

Robust localization algorithms for an autonomous campus tour guide*

Richard Thrapp and Christian Westbrook and Devika Subramanian
Department of Computer Science
Rice University, Houston TX 77005, USA
{richardt, cw, devika}@cs.rice.edu

February 15, 2001

Abstract

This paper describes a robust localization method for an outdoor robot that gives tours of the Rice University campus. The robot fuses odometry and GPS data using extended Kalman filtering. We propose and experimentally test a technique for handling two types of non-stationarity in GPS data quality: abrupt changes in GPS position readings caused by sudden obstructions to line of sight access to satellites, and more gradual changes caused by disparities in atmospheric conditions. We construct measurement error covariance matrices indexed by number of visible satellites and switch them into the localization computation automatically. The matrices are built by sampling GPS data repeatedly along the route and are updated continuously to handle drift in GPS data quality. We demonstrate that our approach performs better than extended Kalman filters that use only a single error covariance matrix. With a GPS receiver that delivers 1 meter accuracy, we have been able to localize good to 40 cm through a challenging route in the Engineering Quadrangle of Rice University.

1 Introduction

Our goal is to build an autonomous mobile robot that gives tours of the Rice University campus (see Figure 2). While there are many successful mobile robots that autonomously navigate and explore *indoor* environments like museums [13, 10] and offices [8] over extended periods of time, there are far fewer examples¹ in the outdoor arena, most being research prototypes [1, 3, 11, 12, 6]. There are several reasons why the design of autonomous outdoor robots is challenging: outdoor environments vary more than those indoors, and we do not have good characterizations for them, and robust solutions to the localization

* We thank the George R. Brown School of Engineering, Rice University for its generous support of our work.

¹The new grass mowing robots guided by externally placed wires at lawn boundaries and around obstacles are not considered here, because they require the environment to be instrumented for them.

problem for outdoor robots are still under development. We believe that the pursuit of the engineering goal of building an autonomous outdoor tour guide for Rice University, a tree-filled oasis in the urban metropolis of Houston, will help us make progress on the scientific goal of characterizing outdoor environments for which reliable mobile robot navigation algorithms can be designed and built.

This paper describes the localization algorithms used by our robot, their integration with the navigation algorithms, and presents preliminary experimental results. An important ground rule we followed to ensure portability of our methods was to allow no instrumentation or modification of our environment. A key issue for outdoor robots is the choice of sensors for performing localization. While indoor robot localization algorithms make extensive use of sonars [14, 13, 5], they are virtually useless for outdoor robots, as the environment consists mostly of vast open spaces where the sensors return no valid range information. They are, however, useful (together with bump sensors) for local obstacle avoidance. With a view to determining the smallest set of sensors needed to provide robust localization and navigation in a campus environment, we have currently limited ourselves to data from odometry, which provides location and orientation information relative to a start position, and GPS, which provides global positioning and heading data. Our objective is to understand experimental limits on the accuracy of localization achievable using just these two sources of information.²

Our robot needs to give tours throughout the day without requiring human intervention. It therefore requires localization accuracies of at least 40 cm at all times during its entire operation. This is about half the width of the walkways around our campus. Unlike office environments in which short term localization errors are generally non-fatal, small errors can have catastrophic consequences for our outdoor robot e.g., falling off of a curb or missing a sidewalk and rolling onto a busy street.

²We plan to add vision to our robot to augment odometry and GPS for localization.

It is well known that dead-reckoning using pure odometry is not a very robust localization technique for robots [2, 6] that cover long distances, and are in continuous operation over extended periods of time. This is because errors in odometry accumulate over time due to inaccuracies in the kinematic model, precision limitations of encoders, and unobservable factors like wheel slippages that are not accounted for in the kinematic equations. Kalman filtering of carefully calibrated odometric data with state measurement signals provided by a redundant sensor (e.g., a gyroscope) can provide significant improvements [6]. However, they still cannot on their own provide localization accuracies of 40 cm over extended periods of time, as needed for our problem.

