On the Representation and Estimation of Spatial Uncertainty

– A Discussion on the Role of Geometry in Robot Maps –

Timothy David Barfoot

University of Toronto Institute for Aerospace Studies

Rev. 202, May 30, 2010

Outline

Puzzling over SLAM

- Imagine doing a jigsaw puzzle where
  - (i) you don't have the picture on the box and,
  - (ii) the pieces only fit together approximately.
- This is a bit like the problem of *simultaneous localization and mapping* (SLAM).

- We want to assemble (localize) the pieces (measurements) into the most likely picture (map) that we can.
- Now ask yourself, how do you solve jigsaws?
- Which strategy works better for very large puzzles?
  - Global: Border first then add pieces to the interior.
  - Local: Grow clusters and eventually join them together.

Framing SLAM

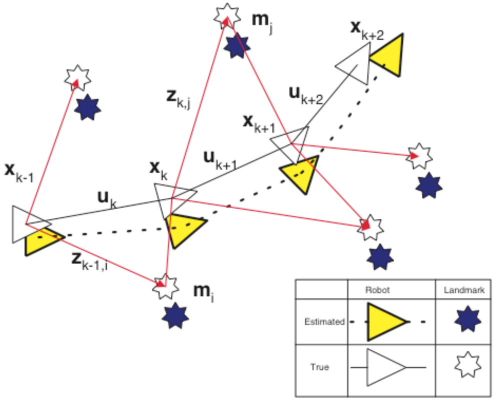

Image: (Durrant-Whyte and Bailey, 2006)

- A robot moves, generating a sequence of frames in the places it's been.
- We receive a sequence of odometric measurements (from frame to frame) and map measurements (frame to landmark).
- We'd like to build a 'map' (frames and landmarks) from the (uncertain) measurements.

But what is a ‘map’?

Brooks (1985) proclaimed,

> *“A representation of the world is not something from which the world need be reconstructible. Rather a representation of the world is a statement of facts deducable from observations, and ideally includes enough facts that anything deducable from past observations is also deducable from the representation. A representation is not an analogous structure to the world; it is a collection of facts about the world.”*

In other words, we should not assume a priori the form of a robot map.

What's so hard about SLAM?

Chatila and Laumond (1985) commented in an early work on mapping that,

> *"The problems addressed in this paper are crucial when dealing with real-world mobile robots as opposed to simulated ones. They are the problems of:*
>
> 1) *constructing and maintaining an accurate enough environment model (of objects and space), that remains consistent as the robot explores new areas or sees again regions that are already modeled, and,*
> 2) *knowing its own position in this environment.*
>
> *The general problem concerns an unstructured environment, and the robot should be able to construct its references by itself. Imprecise mapping is an important issue in current mobile robot research."*

Some 25 years later this last statement still holds true.

Primordial SLAM

- According to Durrant-Whyte and Bailey (2006), the genesis of SLAM occurred during the 1986 IEEE ICRA conference held in San Francisco.
- Many of the soon-to-be SLAM pioneers were present including Peter Cheeseman, Jim Crowley, Hugh Durrant-Whyte, Raja Chatila, Oliver Faugeras, and Randal Smith.
- The next few years saw much progress representing uncertainty in maps (Durrant-Whyte, 1988; Chatila and Laumond, 1985; Crowley, 1989; Smith and Cheeseman, 1986; Ayache and Faugeras, 1988).
- The most influential paper[1] from that time is

  Smith, R. C., Self, M., and Cheeseman, P., "Estimating Uncertain Spatial Relationships in Robotics," in I. J. Cox and G. T. Wilfong, editors, *Autonomous Robot Vehicles*, pages 167-193, Springer Verlag, New York, 1990,

  in which the *stochastic map* framework was introduced.

[1] There was also an 1987 IEEE ICRA paper by the same name but it did not appear in the printed proceedings.

