Exchanging Uncertain Multi-Level Maps

Tom Cooper and Simon Maskell
QinetiQ, St Andrews Road, Malvern, Worcestershire

Abstract

This paper is associated with the SEN016 IF study, 'Exchanging Multi-Level Maps'. This study has successfully built on QinetiQ's Bayesian Modelling Toolkit (BMT) and has applied the BMT to develop a proof-of-principle that shows it is possible to exchange three maps, represent uncertain knowledge of a small 2d multi-resolution world and consider multiple sensors' contribution to the derived atlas. A plan for future research to extend this proof-of-principle to a demonstration with 3D maps is proposed.

Keywords : Bayesian Networks, Robots, Uncertain multi-level maps, exchanging.

Introduction

We are addressing the problem of how uncertain and incomplete maps (eg output by a SLAM algorithm or provided by Ordinance Survey), with different resolutions can be exchanged between robots. To achieve this capability, a system needs the ability to reconcile various differences between maps. In order to demonstrate a proof-of-principle of such a system, we consider 'toy' worlds which allow experiments to be performed without requiring computer algorithms to be optimised. This proof-of-principle involves:

- Exchange of maps: two maps input, one map output
- Exchange of maps at different resolutions and using different sensors
- A small 2-dimensional world

This paper proceeds with a review of the problem, some discussion of the solution, a presentation of results and ends with conclusions and recommendations for future work.

The Problem

The toy world we consider is two-dimensional, and consists of a 5x5 array of pixels. Each pixel can take one of three values: brick, glass and air. Figure 1 shows the world.

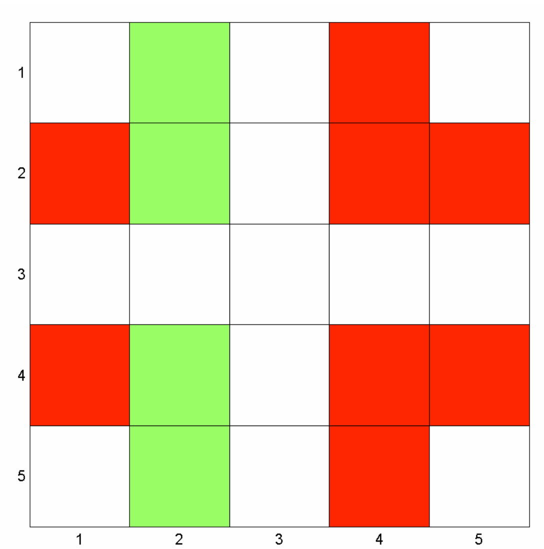

Figure 1: The ternary toy world (red is brick, white is air, green is glass)

Robot A maps part of the area and has a sensor that cannot tell the difference between brick and glass. Figure 2 shows the (perfect) map robot A makes of the ternary world. Robot B maps a different part of the area and has a sensor that cannot tell the difference between air and glass. Figure 3 shows the (perfect) map that robot B makes of the ternary world. The generalisation of the problem to one that involves inputs at multiple scales is assumed to be evident to the reader.