GPS is now a standard technique for obtaining absolute position information for outdoor robots [6, 3, 11, 1, 12]. Our GPS receiver (an Invicta 210S) can provide position information accurate to about a meter ($1\sigma = 100\text{cm}$). It also provides estimations of the current heading by using the Doppler shift of the satellite signals. The accuracy of this estimate depends greatly on the speed at which the GPS antenna is moving, but at the top speed of our robot, we have determined that $1\sigma = 1.84$ radians. Differential GPS systems like RTK GPS can provide centimeter level resolution, however they cost an order of magnitude more, and their performance is very sensitive to the number of visible satellites [3]. GPS information alone is not sufficient to achieve the localization accuracies needed for our application, because the tour guide's route on our campus is largely covered with trees and runs very close to tall buildings which obstruct line of sight access to the GPS satellites. Further, atmospheric conditions degrade the quality of the GPS signals in varying ways at different locations.

In this paper, we adapt Extended Kalman Filtering (EKF) [7, 9] used in [6, 3, 12] for mobile robot localization to fuse odometry and GPS data. Our innovation is the addition of a dynamic mechanism to handle non-stationarities in GPS data quality. Because of the rapid change in quality of GPS data when the view of satellites is obstructed in an urban campus environment, the standard approach of using a single covariance matrix to model measurement errors in GPS data is not adequate. This approach necessitates artificially increasing variances more than is usually needed, causing slower convergence of the Kalman filter. If the data quality diminishes suddenly, the Kalman filter does not account for the reduction in quality quickly enough, and our robot performs poorly because inaccurate GPS data is weighted too heavily. Conversely, if our robot gains a line-of-sight path to additional satellites after having adapted to a reduced quality GPS signal, the increased variances cause the Kalman filter to not take advantage of the higher quality data, so control errors in our robot accrue again. We show that by using a finite number of error

covariance matrices in different GPS quality situations, creating them as needed and updating them dynamically, it is possible to handle changes in GPS signal quality quickly and effectively.

The paper is organized as follows: In Section 2, we give a brief description of the tour guide task as well the mechanical and sensor configuration of our robot. We lay out the basic EKF algorithm upon which our localization algorithm is based. We then describe the extension to the EKF algorithm that dynamically constructs error covariance matrices to handle non-stationarities in the GPS data quality. Section 3 contains a brief discussion of issues in designing control algorithms for navigation that use position and heading estimates produced by the localization algorithm. In Section 4, we provide details of our experiments and demonstrate that the dynamic measurement error covariance matrix generation handles rapid changes in GPS signal quality well. We conclude with a brief summary of our ongoing work in designing methods for human interaction with our robot.

2 The Tour Guide Task and Robot

2.1 The task and the robot

The tour guide task requires the ability to navigate in a dynamic, uninstrumented, potentially dangerous (vehicular traffic on streets, sharp curbs, moving obstacles such as animals and people, etc.) urban environment. In addition, the robot needs to interact with a tour group in an interesting and informative manner. Interaction is tightly interwoven with navigation: the robot needs to be aware of its location so it can use its location context to answer questions appropriately.

Our tour guide robot is an ATRV Jr. from RWI Inc, named Virgil³ (Figure 1). It is a four-wheeled robot designed for outdoor use and comes equipped with an array of sonars and odometry. We added a GPS receiver used typically in marine applications which receives real-time corrections from the Coast Guard station at Galveston. We also added touch-sensitive bumpers for obstacle detection and avoidance. The wheels on the same side of the robot are mechanically coupled. The raw encoder information is not directly available, neither do we have access to the full kinematic model of the robot which is used by the on-board odometry computation to provide integrated measures like distance traveled and change in orientation in a given sampling interval. Both odometric and GPS data are sampled at 10Hz in our robot.

³Virgil is named after the guide in Dante's *Inferno*.



Figure 1: Virgil: The Rice campus tour guide

2.2 Odometry

Before integrating data from different sources, we calibrated the odometry using the GPS receiver. Because odometry measures the number of rotations in the wheels rather than the actual distance traveled, several sources of error can accumulate. Largely, systematic error is due to tire size miscalculations: as the tires wear down, the amount of linear distance traversed reduces in comparison to the number of rotations the tires travel. In addition, because the tires are constantly being worn down, this analysis must be reperformed periodically to estimate new systematic error values. This error can be directly compensated for by scaling the commands given to the drive system. In addition, other sources of error such as slippage and surface imperfections result in a random component to the error whose variance can also be approximated through repeated trials.

