consistent range scan alignment for environment mapping," *Autonomous Robots*, vol. 4, pp. 333-349, 1997.




Fig. 7. (a) USC campus rendered as a point cloud (elevated view). (b) Detail view (ground level); note the palm tree and lamppost. (c) Photograph of the scene (the photograph and the map were taken at different times, so not all features are identical).

- [4] A. Howard, "Multi-robot mapping using manifold representations," in *IEEE International Conference on Robotics and Automation*, New Orleans, Louisiana, Apr 2004, pp. 4198–4203.
- [5] S. Borthwick and H. Durrant-Whyte, "Simultaneous localisation and map building for autonomous guided vehicle," in *Proceedings of the IEEE/RSJ/GI International Conference on Intelligent Robots and Systems*, vol. 2, 1994, pp. 761–8.
- [6] J. J. Leonard and H. Durrant-Whyte, "Simultaneous map building and localization for an autonomous mobile robot," in *Proceedings of the IEEE/RSJ International Workshop on Intelligent Robots and Systems*, vol. 3, 1991, pp. 1442–1447.
- [7] S. Thrun, D. Fox, and W. Burgard, "A probabilistic approach to concurrent mapping and localisation for mobile robots," *Machine Learning*, vol. 31, no. 5, pp. 29–55, 1998, joint issue with Autonomous Robots.
- [8] S. Thrun, "A probabilistic online mapping algorithm for teams of mobile robots," *International Journal of Robotics Research*, vol. 20, no. 5, pp. 335–363, 2001.
- [9] D. Haehnel, W. Burgard, D. Fox, and S. Thrun, "An efficient FastSLAM algorithm for generating maps of large-scale cyclic environments from raw laser range measurements," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, Las Vegas, Nevada, U.S.A., Oct 2003.
- [10] M. Bosse, P. Newman, J. Leonard, M. Soika, W. Feiten, and S. Teller, "An Atlas framework for scalable mapping," in *IEEE International Conference on Robotics and Automation*, Taipei, Taiwan, September 2003.
- [11] M. Montemerlo and S. Thrun, "Large-scale robotic 3d mapping of urban structures," in *9th International Symposium on Experimental Robotics*, Singapore, June 2004, to appear.
- [12] H. Zhu and R. Shibasaki, "Reconstructing urban 3d model using vehicle-borne laser range scanners," in *Third International Conference on 3-D Digital Imaging and Modelling*, Quebec City, Quebec, Canada, May 2001, pp. 349–356.
- [13] F. Dellaert, D. Fox, W. Burgard, and S. Thrun, "Monte carlo localization for mobile robots," in *IEEE International Conference on Robotics and Automation (ICRA99)*, May 1999.
- [14] S. Thrun, D. Fox, W. Burgard, and F. Dellaert, "Robust Monte Carlo localization for mobile robots," *Artificial Intelligence Journal*, vol. 128, no. 1–2, pp. 99–141, 2001.
- [15] D. Fox, "KLD-sampling: Adaptive particle filters," in *Advances in Neural Information Processing Systems 14*. MIT Press, 2001.
- [16] C. Martin and S. Thrun, "Real-time acquisition of compact volumetric maps with mobile robots," in *Proceedings of the IEEE International Conference on Robotics and Automation*, San Francisco, California, U.S.A., April 2000.
- [17] M. A. Fischler and R. C. Bolles, "Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography," *Communications of the ACM*, vol. 24, pp. 381–395, 1981.