Session Overview
Simultaneous Localisation and Mapping

Paul Newman[1] and Henrik I. Christensen[2]

[1] University of Oxford pnewman@robots.ox.ac.uk
[2] Royal Institute of Technology, Sweden hic@nada.kth.se

1 Introduction

The Simultaneous Localisation and Mapping (SLAM) problem remains a prominent area of research in the mobile robotics community. The ISRR symposia have borne witness to marked progress of the field since its conception almost 20 years ago. This year, once again, the question "is the SLAM problem now solved?" was posed. Well the answer to that question probably lies in the definition of "solved". We still do not have the persistent SLAM-enabled machines that we strive for, so in that sense, perhaps it isn't solved, but we do have a firm understanding of the problem now. We do appreciate the limits of performance, we can handle uncertainties in a principled way and recognize the penalties for failing to do so. We also have several solutions to the scaling problem that so dogged the field for several years. To these probabilistic frameworks we are able to attach any of several representational schemes to represent both maps and vehicle trajectories. We run these "solutions" on vehicles equipped with various sensors, cameras, radars, sonars and of course the ubiquitous laser range finder.

One crucial missing component is that of operational robustness. Broadly, the issue can be split into two categories: firstly robustness in the face of erroneous manipulation and insufficient representations of the underlying pdfs and secondly robustness in presence of perceptual ambiguity. The later problem is receiving substantial attention under the guise of the "data association" and "loop closing" problems within the SLAM context. Failing to obtain persistent, long-term SLAM deployments because of accumulating errors in pdf representations is, of course, a closely related problem (bad data association can be caused by incorrect probabilistic representations). A common, although not blanket, criticism of contemporary SLAM techniques is their lack of introspection, they tend to be passive both in data acquisition and data processing. There seems to be a significant scope for planning, acting, and perceiving to aid the SLAM estimation process itself and be more pro-active in assessing the quality of the estimation results.

Perhaps the greatest challenges to contemporary SLAM techniques become clear when trying to apply them in the great outdoors. The benign, distinct surfaces of the flat indoor domain are no more, the world is now truly 3D and single-plane laser scanners are inadequate. The local scene is frequently orders of magnitude larger and may need multiple sensor modalities to access it - cameras, radar, 3D laser and in the underwater domain, beam-steerable sonars. Then there is the issue of performing SLAM in highly dynamic environments that outdoor settings typically demand. The overwhelming majority of SLAM research has relied upon the static world assumption - with varying but typically small degrees of tolerance to scene dynamics. This begs the question how should a principled SLAM system cope with substantial and unexpected scene changes - how can it differentiate this from a catastrophic estimation failure?

2 Summary of papers presented at ISRR

The paper by Bowling et al. addresses the problem of localisation without an *a-priori* choice of representation or specification of process and observation models. The paper hinges on the concept of Action Respecting Embedding a technique similar to Local Linear Embedding, that learns a low dimension manifold within a high dimensional measurement input space. Crucially this operation preserves the local topology originally present when the measurement sequence was gathered. While not addressing the SLAM problem in a familiar way, the paper does illustrate the opportunities that techniques being established in the machine learning domain offer the SLAM research community.

The paper by Wang et al. is a presentation of decoupling in SLAM. Traditionally there is a correlation between robot motion and sensory readings which results in a correlation of all data in a SLAM model. The correlation results in an overall complexity of SLAM which is $O(N^2)$, where N is the number of map features. Various approaches to address the scalability problem have presented in the literature, including the C-EKF by Nebot et al [3], FastSLAM by Montemerlo [2] and the Atlas framework by Bosse et al [4]. In this paper it is demonstrated how a careful relative formulation of the problem, combined with the information filter framework allows decoupling of mapping and localisation — providing a SLAM algorithm with good scaling properties that still allows each feature estimate to be improved with each observation.

Another approach which addresses the scaling problem is presented by Walter et al. The paper again uses the information formulation of the SLAM problem and, like the SEIF proposal [1] manages the scaling problem by maintaining an active set of features with substantial correlations to the vehicle. The suggestion here is to use the act of deleting and re-initialising the vehicle states to create and manage this active subset of features in a consistent fashion. The paper analyses the new proposal (ESEIF) and compares it to the SEIF formulation concluding with a side by side comparison of the two algorithms working on two well known data sets.

3 Wrap-up

So it seems that while it is indisputable that the state of the art SLAM has moved on substantially over the past decade there is still interest research going on, much to do and many interesting questions left un-answered. It is not a solved problem but we do know what questions we should be asking.

References

1. S. Thrun, and Y. Liu. Multi-robot SLAM with sparse extended information filters. "Proceedings of the 11th International Symposium of Robotics Research", Dario and Chatila (Eds.), Springer Verlag, Sienna, October 2003.
2. M. Montemerlo, S. Thrun, D. Koller, B. Wegbreit. FastSLAM A Factored Solution to the Simultaneous Localization and Mapping Problem, "Proceedings of the AAAI National Conference on Artificial Intelligence", Edmonton, Canada, 2002.
3. J. Guivant and E. Nebot. Optimization of the Simultaneous Localization and Map Building Algorithm for Real Time Implementation. *IEEE Transactions on Robotic and Automation*, vol. 17, No. 3, pp. 242-257, June 2001.
4. M. Bosse and P. Newman and J.J. Leonard and S. Teller. SLAM in Large-scale Cyclic Environments using the Atlas Framework. *International Journal of Robotics Research*. vol. 23, pp. 1113-1139, 2004