By traveling, according to odometry, in straight lines for fairly large distances (20 meters), and comparing the odometry's results for distance traveled with GPS data averaged over 100 readings, we are able to determine approximate values for systematic and random errors in odometry. In addition, to limit the effect of systematic inaccuracies in translating GPS coordinates into local coordinates, we performed the test from many different starting positions and headings. It is possible that this method reports a slightly higher than actual random error rate due to GPS inaccuracies. To determine turning error, we follow a similar procedure of moving forward a distance and using GPS data at the endpoints of that movement to approximate the current heading, using odometry to turn a preset angle, and then move forward again to calculate the true angle turned using averaged GPS data. This is inherently less accurate than calculating distance over a straight path, but gives a reasonably good approximation of the true errors. Again, because of GPS inaccuracies, the random component in the measured error will likely be larger than the true random error.

After performing these tests, we observed that our robot's odometry consistently under-maneuvered both

Movement	Systematic Error	Random Error
Translational	-0.0290	0.0036
Rotational	-0.0492	0.0589

Table 1: Systematic errors and random error variances measured in odometry related to the two types of motion supported.

while traveling in straight lines and while turning as a result of the smaller than expected size of the worn tires. These results are presented in Table 1 as a ratio of error to distance traveled for 29 trials of both the translational and rotational measurements which approximates the actual error. Systematic error is accounted for directly in the odometry system, increasing the robot's perception of how far it has traveled by the appropriate ratio, and random error is handled as uncertainty in the data fusion process.

2.3 Extended Kalman Filter

Kalman filtering is a well known technique for state and parameter estimation [7, 9]. The standard Kalman filter assumes that the controlled process is governed by a linear stochastic difference equation. An extended Kalman filter handles non-linear stochastic processes by linearizing about the current mean and covariance.

In the 2D outdoor robot localization problem, the state of the robot is its position and orientation (x, y, θ) in a fixed frame of reference. The state $(0,0,0)$ is the geographic center of the Rice campus (Baker Fountain). All (x, y) positions are measured in centimeters north and east relative to this location, and the orientation θ is the angle from due north.

The robot's state evolves according to the following system of non-linear stochastic difference equations. The state of the system at time k is (x_k, y_k, θ_k) . The wheel encoders yield, at each sampling period the translation u_{tk} along the heading θ_k and a rotation $u_{\theta k}$. These equations relate the state at time $k + 1$ to the state at time k , and the internally sensed translation u_{tk} in the direction θ_k and rotation $u_{\theta k}$ in the interval between times k and $k + 1$. The zero mean vector $w_k = N(0; Q_k)$ represents (normally distributed) noise in the state evolution process.

$$\begin{aligned} x_{k+1} &= x_k + \cos\theta_k u_{tk} + w_{kx} \\ y_{k+1} &= y_k + \sin\theta_k u_{tk} + w_{ky} \\ \theta_{k+1} &= \theta_k + u_{\theta k} + w_{k\theta} \end{aligned}$$

These equations can be summarized as follows, where s_k is the state of the system at time k and $u_k = (u_{tk}, u_{\theta k})$ is the vector obtained from the encoders for the period from time k to time $k + 1$.

$$s_{k+1} = f(s_k, u_k, w_k)$$

We use the GPS signal to determine the measurement error between the actual state and the internally computed state above. We model the measurement process as follows

$$z_k = s_k + v_k$$

where z_k is the measurement of the actual state s_k at time k , $v_k = N(0; R_k)$ is a zero mean measurement noise vector. The measurement and process noise vectors are assumed to be independent of one another, and to have normal probability distributions represented by the vectors Q_k and R_k of variances for the three state components x , y and θ . There is an important subtlety here caused by the fact that the GPS data is in a different coordinate system from the one maintained by our robot. We use the equations in [4] to convert global GPS coordinates into local (Baker Fountain relative) coordinates. These equations perform very accurate translations taking the curvature of the earth's surface into account, and are in wide use in the agricultural world. As a result, we do not observe the loss of precision in translating between local and global coordinates noted in [6].