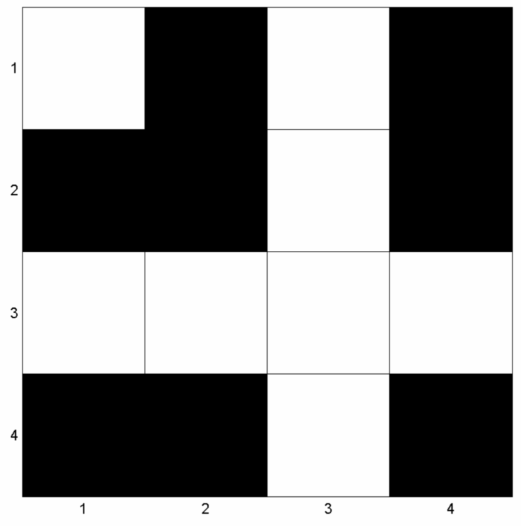

Figure 2: Robot A's map of the ternary world

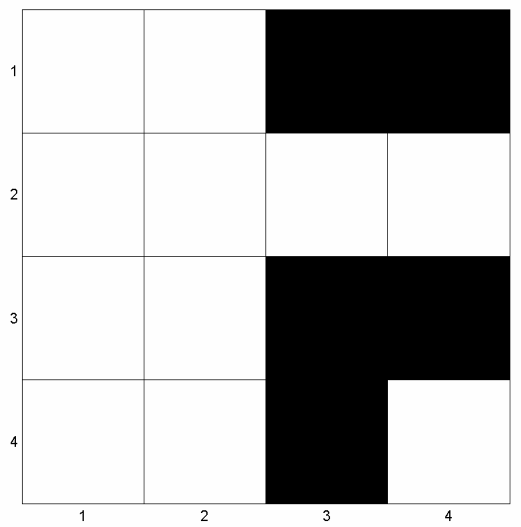

Figure 3: Robot B's map of the ternary world

In the maps we have shown, the robots have made no errors in the values of pixels that they have measured. In the real world, such errors are inevitable, and a map repository should behave robustly when it encounters them.

The problem we are tackling is how to develop a repository architecture that is able import and export maps from robots with uncertain transformations and errors relating the maps with the world, at various resolutions, and for different sensor modalities.

Solution

A Bayesian network [1,2] is a method of encoding the dependencies between random variables. Roughly speaking, nodes of the network correspond to random variables, and arrows between the nodes show that the variables are related.

Our approach is to use a Bayesian network to describe:

- the world,
- maps of the world, and
- the relations between them.

Specifically, our approach is to use QinetiQ's Bayesian Modelling Toolkit (BMT), which can be used to describe and analyse generic Bayesian networks. The BMT can therefore be used to perform certain tasks that the repository needs to perform: that is, it can be used to develop algorithms to allow the exchange of multi-level maps.

To represent the world, we propose a model that has sets of nodes, where each set corresponds to the world at a given resolution. The highest resolution level is parameterised with an Ising model. The lower resolution levels use a quad-tree dependence structure, such that lower-resolution nodes are a summary of the higher-resolution nodes beneath. The BMT's graphical representation is exemplified by the 2-layer model shown in Figure 4.

