

A System for Volumetric Robotic Mapping of Abandoned Mines

Sebastian Thrun[†], Dirk Hähnel^{†‡}, David Ferguson[†], Michael Montemerlo[†], Rudolph Triebel[†],
Wolfram Burgard[‡], Christopher Baker[†], Zachary Omohundro[†], Scott Thayer[†], William Whittaker[†]

[†] School of Computer Science
Carnegie Mellon University
Pittsburgh, PA, USA

[‡] Department of Computer Science
University of Freiburg
Freiburg, Germany

Abstract—This paper describes two robotic systems developed for acquiring accurate volumetric maps of underground mines. One system is based on a cart instrumented by laser range finders, pushed through a mine by people. Another is a remotely controlled mobile robot equipped with laser range finders. To build consistent maps of large mines with many cycles, we describe an algorithm for estimating global correspondences and aligning robot paths. This algorithm enables us to recover consistent maps several hundreds of meters in diameter, without odometric information. We report results obtained in two mines, a research mine in Bruceton, PA, and an abandoned coal mine in Burgettstown, PA.

I. INTRODUCTION

The lack of accurate maps of inactive, underground mines poses a serious threat to public safety. According to a recent article [3], “tens of thousands, perhaps even hundreds of thousands, of abandoned mines exist today in the United States. Not even the U.S. Bureau of Mines knows the exact number, because federal recording of mining claims was not required until 1976.” The lack of accurate mine maps frequently causes accidents, such as a recent near-fatal accident in Quecreek, PA [18]. Even when accurate maps exist, they provide information only in 2-D, which is usually insufficient to assess the structural soundness of abandoned mines.

Hazardous operating conditions and difficult access routes suggest that robotic exploration and mapping of abandoned mines may be a viable option. The idea of mapping mines with robots is not new. Past research has predominantly focused on acquiring maps for autonomous robot navigation in active mines. For example, Corke and colleagues [8] have built vehicles that acquire and utilize accurate 2-D maps of mines. Similarly, Baily [1] reports 2-D mapping results of an underground area using advanced mapping techniques. None of these techniques generate volumetric maps of mines.

In general, the mine mapping problem is made challenging by the lack of global position information underground. As a result, mine mapping must be approached as a *simultaneous localization and mapping*, or SLAM, problem [10], [15], [20]. In SLAM, the robot acquires a map of its environment while simultaneously estimating its own position relative to this map. The SLAM problem is known to be particularly difficult when the environment possesses cyclic structure [5], [6], [13], [21]. This is because cycles pose

hard correspondence problems that arise due to the (relatively) large position error accrued by a vehicle when closing cycles. Mines often contain a large number of cycles, hence the ability to handle cycles is essential for successful approaches to mapping mines.

This paper describes a SLAM algorithm for acquiring 3-D models of underground mines that can accommodate multiple cycles. Our algorithm uses a scan matching algorithms for constructing 2-D mine maps described in [14]. To close cycles, however, it utilizes an iterative correspondence algorithm based on the *iterative closest point* algorithm (ICP) [4], adapted to the problem of establishing correspondence in cyclic maps. 3-D maps are generated by applying scan matching to 3-D measurements after the 2-D mapping is complete.

Our algorithm has successfully enabled two robotic systems to acquire 3-D maps of mines. The first such system consists of an instrumented cart, which is pushed manually through a mine. This system is a low-cost solution to the mine mapping problem, but it can only be brought to bear in environments accessible to people. Our second system consists of a rugged robotic platform equipped with laser range sensors. Abandoned mines, when dry, are often subject to low oxygen levels, poisonous gases, and they may be structurally unstable. Since bringing humans into such mines exposes them to a serious danger of life, the employment of autonomous robotic systems appears to be natural solution. This paper provides results obtained in two different mines, both located in Pennsylvania, USA. One of these mines is a research mine, accessible to people. Another is a former deep mine turned into a strip mine, inaccessible to people but accessible to robotic vehicles.