We now show the prediction and update steps in the EKF which combines internal state (odometric) and external measurement (GPS) data. The equations for the prediction step are:

$$\begin{aligned} s_{k+1/k} &= f(s_{k/k}, u_k, 0) \\ P_{k+1/k} &= A_k P_{k/k} A_k^T + Q_k \end{aligned}$$

$s_{k+1/k}$ stands for our prediction of the state vector for time $k + 1$ given internal sensor information from time k and knowledge of the state at time k . $P_{k+1/k}$ is the a priori estimate of the error covariance, i.e. the covariance of the difference between the actual state and the state predicted on the basis of measurements till time k . A_k is the Jacobian of the process $f(\cdot)$ with respect to the state vector $s = (x, y, \theta)$. By differentiating the state evolution equation with respect to s , we obtain the following matrix:

$$A_k = \begin{bmatrix} 1 & 0 & -u_{tk} \sin(\theta_k) \\ 0 & 1 & u_{tk} \cos(\theta_k) \\ 0 & 0 & 1 \end{bmatrix}$$

Note that Q_k is the process error covariance matrix for time k . For our robot, the prediction equations reduce to:

$$\begin{aligned} s_{k+1/k} &= f(s_{k/k}, \mathbf{u}_k, 0) \\ P_{k+1/k} &= A_k P_{k/k} A_k^T + Q_k \end{aligned}$$

Essentially, these two prediction equations project the state and covariance estimates from time step k to $k + 1$.

The update equations correct the state and covariance estimates with the measurement z_k . m_k is the GPS reading at time k converted into the local coordinate system for the robot. We compute the Kalman gain K and use it to

correct s and P as follows. R_k is the measurement noise covariance matrix.

$$\begin{aligned} K &= P_{k+1/k} (P_{k+1/k} + R_k)^{-1} \\ s_{k+1/k+1} &= s_{k+1/k} + K(m_k - s_{k+1/k}) \\ P_{k+1/k+1} &= (I - K)P_{k+1/k} \end{aligned}$$

K is the weight allocated to the state and error covariance correction. The accuracy of R_k determines the effectiveness of the EKF, and because the accuracy of the GPS measurement changes based on a variety of factors, a single predetermined R_k in many cases does not achieve good localization accuracy. Furthermore, because the accuracy of the GPS receiver can change almost instantaneously as the robot Virgil passes below a tree or near a building when one or more satellites are obstructed from view, slowly updating measurement error variance using new data points is not effective. If we do not model the temporal variations in the measurement error R_k correctly, then the filter will be unable to respond quickly enough to sudden deterioration or sudden improvement in GPS data quality.

2.4 Handling non-stationarity in R_k

There are two kinds of changes of R_k in time, thus we have developed two schemes to handle these changes. One change occurs abruptly and is caused by the number of visible satellites changing; the other is more gradual and represents a drift as atmospheric conditions and other factors affect signals. By analyzing the GPGGA NMEA string returned by the receiver, we can determine when additional satellites are acquired or lost, and keep a distinct R_k for each number of satellites. Over time, each of these values is updated dynamically based on new data. We switch between these R_k 's as determined by the number of visible satellites and are able to obtain much faster convergence to the true state than by using a single measurement error vector which averages them all. In addition to the GPGGA string, the GPS receiver supplies the GPRMC message which provides the current direction of travel of the robot, but the accuracy of this value varies depending on the speed of the robot. We attack this problem using a similar method of swapping R_k 's depending on ranges of speed, almost completely ignoring data that is read at very low speed.

To handle the second kind of change, we examine the evolution of the error term $(m_k - s_{k+1/k})(m_k - s_{k+1/k})^T$ over N time steps, and correct R_k to approach the averaged error term measured over the N time steps. This process tracks the drift in R_k and yields better convergence properties for the extended Kalman filter, as documented in our experimental section.

3 Integration of localization with navigation

Our localization algorithm provides estimates of the current position and orientation of the robot in the campus coordinate frame centered at Baker Fountain. We specify the tour as a list of way points in this reference frame. We have built a simple proportional controller that drives the robot from the start point through the way points in order of occurrence. This controller servos on the difference between the current and desired position and orientation. Low level obstacle detection and avoidance is performed with sonars and two bump sensors.

The control correction choice we explore is whether to turn in place (forward-speed = 0) to achieve a heading correction, or whether to turn and move forward at the same time. The correction frequency controls the rate or conditions under which corrective motor commands (turn and speed) are issued. All of our controllers make active use of the stability of the Kalman filter state estimates to decide when to issue corrective actions. They wait for the confidence in the state estimates to achieve a certain threshold (measured by the difference in successive Kalman filter estimates) before changing turn and speed.