The Stochastic Map

Smith et al. (1990) wrote,

> *"This paper presents a representation that makes explicit the uncertainty of each degree of freedom in the spatial relationships of interest. A method is given for combining uncertain information regardless of which frame it is presented in, and it allows the description of the spatial uncertainty of one frame relative to any other frame."*

but later in the paper wrote,

> *"...our 'map' consists of the current estimate of the mean of the system state vector, which gives the nominal locations of objects in the map with respect to the world reference frame..."*

Seemingly innocuous, the adoption of a **single privileged coordinate frame** in the stochastic map framework was an interesting by-product of this landmark paper on representing uncertainty in robot maps.

Prequel

Interestingly, a few years earlier Smith and Cheeseman (1986) wrote another paper,

Smith, R. C. and Cheeseman, P., "On the Representation and Estimation of Spatial Uncertainty," *The International Journal of Robotics Research*, 5(4):56-68, 1986.

in which they note,

> *"Brooks (1985) argues that it is not appropriate for mobile robots to use a global reference frame. He feels that a set of local reference frames linked via uncertainty transformations is better. We show how the uncertainty of a frame relative to another can be estimated and how the reduction in uncertainty due to sensing can be mapped into any frame, regardless of where the sensing was performed. Because of this flexibility, no particular frame is necessary as an absolute reference."*

This paper was the inspiration for the title of this talk, as it marks a critical decision point in SLAM's history.

Brooks' Dogma

What Brooks (1985) actually wrote was,

> *"The underlying problem is that worse case error needs to be assumed in placing things in an absolute coordinate system, and cumulative worse cases soon lead to useless models globally.*
>
> *We use no global or absolute coordinate system. We do not ignore errors nor do we use beacons or inertial navigation systems. Instead we will use only local coordinate systems with relative transforms and error estimates."*

This statement actually tells the whole story about why the stochastic map framework is not compatible with a single privileged coordinate frame, but we'll return to that a bit later.

Despite these cautions, most SLAM research adopted the single privileged coordinate frame for the next several years.

SLAM Proper

According to Durrant-Whyte and Bailey (2006),

> *"The conceptual break-through came with the realisation that the combined mapping and localisation problem, once formulated as a single estimation problem, was actually convergent. Most importantly, it was recognised that the correlations between landmarks, that most researchers had tried to minimize, were actually the critical part of the problem and that, on the contrary, the more these correlations grew, the better the solution. The structure of the SLAM problem, the convergence result and the coining of the acronym SLAM was first presented in a mobile robotics survey paper presented at the 1995 International Symposium on Robotics Research."*

SLAM SNAFU

- For several years, SLAM research progressed well: EKF-SLAM, SEIF, Rao-Blackwellized Filter (e.g., FastSLAM).
- Until about the year 2000, the focus was on important issues such as data association, environment representation (i.e., to deal with range-only or bearing-only sensors), moving from 2D to 3D, and scaling up to ever-larger environments.
- At this point, scaling up to larger environments brought the computational complexity issue to the forefront of SLAM research. How could these methods be made to work in large-scale environments with loop closures?
- A flurry of new research ensued, focussing on two critical aspects of scaling up SLAM: **consistency** and **complexity**.

Rise of the Submaps

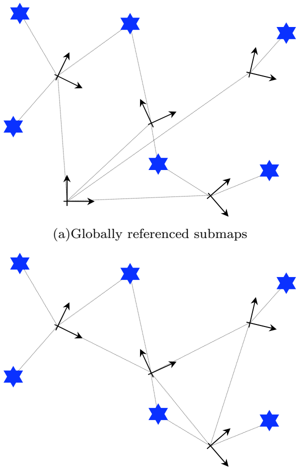

(a)Globally referenced submaps

(b)Locally referenced submaps

Image: (Bailey and Durrant-Whyte, 2006)

- Both global and local submap techniques were introduced to tackle complexity and consistency in SLAM.
- Most of these techniques work in a two-level architecture. A typical stochastic map is used to create each local submap. The submaps are locked and then aligned with respect to one another or a global frame.
- The problem with these techniques is that it is difficult to decide how to break the problem into discrete submaps. However, with some tuning, these techniques enjoyed some limited success.