II. THE ROBOT SYSTEMS

Figure 1 shows the two robotic systems used in our research. On the left is a cart, equipped with four 2-D laser range finders. The laser range finders provide information about the mine cross section ahead of the vehicle, and the ground and ceiling structure. The center panel in Figure 1 shows the Groundhog robot, a tele-operated device constructed from the chassis of two ATVs [2]. The robot is equipped with two 2-D laser range finders, one pointed forward for 2-D mapping and one pointed towards the ceiling for 3-D



Fig. 1. From left to right: Mine mapping cart with four laser range finders, pushed manually through a mine. Groundhog robot used for breaching difficult mine environments. Strip mine in Burgettstown, PA. None of the vehicles provide any odometry information.

mapping. The right panel of this figure shows Groundhog’s descent into an abandoned mine in Burgettstown, PA. Unfortunately, neither of these systems possess odometers or inertial sensors. Thus, the location of the vehicles relative to their points of entry can only be recovered from the range scan data.

III. MINE MAPPING ALGORITHM

A. 2-D Scan Matching

In a first processing stage, our approach applies the scan registration technique described in [14] to recover locally consistent pose estimates, which is reminiscent of prior work in [4], [13], [17]. This algorithm aligns scans by iteratively identifying nearby points in pairs of consecutive range scans, and then calculating the relative displacement and orientation of these scans by minimizing the quadratic distance of these pairs of points. The result of registering scans in this way is a relative displacement and orientation between two consecutive scans:

$$\delta_t = \begin{pmatrix} \Delta x_t & \Delta y_t & \Delta \theta_t \end{pmatrix}^T \quad (1)$$

This relative information makes it possible to recover an estimate of the global coordinates at which a scan was acquired. We will denote such global coordinates by

$$\xi_t = \begin{pmatrix} x_t & y_t & \theta_t \end{pmatrix}^T \quad (2)$$

where x_t and y_t are Cartesian coordinates in 2-D, and θ_t is the robot’s orientation relative to the global coordinate system at time t . The global coordinates are recovered by applying the following recursive estimation equation:

$$\xi_t = f(\xi_{t-1}, \delta_t) \quad (3)$$

with

$$f(\xi_{t-1}, \delta_t) = \begin{pmatrix} x_{t-1} + \Delta x_t \cos \theta_t + \Delta y_t \sin \theta_t \\ y_{t-1} - \Delta x_t \sin \theta_t + \Delta y_t \cos \theta_t \\ \theta_{t-1} + \Delta \theta_t \end{pmatrix} \quad (4)$$

Unfortunately, the pairwise scan registration technique is unable to recover the *global* structure of an environment. This is specifically problematic in environments that contain cyclic structure. Figure 3a shows an occupancy grid

map [11] after executing the local ICP scan registration. While this map is consistent at the local level, it is inconsistent at the global level due to inconsistencies that arise from the accumulation of small errors in the ICP scan matching procedure. The remaining problem is one of correspondence. To acquire globally consistent maps, we need to know the points in time the robot traversed the same mine segment. This problem is generally considered one of the most challenging problems in robotics, and has been addressed by several researchers [5], [6], [13], [21].

B. Building Consistent Maps With Many Cycles

Our approach uses a modified version of the iterative closest point algorithm (ICP) to estimate the correspondence between robot poses at different points in time. To obtain a globally consistent map, our approach iterates a step in which correspondences are identified, and a step in which a path is recovered from the hypothesized correspondences. This iterative optimization procedure is familiar from the literature on ICP [4], the expectation maximization [9], and the RANSAC [12] algorithm in computer vision (see also [22]); the inner-loop optimization is reminiscent of an approach in [17]. The iteration of both steps leads to a sequence of poses $\xi_t^{[0]}, \xi_t^{[1]}, \dots$ of increasing global consistency.

The initial poses are obtained from the local scan matcher described in the previous section: $\xi_t^{[0]} := \xi_t$. Figure 2a shows the sequence of poses, subsampled in five-meters intervals for computational efficiency. In a first step, possible correspondences are identified. Our algorithm identifies pairs of poses $\xi_{s_i}^{[n]}$ and $\xi_{t_i}^{[n]}$, indexed by s_i and t_i , which fulfill multiple criteria: they have to be nearby; they have to lie on approximately parallel path segments; and the line connecting them has to be approximately orthogonal to their respective paths. Figure 2b shows the pose pairs identified by our algorithm in the first iteration.