The first controller we built attempts to simultaneously reduce position and heading differences, and therefore turns as fast as possible to reduce the difference between current and desired heading without changing the forward speed. This controller is very sensitive to choice of the action confidence threshold. If it is set too high, the robot continues to move forward at its current speed in the current direction for a longer period of time. If the robot is off course to begin with, it moves farther from the approved path before correcting, and the correction turn is larger. Lowering the confidence threshold results in a greater number of incorrect turns causing small movements away from the desired path, with the robot drunkenly weaving around the desired path.

To correct this problem, we built a second controller identical to the first except that it cuts forward speed to zero while turning. For a set time interval (half a second) it turns toward the goal, reducing the difference between current and desired heading. It interleaves forward motion (to reduce difference between (x, y) locations) with in-place turns (to reduce heading differences). The advantage of eliminating forward motion during turns is that our skid-steer robot turns more accurately and stays closer to the planned route. However, the jerkiness of consecutive turn and forward motion phases proved distracting for people in tour groups.

In our final and most successful attempt, we returned to making course corrections during forward movement. However, instead of attempting to reduce the difference between current and goal heading as soon as possible, we re-

	Distance	Heading	$\sigma_{Heading}$
Uncorrected GPS	1.16	1.69	1.58
EKF: Static R_k	1.33	2.05	1.64
EKF: Dynamic R_k	1.08	1.33	1.54

Table 2: For each data fusion algorithm used, this table presents the ratio of the distance traveled to the total straight line distance of the tour as well as the average and standard deviation of the difference between the heading of the straight line path and the heading chosen by the robot controller.

duce the heading difference slowly, causing the robot to gently arc its way to the goal state. By approaching the goal heading slowly, we reduce our reliance on knowing the exact difference between the current and goal headings; knowing the sign and order of magnitude of the difference suffices. This causes drifts and over-corrections to drastically reduce and results in smooth and fluid motion between way points on the tour.

4 Experimental Results

The way points in the tour of the Engineering Quadrangle are shown in Figure 2. This is a challenging route that runs close to several tall buildings, large granite sculptures and several rows of trees. In our experiments we compare the behavior of (1) a dual controller that uses GPS when available, and odometry when it is not, (2) an EKF based controller that fuses odometry and GPS data without considering non-stationarities in the measurement errors, (3) our EKF based controller that tracks and updates measurement error covariances in time.

At particularly troubling times for the GPS receiver, the first method falls back entirely to odometry, until GPS and odometry can be resynchronized according to average velocity readings. The result of this is that the robot sometimes travels without any GPS data for extended periods of time, and localization errors accumulate until it is completely off course. The second method works well when many satellites are visible and the quality of the GPS data is good. However, when satellites are lost, which often happens in the last stretch of the tour, the robot's ability to localize correctly degrades significantly. Tracking changes in the measurement error covariance matrices dramatically improves localization accuracy, because inaccurate GPS data is weighted less heavily with respect to odometry, and their combination provides information that is able to keep the robot within 40 cm of the desired route.

Table 2 presents results averaged over three runs of the three methods. For each method we measure the deviation from a straight line path between the way points on




Figure 2: A top view of Rice's Engineering Quadrangle. The sculptures in the grassy centers are giant granite monoliths called 45, 90, and 180, corresponding to their angles of inclination. The tour path is given by points a through h .

the tour. We take the ratio of the length of the actual path covered to the length of the shortest path between the way points. We also provide the mean and standard deviation of the differences in heading (sampled at 10Hz) between the true heading between two way points and the orientation of the robot between those two way points. These two measures give us a sense of the accuracy of localization and the smoothness of the control policies that use the localization measures. Figures 3 and 4 show the actual trajectories of the robot over the above three runs with static and dynamic R_k 's. These runs were gathered when the variability in the GPS data quality was not as high, so the difference between the two approaches surfaces in the consistency of the trajectories generated with dynamic R_k 's.

It should be noted, however, that in areas with fairly steep or inconsistent inclines, our robot's odometry is incapable of determining the angle of the incline at which the robot is traveling, and thus can not determine the horizontal speed of travel. Furthermore, this information is difficult to deduce from the GPS readings; while altitude can be determined from GPS, the value is significantly less reliable than latitude and longitude readings. To compensate for this uncertainty, the variances associated with odometry would increase greatly. It is conceivable that under these conditions the extended Kalman filter even with the mechanisms for handling non-stationarity in GPS signal quality will be unable to converge on the correct location and heading of the robot. However, the addition of an inertial navigation system (INS) with six degrees of freedom would aid in correcting these errors. We would need to construct a new




Figure 3: Trajectories with EKF fusion of GPS and odometry using static R_k .