Submap Prodigy

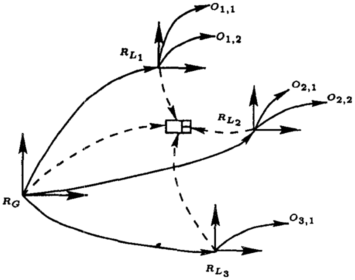

Figure 1: World representation with 3 local frames R_{L_m} related in the global frame R_G.

Image: (Hébert et al., 1996)

- An early version of submapping appeared in 1996:

  Hebert, P., Betge-Brezetz, S., and Chatila, R., "Decoupling odometry and exteroceptive perception in building a global world map of a mobile robot: the use of local maps," in *Proceedings of the IEEE International Conference on Robotics and Automation* (ICRA), volume 1, pages 757-764, 1996.

- This idea was reborn as the 'manifold' later on.

Enter Relative Submaps

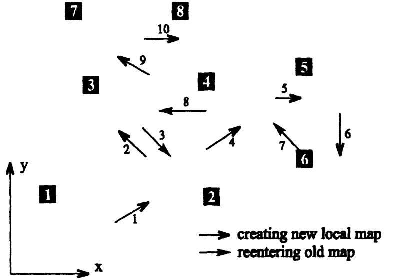

Fig. 7. Representation of a large environment as an inter-connected set of local maps.

Image: (Chong and Kleeman, 1999)

- The earliest example of a true two-tiered submapping approach is

  Chong, K. S. and Kleeman, L., "Feature-Based Mapping in Real, Large Scale Environments Using an Ultrasonic Array," *International Journal of Robotics Research*, 18(1):319, 1999.

- Correlations between features in each local maps were kept as well as correlations between local maps.
- This was further developed by Williams (2001), but due to the decoupling between submaps, neither version is globally convergent.

The Atlas Framework

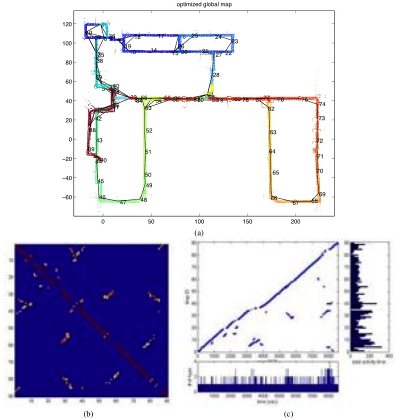

Fig. 16. (a) Global optimized map and Atlas graph for processing of the Killian Court data set using laser scan-matching as the local mapping method (Appendix C). (b) Map adjacency matrix. (c) Map times.

Image: (Bosse et al., 2004)

- A hybrid metric (local)-topological (global) approach was introduced[a] in

  Bosse, M., Newman, P., Leonard, J., and Teller, S., "Simultaneous Localization and Map Building in Large-Scale Cyclic Environments Using the Atlas Framework," *International Journal of Robotics Research*, 23(12):1113-1139, 2004.

- The approach is computationally-efficient, but weakly globally convergent due to its use of the covariance intersection algorithm.

[a] There is also an earlier ICRA 2003 version.

Hierarchical SLAM

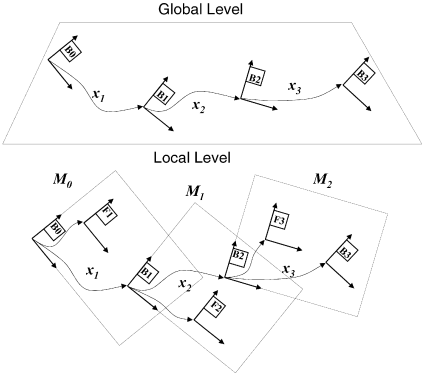

Fig. 1. Two-level hierarchical SLAM model.