Next, a new set of poses is calculated that matches these correspondences. To calculate such poses in closed form, our approach transforms the relative pose information δ_t into quadratic constraints between adjacent poses. More specifically, our approach applies the following Taylor expansion

$$\xi_t \approx \hat{\xi}_t^{[n]} + F_{t-1}^{[n]} (\xi_{t-1} - \hat{\xi}_{t-1}^{[n]}) \quad (5)$$




Fig. 2. Global correspondence: (a) Path of the robot, with a node placed every five meters; (b) initial set of correspondences; (c) path obtained under these correspondences; (d) new set of correspondences obtained using the new path; (e) optimal path under these new correspondences; (f) final path and correspondences after three full iterations of the algorithm.



Fig. 3. (a) Map of the coal mine in Bruceton, PA, based on scan matching. This map is obtained by incremental scan matching, and the resulting poses form the input to our loop closing algorithm. (b) Map obtained using our loop closing routines. This map measures approximately 250 by 200 meters in size and contains three large loops. (c) 2-D map of a mine in Burgettstown, PA.

Here $\hat{\xi}_t^{[n]}$ denotes the pose “prediction” $f(\xi_{t-1}^{[n]}, \delta_t)$, and $F_{t-1}^{[n]}$ is the tangent to the function f at $\xi_{t-1}^{[n]}$:

$$\begin{aligned}
 F_{t-1}^{[n]} &= \frac{\partial f(\xi_{t-1}^{[n]}, \delta_t)}{\partial \xi_{t-1}^{[n]}} \\
 &= \begin{pmatrix} 1 & 0 & -\Delta x_t \sin \theta_t^{[n]} + \Delta y_t \cos \theta_t^{[n]} \\ 0 & 1 & -\Delta x_t \cos \theta_t^{[n]} - \Delta y_t \sin \theta_t^{[n]} \\ 0 & 0 & 1 \end{pmatrix}
 \end{aligned} \quad (6)$$

Both $\hat{\xi}_t^{[n]}$ and $F_{t-1}^{[n]}$ are constants in the optimization to follow. The goal of the optimization is to identify poses ξ_t that minimize the quadratic distance to the approximation in (5). This is achieved by minimizing the following quadratic error:

$$\sum_t (\xi_t - \hat{\xi}_t^{[n]} - F_{t-1}^{[n]} (\xi_{t-1} - \xi_{t-1}^{[n]}))^T H_t (\xi_t - \hat{\xi}_t^{[n]} - F_{t-1}^{[n]} (\xi_{t-1} - \xi_{t-1}^{[n]})) \quad (7)$$

Turning the exact calculation in (4) into an optimization problem enables us to “bend” the path of the robot. The matrix H_t measures the penalty associated with bending the path. Mathematically, H_t characterizes the negative log-likelihood of a Gaussian noise model of the ICP scan matcher. Ideally, H_t should be extracted by analyzing the curvature of the ICP target function under translation and rotation. In our software, we simply use a fixed diagonal matrix for the penalty H_t .

The correspondences are incorporated into the optimization through an additional quadratic penalty function. Each pair (t_i, s_i) in the set of pairwise correspondences is mapped into a quadratic cost function of the type:

$$(\xi_{t_i} - \xi_{s_i})^T Z (\xi_{t_i} - \xi_{s_i}) \quad (8)$$

Here Z is a diagonal penalty matrix. Technically, our approach does not enforce $\xi_{t_i} = \xi_{s_i}$; instead, it minimizes the quadratic distance between these poses, with the penalty z . The total cost function of incorporating all correspondences is given by

$$\sum_i (\xi_{t_i} - \xi_{s_i})^T Z (\xi_{t_i} - \xi_{s_i}) \quad (9)$$

To solve the coupled quadratic optimization problem, we now conveniently reorder the terms in (7) and (9). All terms linear in ξ_t in (7) are collected in a large matrix $A^{[n]}$, and all remaining constants into the vector $c^{[n]}$. Similarly, all linear terms in (9) are subsumed in a matrix $B^{[n]}$. The sum of (7) and (9) is then of the following quadratic form:

$$J^{[n]} = (A^{[n]}\xi - c^{[n]})^T \mathcal{H} (A^{[n]}\xi - c^{[n]}) + (B^{[n]}\xi)^T \mathcal{Z} B^{[n]}\xi$$