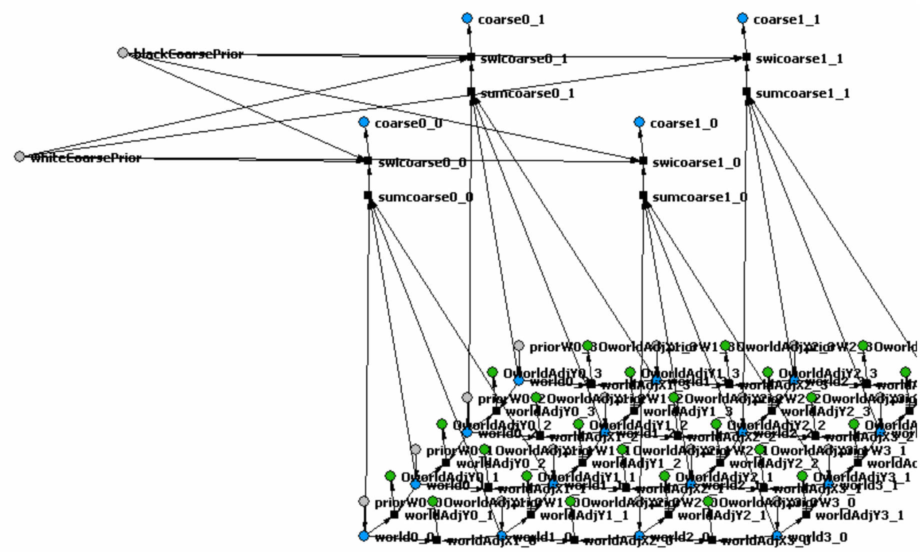

Figure 4: A two-level Bayesian network model of a 5x5 world

We model the ingest of a map through the addition of 'observed' nodes. Each observed node corresponds to a pixel in the map. The relation of the node to the atlas is parameterised by an unobserved transformation (here assumed to be a simple translation) and a stochastic dependence to model sensor errors. Map nodes are connected to world nodes from the level of the atlas of the appropriate resolution. Figure 5 illustrates how a single map connected to the atlas is represented in the BMT.

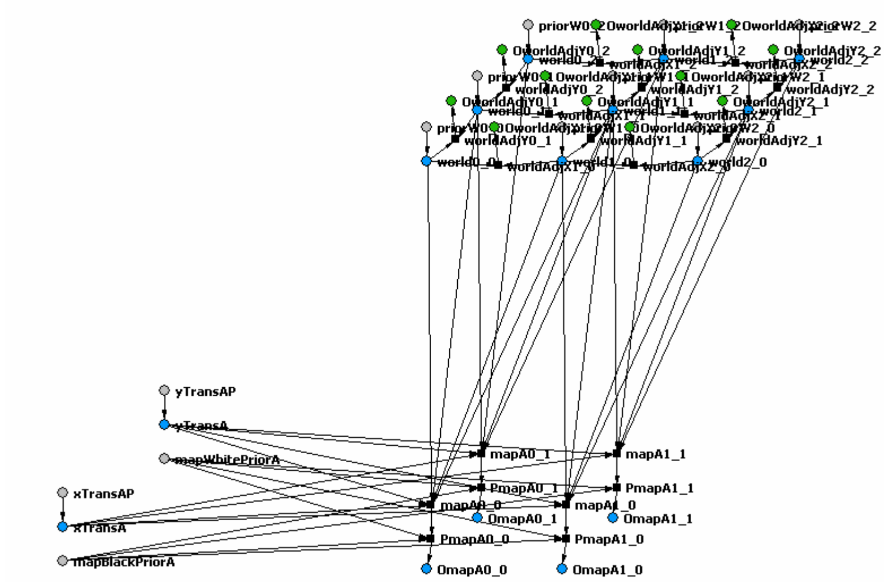

Figure 5: The world connected to a map with an unknown translation

These model definitions enable the solution for the proof-of-principle to be described in the BMT. This is illustrated in Figure 6. This figure shows a two-level atlas of a 5x5 world connected, at the highest resolution level, to two 4x4 maps. This is the model that corresponds to the results that now follow.