Image: (Estrada et al., 2005)

- Hierarchical SLAM is presented in

  Estrada, C., Niera, J., and Tardos, J. D., "Hierarchical SLAM: Real-Time Accurate Mapping of Large Environments," *IEEE Transactions on Robotics*, 21(4):588-596, 2005.

- Local maps are guaranteed to be statistically independent.
- The global map is a **relative stochastic map**.
- The approach explicitly accounts for loop constraints by using SQP at the global level.
- Loops are detected using external data association between local maps.

Drawbacks of Submaps

- Much work has continued since 2005 on submapping methods (see Sibley et al. (2010) for a review).
- According to Sibley et al. (2010), there are several outstanding issues common to many submapping methods:
  - i.e., “map overlap, data duplication, map fusion and breaking, map alignment, optimal sub-map size, and consistent global estimation in a single Euclidean frame”
- Avoiding the need for submaps would appear advantageous. More on this later.

Robots on Manifolds

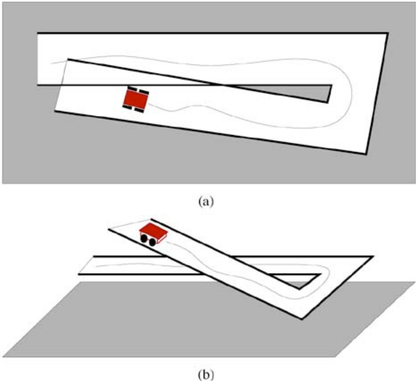

Fig. 1. ***Illustration of a partially closed loop. (a) Planar representation. (b) Manifold representation.***

Image: (Howard et al., 2006)

- In 2004, Andrew Howard gave a standing-room-only talk at ICRA:

  Howard, A., "Multi-robot mapping using manifold representations," in *Proceedings IEEE International Conference on Robotics and Automation (ICRA)*, volume 4, pages 4198-4203, New Orleans, LA, 2004.

- The idea was that by embedding a robot map in a manifold of higher dimension than the usual Euclidean space, it was unnecessary to have global consistency to be useful for robot navigation!
- Howard et al. (2006) provides further details.

Manifold 2D Example

- The manifold idea was used in a teach-and-repeat system for underground mining:

  Marshall, J., Barfoot, T., and Larsson, J., “Autonomous Underground Tramming for Center-Articulated Vehicles,” *Journal of Field Robotics*, 25(6-7):400-421, 2008.

- Now in operation in a mine in Finland!

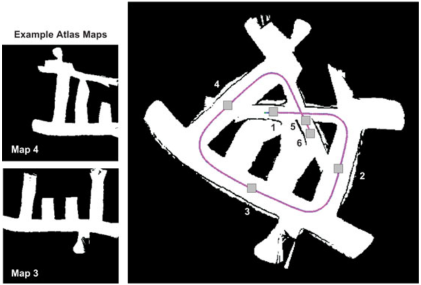

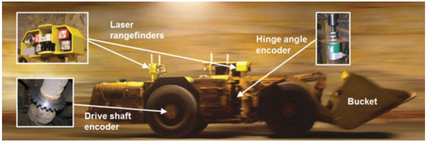

Figure 1. 10-t-capacity Atlas Copco ST1010c LHD, with sensor layout.

Images: (Marshall et al., 2008)

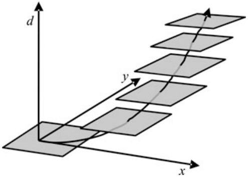

Figure 7. Vehicle path in a manifold described locally by atlas maps.

Manifold 2D Example

(Open in Acrobat Reader to play this movie.)