Here $\xi = \xi_1, \xi_2, \dots$ is the vector of all poses, and \mathcal{H} and \mathcal{Z} are high-dimensional versions of H and Z , respectively. Minimizing this quadratic expression is now straightforward. In particular, we calculate its first derivative

$$\begin{aligned} \frac{\partial J}{\partial \xi_t} &= A^{[n]T} \mathcal{H} (A^{[n]}\xi - c^{[n]}) + B^{[n]T} \mathcal{Z} B^{[n]}\xi \\ &= \left[A^{[n]T} \mathcal{H} A^{[n]} + B^{[n]T} \mathcal{Z} B^{[n]} \right] \xi \\ &\quad - A^{[n]T} \mathcal{H} c^{[n]} \end{aligned} \quad (10)$$

Setting this expression to zero gives us the new set of poses $\xi^{[n+1]}$:

$$\xi^{[n+1]} = \left[A^{[n]T} \mathcal{H} A^{[n]} + B^{[n]T} \mathcal{Z} B^{[n]} \right]^{-1} A^{[n]T} \mathcal{H} c^{[n]}$$

This calculation involves multiplying and inverting matrices whose dimensions are linear in the number of robot poses. These matrices are sparse; however, they can still be humongous. Our software therefore subsamples the set of all poses: As indicated in Figure 2, only a single pose is included for every five meters of robot motion. In the specific data set shown in Figure 2, this reduces the number of pose variables from 13,116 to 381, a dimension that is easily handled

by efficient linear algebra libraries. Adjusted poses for those poses not included in the optimization are easily recovered through linear interpolation. Finally, we note that the linearization is only an approximation, and multiple iterations of the minimization may be required. In our experiments, we always obtained good results in the first two iterations of the optimization.

Figure 2c shows the resulting alignment for the previously calculated correspondences. While the path is now globally consistent in the area where correspondences were identified, it is still inconsistent in other areas. Iterating the basic algorithm leads to the remaining panels in Figure 2. As is easily seen, our approach succeeds in recovering a globally consistent map. The algorithm converges when the correspondences are identical to the ones estimated in the previous iteration.

C. Globally Consistent 2-D Occupancy Maps

Based on the pose estimates obtained in the previous step, our approach extracts an occupancy grid map from the results of the path alignment. It does so by applying once again the scan matching algorithm used to establish the initial relative pose estimates [14], but this time using the poses $\xi_t^{[n]}$ obtained in the global alignment step as an additional constraint. As above, this constraint is represented by a quadratic penalty function, which is easily incorporated into the classical scan matching algorithm (which also optimizes a quadratic function).

Figure 3b shows the map obtained from data acquired in the Bruceton research mine. The map measures 250 by 200 meters in size, and has been constructed without any odometry information.

D. Volumetric 3-D Mapping

In a final step, our approach recovers a 3-D map of the mine. This map is obtained by utilizing the upward pointed 2-D laser and (in the case of the instrumented cart) the downward pointed 2-D laser. Good initial maps are obtained by using the 2-D pose information to construct a 3-D map, via the obvious geometric projections, as described in [16]. Unfortunately, such a reconstruction is only valid for planar environments; in non-planar environments, both volumetric lasers may be tilted, and estimating the tilt is essential for the accuracy of the resulting maps.

Our approach utilizes a forward-pointing vertical laser, presently only available on the robotic cart, which provides a vertical cross-section of the mine as the robot moves. This cross-section enables the robot to register its ceiling and ground scans while simultaneously recovering its pitch (the roll cannot presently be recovered). This estimation is performed using a 3-D variant of the scan matching technique described in [14], using the results of the 2-D pose estimation as a starting point.




Fig. 6. Sequence of 3-D visualizations of a volumetric mine map. Shown in red are the sensor measurements used for generating the mine map.