Figure 4: Trajectories with EKF fusion of GPS and odometry using dynamic R_k that accounts for non-stationarities in GPS data quality.

EKF that fuses odometry, GPS and INS data appropriately. Since the Rice campus is relatively flat, our experimental results do not reflect this problem.

5 Discussion and Conclusions

Our work builds on several existing pieces of work in designing mechanisms for fusing odometry and GPS data for localizing outdoor mobile robots. The idea of using short term localization based on dead reckoning when GPS data is available in several papers including [6, 1, 12]. The use of extended Kalman filtering is proposed in [6]. However, they advocate the use of a dual controller where at any given time, one of a pure GPS localizer or a Kalman filter that fuses gyro and odometry data is used. For our problem, combining the GPS data with odometry using different measurement error covariance matrices that track the

non-stationarities in the quality of the GPS signal proved to be the more effective technique. The use of differential GPS (RTK GPS) is advocated in [3]; we believe their technique can also benefit from our schemes for handling non-stationarities in the GPS signal quality.

Our current work is in extending the range of the tour to go beyond the Engineering Quadrangle to cover the rest of the Rice University campus. The new route will require the robot to cross a busy campus street and we are working on mechanisms to ensure our robot's safety during this crossing. We are also incorporating voice recognition with dynamic grammars to enable our robot to recognize questions put to it by members of tour groups and to respond in a location-aware manner to such questions.

References

- [1] T. Aono, K. Fujii, S. Hatsumoto, and T. Kamiya. Positioning of vehicle on undulating ground using gps and dead reckoning. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 3443–3448, 1998.
- [2] B. Barshan and H. F. Durrant-Whyte. Inertial navigation systems for mobile robots. *IEEE Transactions on Robotics and Automation*, 11(3):328–342, 1995.
- [3] D. Bouvet and G. Garcia. Civil-engineering articulated vehicle localization: solutions to deal with gps masking phases. In *Proceedings of the IEEE International Conference on Robotics and Automation*, 2000.
- [4] C. G. Carlson. What do latitude and longitude, readings from a gps receiver mean? Technical report, South Dakota State University, 1999.
- [5] A. Elfes. Sonar-based real-world mapping and navigation. In *IEEE Journal of Robotics and Automation*, volume RA-3(3), pages 294–265, 1987.
- [6] P. Goel, S. I. Roumeliotis, and G. S. Sukhatme. Robust localization using relative and absolute position estimates. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 1999.
- [7] R. E. Kalman. A new approach to linear filtering and prediction problems. In *Transaction of the ASME—Journal of Basic Engineering*, pages 33–45, March 1960.
- [8] K. Konolige and K. Myers. The saphira architecture for autonomous mobile robots. Technical Report downloadable from <http://www.ai.sri.com/> konolige, to appear in an MIT Press collection, SRI International, 2000.
- [9] P. S. Maybeck. *Stochastic Models, Estimation, and Control*. Academic Press, New York, NY, 1979.
- [10] I. Nourbakhsh, J. Bobenage, S. Grange, R. Lutz, R. Meyer, and A. Soto. An affective mobile educator with a full-time job. *Artificial Intelligence*, 114((1-2)):95–124, 1999.
- [11] L. Pampagnin, F. Peyret, and G. Garcia. Architecture of a gps-based guiding system for road compaction. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 2422–2427, 1998.
- [12] S. Sukkarieh, E. M. Nebot, and H. F. Durrant-Whyte. Achieving integrity in an ins/gps navigation loop for autonomous land vehicle applications. In *Proceedings of the 1998 IEEE International Conference on Robotics and Automation*, pages 3437–3443, 1998.
- [13] S. Thrun, M. Beetz, M. Bennewitz, W. Burgard, A. Cremers, F. Dellaert, D. Fox, D. Haehnel, C. Rosenberg, N. Roy, J. Schulte, and D. Schulz. Probabilistic algorithms and the interactive museum tour-guide robot minerva. *International Journal of Robotics Research*, 2000.
- [14] S. Thrun, A. B. W., B. D. Fox, T. Frohlinghaus, D. Hennig, T. Hofmann, M. Krell, and T. Schmidt. Map learning and high-speed navigation in rhino. *MIT Press, Cambridge, MA*, 2000.