Manifold 3D Example

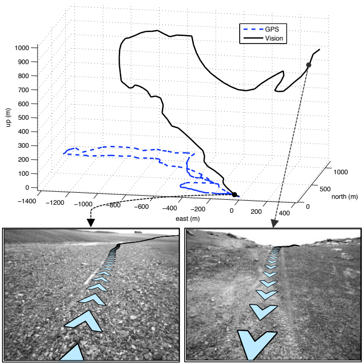

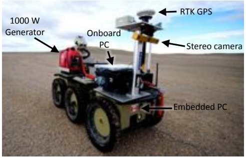

Images: (Furgale and Barfoot, 2010)

- Manifold also used in a teach-and-repeat system for outdoor robots:

  Furgale, P. T. and Barfoot, T. D., “Visual Teach and Repeat for Long-Range Rover Autonomy,” 2010, submitted to the *Journal of Field Robotics*, special issue on “Visual Mapping and Navigation Outdoors”.
- 32 km of testing, 99.6% autonomous!
- ICRA 2010 Kuka Service Robotics Best Paper Award

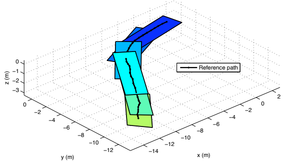

Figure 9: A view of six overlapping submaps with the reference path plotted above.

Manifold 3D Example

(Open in Acrobat Reader to play this movie.)

What do we know up to here?

- Using a single privileged coordinate system does not scale up in terms of both consistency and complexity.
- Manifolds allow useful robot behaviour (even without enforced global consistency) by using relative coordinates. A robot should always retain the ability to reverse its own path.
- Submaps partially alleviate the problems of single privileged coordinate system (at least at the global level); but they still use privileged coordinate systems within each local map and there are several new submap issues introduced.
- **Why not go all the way to a fully-relative framework (i.e., relative stochastic map), and avoid the need for submaps altogether?**
  - No privileged frame, manifold of poses, no submap issues!

Relative Bundle Adjustment

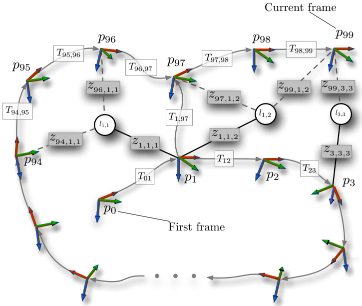

Figure 1: Example trajectory starting at the first frame, p_0, and ending at the current frame, p_{99}. The loop closure between frame 1 and frame 97 adds an extra edge to the graph. Landmark base-frames are indicated with solid lines. Each transform includes an infinitesimal delta-transform defined about $t_j = 0$ — that is, $T_{\alpha j} = \hat{T}_{\alpha j} T_{(t_j)}$, where $\hat{T}_{\alpha j}$ is the current estimate of the relative transform between frame α and frame j.

Image: (Sibley, 2009)

- A fully-relative approach is introduced[a] in:

  Sibley, G., Mei, C., Reid, I., and Newman, P., "Vast Scale Outdoor Navigation Using Adaptive Relative Bundle Adjustment," *International Journal of Robotics Research*, to appear, 2010.

- Current formulation handles exteroceptive measurements (i.e., stereo camera) but not interoceptive measurements (e.g., wheel odometry).

[a] Also an Oxford Tech Report and RSS paper from 2009.

Adaptive Relative Bundle Adjustment

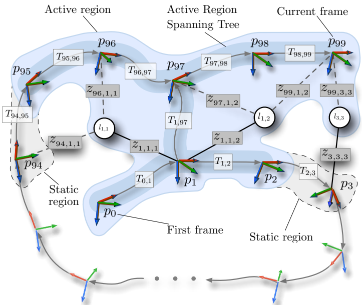

Figure 5: Discovery of local active region. In this example, re-projection errors have changed by more than $\Delta\epsilon$ in the local frames p_{95}, p_{96}, p_{97}, p_{98}, p_{99}, p_0, p_1, and p_2. This local active region is discovered by a weighted breadth-first-search starting at the current frame, p_{99}. All landmarks visible from active frames are optimized for. Any non-active frames that have measurements of active landmarks are added to the static region. Measurements from the static region contribute to the objective function, but the associated edges are not solved for (the frames are fixed).