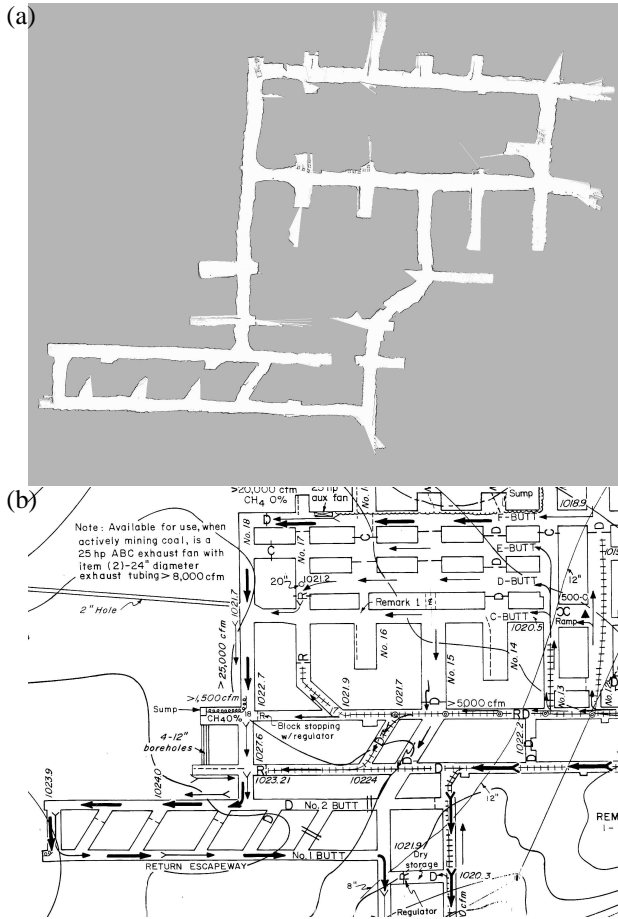


Fig. 4. (a) 2D Mine map acquired by the robotic cart; (b) a hand-drawn map of the corresponding mine segment for comparison.

IV. RESULTS

We obtained all our data in segments of two abandoned coal mines in Pennsylvania. The Bruceton Mine is geographically close to the Quecreek Mine in Somerset County. It is operated as a research mine by the U.S. Bureau of Mines, enabling us to enter robotic equipment without the need for explosion-proof certification. The Burgettstown Mine is an abandoned mine in a dangerously unstable state. Human access is prohibited and the floor of the mine is covered in a thick toxic sludge known as “yellow boy.” The entrance to this mine was discovered only days before the




Fig. 5. (a) A section of the 3-D map produced using the mine-cart and our 2-D mapping algorithms; (b) similar section, using 3-D scan matching for post-processing the mine maps.

robotic mapping expedition, at which point it was fully submerged. In preparation for the robot mapping experiment, water was pumped out of the mine. Mines of this type pose threats to people due to the low oxygen levels and the danger of collapse.

We already discussed example 2-D maps acquired in both mines, and shown in Figure 3. A second map is shown in Figure 4, along with a hand-drawn map of the corresponding mine section. This map was acquired by the robotic cart. Figure 6 shows example views of a 3-D volumetric maps obtained using this system. The lower hemispheres of the maps are missing because our robot has no downward pointing laser. Views of a full 3-D map acquired by the robotic cart is shown in Figure 5. Here we illustrate the effect of the final scan registration step in the full 3-D model—a step




Fig. 7. A section of the 3-D map produced by the Groundhog robot in the Burgettstown mine. This robot possesses no downward pointed laser; hence the map only shows the ceiling and upper side walls of the mine.

that requires a total of four laser range finders. From this map, the total volume of the mine is easily calculated; information that is typically not available from existing mine maps.

Maps of the Burgettstown mine are shown in Figures 3c and 7. These maps are much smaller than those of the Bruce-ton mine. However, their significance lies in the fact that they have been acquired in an environment inaccessible to people. The entrance of the mine is shown in Figure 1c. Figure 7 shows a view of a 3-D map acquired by our Groundhog robot. As before, only the upper half of the mine has been mapped, since the robot possesses no downward pointed sensor.

V. CONCLUSION

We have presented systems and algorithms for robotic mapping of underground mines. Both of our systems are equipped with laser range finders to recover ego-motion and to build accurate maps. Our approach relies on 2-D scan matching to recover a locally consistent map, and on a 2-D global alignment algorithm for generating globally consistent maps. The resulting maps and robot paths form the basis for integrating the 3-D information, acquired by additional scanners pointed at the ceiling and the floor of a mine. A final optimization step further improves the spatial consistency of the resulting 3-D mine map.

While we find that in the mines explored so far, our approach consistently produces accurate maps, the greedy nature of this algorithm makes it possible to get stuck in local minima. Algorithms such as RANSAC [12] are applicable to reduce the danger of getting stuck in a local minimum, at the expense of increased computational complexity.

We believe that existing techniques for mobile robot exploration [7], [19], [23] can be adapted for the purpose of autonomously exploring mines. Such an extension would overcome a crucial limitation of the present approach, namely its reliance on human tele-operation.

ACKNOWLEDGMENTS

This research was motivated and supported by an ongoing subterranean mapping course at CMU. Pictures of the mine-cart were taken by Tom Stepleton. We thank Ben Wegbreit, John Leonard and Michael Bosse and the CMU Mine Mapping Class for their input on our algorithm design. The research has been sponsored by DARPA's MARS Program (contract N66001-01-C-6018 and contract NBCH1020014), which is gratefully acknowledged.

REFERENCES

- [1] T. Bailey. *Mobile Robot Localisation and Mapping in Extensive Outdoor Environments*. PhD thesis, University of Sydney, 2002.
- [2] C. Baker, Z. Omohundro, S. Thayer, W. Whittaker, M. Montemerlo, and S. Thrun. A case study in robotic mapping of abandoned mines. Submitted to FSR'03.
- [3] J.J. Belwood and R.J. Waugh. Bats and mines: Abandoned does not always mean empty. *Bats*, 9(3), 1991.
- [4] P. Besl and N. McKay. A method for registration of 3d shapes. *Trans. Pattern Analysis and Machine Intelligence*, 14(2):239–256, 1992.
- [5] M. Bosse, J. Leonard, and S. Teller. Large-scale CML using a network of multiple local maps. In [15].
- [6] M. Bosse, P. Newman, M. Soika, W. Feiten, J. Leonard, and S. Teller. An Atlas framework for scalable mapping. *ICRA-03*.
- [7] W. Burgard, D. Fox, M. Moors, R. Simmons, and S. Thrun. Collaborative multi-robot exploration. *ICRA-00*.
- [8] P. Corke, J. Cunningham, D. Dekker, , and H. Durrant-Whyte. Autonomous underground vehicles. *CMTE Mining Technology Conf-96*.
- [9] A.P. Dempster, A.N. Laird, and D.B. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society, Series B*, 39(1):1–38, 1977.
- [10] G. Dissanayake, H. Durrant-Whyte, and T. Bailey. A computationally efficient solution to the simultaneous localisation and map building (SLAM) problem. In [15].
- [11] A. Elfes. *Occupancy Grids: A Probabilistic Framework for Robot Perception and Navigation*. PhD thesis, CMU, 1989.
- [12] 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*, 24:381–395, 1981.
- [13] J.-S. Gutmann and K. Konolige. Incremental mapping of large cyclic environments. *CIRA-00*.
- [14] D. Hähnel, D. Schulz, and W. Burgard. Map building with mobile robots in populated environments. *IROS-02*.
- [15] J. Leonard, J.D. Tardós, S. Thrun, and H. Choset, editors. *Notes of the ICRA Workshop on SLAM*, ICRA, 2002.
- [16] Y. Liu, R. Emery, D. Chakrabarti, W. Burgard, and S. Thrun. Using EM to learn 3D models with mobile robots. *ICML-01*.
- [17] F. Lu and E. Milius. Globally consistent range scan alignment for environment mapping. *Autonomous Robots*, 4:333–349, 1997.
- [18] E. Pauley, T. Shumaker, and B. Cole. Preliminary report of investigation: Underground bituminous coal mine, non-injury mine inundation accident (entrapment), July 24, 2002, Quecreek, Pennsylvania, 2002. Black Wolf Coal Company.
- [19] R. Simmons, D. Apfelbaum, W. Burgard, M. Fox, D. an Moors, S. Thrun, and H. Younes. Coordination for multi-robot exploration and mapping. *AAAI-00*.
- [20] R. Smith, M. Self, and P. Cheeseman. Estimating uncertain spatial relationships in robotics. In *Autonomous Robot Vehicles*: 167–193. Springer, 1990.
- [21] S. Thrun. A probabilistic online mapping algorithm for teams of mobile robots. *International Journal of Robotics Research*, 20(5):335–363, 2001.
- [22] B. Triggs, P. McLauchlan, R. Hartley, and A. Fitzgibbon. Bundle adjustment—A modern synthesis. In *Vision Algorithms: Theory and Practice*, Springer, 2000.
- [23] B. Yamauchi, P. Langley, A.C. Schultz, J. Grefenstette, and W. Adams. Magellan: An integrated adaptive architecture for mobile robots. TR 98-2, ISLE, Palo Alto, CA.