Image: (Sibley, 2009)

- A batch nonlinear least-squares approach is taken to solve for all the relative state variables, making the solution accurate.
- By adaptively selecting which terms are active in the objective function, the system is efficient.
- Loops can be detected and additional constraints added. The resulting pose graph can be used for planning just like a manifold.
- The pose graph can also be re-expressed in a single privileged coordinate frame, but this requires some offline computation.

Real-time Stereo Mapping

(Open in Acrobat Reader to play this movie.)

What's so special about relative coordinates?

- There seem to be several big advantages to using relative coordinates over absolute coordinates:
    1. By encoding states and their uncertainty relatively, the uncertainty on each state variable is never very large and this allows the Gaussian representation of uncertainty to work well.
    2. An adaptive scheme can be used that limits the number of relative variables involved in state updates.
    3. We inherently have a manifold (a.k.a., pose graph) that can be used for robot navigation at any time.
- Using relative coordinates for *every* pose also avoids the issues associated with submaps.
- Using relative coordinates makes SLAM an 'observable' estimation problem (i.e., using absolute coordinates it is unobservable).

Relative coordinates keep uncertainty small

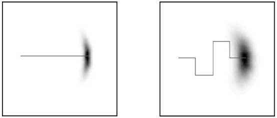

Fig. 1. The density $p(x' \mid x, a)$ after moving 40 meter (left diagram) and 80 meter (right diagram). The darker a pose, the more likely it is.

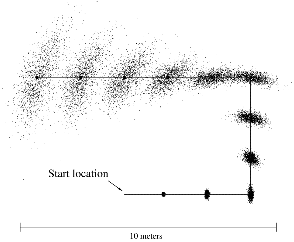

Fig. 2. Sampling-based approximation of the position belief for a robot that only measures odometry. The solid line displays the actions, and the samples represent the robot's belief at different points in time.

Images: (Thrun et al., 2001)

- In SLAM, the further a robot travels from the start, the larger the uncertainty in the pose/map.
- The Gaussian representation of uncertainty used in many estimation paradigms becomes terrible when the states are absolute.
- Relative states keep the uncertainties small and better approximated by a Gaussian; this allows batch nonlinear least-squares to work better.
- Another way to think of it is that we are avoiding the compounding of linearization errors by keeping states local.

Recall Brooks' Dogma

Brooks (1985) was prophetic in saying,

> *"The underlying problem is that worse case error needs to be assumed in placing things in an absolute coordinate system, and cumulative worse cases soon lead to useless models globally.*
>
> *We use no global or absolute coordinate system. We do not ignore errors nor do we use beacons or inertial navigation systems. Instead we will use only local coordinate systems with relative transforms and error estimates."*

Gaussian error modeling does not account for the worst case, but moving to relative state variables may allow us to get away with this assumption.

What does the future hold?

Some predictions:

- A full relative stochastic map framework that handles interoceptive and exteroceptive measurement types.
- A means to incorporate global measurements in the relative framework when available.
- Real-time implementations.
- Loop closure mechanisms.
- ⋮
- All of Brooks' Dogma will turn out to be accurate ;-)

Background Info

- For some background information on state estimation and datasets visit

  http://asrl.utias.utoronto.ca/∼tdb/aer1513/

  Email tim.barfoot@utoronto.ca for password.

Image: robot in a forest of landmarks

Questions?

Image: robot measuring ice thickness with GPR on Ottawa canal (circa 2003)

References I

Ayache, N. and Faugeras, O., “Building, Registering, and Fusing Noisy Visual Maps,” *International Journal of Robotics Research*, 7(6):45–65, 1988.

Bailey, T. and Durrant-Whyte, H., “Simultaneous Localisation and Mapping (SLAM): Part II State of the Art,” *IEEE Robotics and Automation Magazine*, 13(3):108–117, 2006.

Bosse, M., Newman, P., Leonard, J., and Teller, S., “Simultaneous Localization and Map Building in Large-Scale Cyclic Environments Using the Atlas Framework,” *International Journal of Robotics Research*, 23(12):1113–1139, 2004.

Brooks, R., “Visual Map Making for a Mobile Robot,” in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, volume 2, pages 824–829, 1985.

Chatila, R. and Laumond, J.-P., “Position Referencing and Consistent World Modeling for Mobile Robots,” in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, volume 2, pages 138–145, 1985.

Chong, K. S. and Kleeman, L., “Feature-Based Mapping in Real, Large Scale Environments Using an Ultrasonic Array,” *International Journal of Robotics Research*, 18(1):3–19, 1999.

Crowley, J. L., “World Modeling and Position Estimation for a Mobile Robot using Ultra-Sonic Ranging,” in *Proceedings IEEE International Conference on Robotics and Automation (ICRA)*, volume 2, pages 674–680, 1989.

Durrant-Whyte, H. and Bailey, T., “Simultaneous Localisation and Mapping (SLAM): Part I The Essential Algorithms,” *IEEE Robotics and Automation Magazine*, 11(3):99–110, 2006.

References II

Durrant-Whyte, H. F., "Uncertain Geometry in Robotics," *IEEE Journal of Robotics and Automation*, 4(1):23–31, 1988.

Estrada, C., Niera, J., and Tardos, J. D., "Hierarchical SLAM: Real-Time Accurate Mapping of Large Environments," *IEEE Transactions on Robotics*, 21(4):588–596, 2005.

Furgale, P. T. and Barfoot, T. D., "Visual Teach and Repeat for Long-Range Rover Autonomy," 2010, submitted to the *Journal of Field Robotics*, special issue on "Visual Mapping and Navigation Outdoors" on September 30, 2009. Manuscript # ROB-09-0081.

Hébert, P., Betgé-Brezetz, S., and Chatila, R., "Decoupling odometry and exteroceptive perception in building a global world map of a mobile robot: the use of local maps," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, volume 1, pages 757–764, 1996.

Howard, A., "Multi-robot mapping using manifold representations," in *Proceedings IEEE International Conference on Robotics and Automation (ICRA)*, volume 4, pages 4198–4203, New Orleans, LA, 2004.

Howard, A., Sukhatme, G. S., and Matarić, M. J., "Multirobot Simultaneous Localization and Mapping Using Manifold Representations," *Proceedings of the IEEE*, 94(7):1360–1369, 2006.

Marshall, J., Barfoot, T., and Larsson, J., "Autonomous Underground Tramming for Center-Articulated Vehicles," *Journal of Field Robotics*, 25(6-7):400–421, 2008.

Sibley, G., "Relative Bundle Adjustment," Technical Report 2307/09, University of Oxford, Robotics Research Group, 2009.

References III

Sibley, G., Mei, C., Reid, I., and Newman, P., “Vast Scale Outdoor Navigation Using Adaptive Relative Bundle Adjustment,” *International Journal of Robotics Research*, to appear, 2010.

Smith, R. C. and Cheeseman, P., “On the Representation and Estimation of Spatial Uncertainty,” *The International Journal of Robotics Research*, 5(4):56–68, 1986.

Smith, R. C., Self, M., and Cheeseman, P., “Estimating Uncertain Spatial Relationships in Robotics,” in I. J. Cox and G. T. Wilfong, editors, *Autonomous Robot Vehicles*, pages 167–193, Springer Verlag, New York, 1990.

Thrun, S., “Learning Metric-Topological Maps for Indoor Mobile Robot Navigation,” *Artificial Intelligence*, 99(1):21–71, 1998.

Thrun, S., “Robotic Mapping: A Survey,” in G. Lakemeyer and B. Nebel, editors, *Exploring Artificial Intelligence in the New Millenium*, Morgan Kaufmann Publishers, San Francisco, CA, 2003.

Thrun, S., Fox, D., Burgard, W., and Dellaert, F., “Robust Monte Carlo localization for mobile robots,” *Artificial Intelligence*, 128(1-2):99–141, 2001.

Williams, S. B., *Efficient Solutions to Autonomous Mapping and Navigation Problems*, Ph.D. thesis, University of Sydney, Austrialian Centre for Field Robotics, 2001.