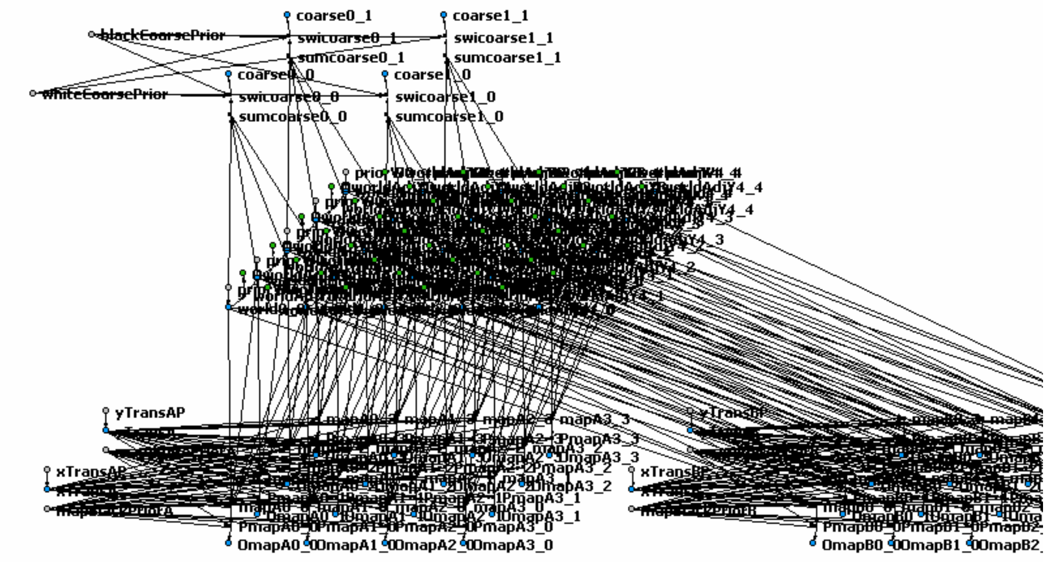

Figure 6: The two-level model of the world connected to two maps

Results

We applied the model shown in Figure 6 to noisy single-sensor versions of the maps shown in Figure 2 and Figure 3. The perception of the noise assumed by the BMT in this proof-of-principle is significant (40% of the pixels in the map are assumed to be incorrectly classified).

The values of the x and y translations for the first map, and the x and y translations for the second map, as calculated by the BMT over 1000 iterations of Gibbs Sampling give the estimated translations shown in Table 1.

Figure 7 shows the mean values (from the same Gibbs sampling process) of the nodes associated with the higher resolution level of the atlas.

Table 1: Estimated Offsets

Robot	P(correct Δx)	P(correct Δy)
A	86%	82%
B	87%	83%

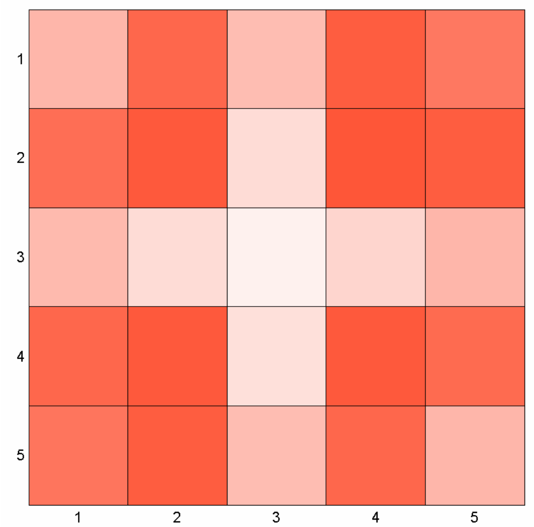

Figure 7: The mean of atlas pixel values

Note that the top-right and bottom-left pixels of the atlas are not observed in either map, but are being estimated solely on the basis of their priors and the spatial dependence structure imposed by the Ising model.

Note also that the atlas is uncertain: the pixels are not all red or white, but are all (differing shades of) pink.

The BMT represents such uncertainty through the parameterisation of a probability density over the joint space of all the unobserved nodes in the model. Bayesian Network algorithms are a mature technique for tackling large non-linear non-Gaussian problems with a computational cost that scales linearly with the number of nodes. However, if the output from the BMT is to provide an input to another subsystem, it is crucial that this interface capitalises on more than just the mean atlas estimate. Exactly how one should interface to and from a multi-level mapping system (particularly when there is the potential for data incest) is motivation for careful system design.

Conclusions

This study has succeeded in developing a proof-of-principle that shows it is possible to exchange multi-level maps (two maps were input, one was output). Uncertain knowledge of a small 2d world was represented and multi-sensor and multi-resolution maps were considered. In particular, the strategy used can model two types of uncertainty: uncertainty of the location of measurements, and uncertainty due to errors in the measurements.

The development of the BMT software, made it possible to develop this proof-of-principle.

The process of conducting an IF study as a precursor to a larger programme has succeeded in refining the scope of such a larger programme.

Recommendations

During this study, a number of specific areas have been identified as being in need of further research. These areas are as follows:

- Algorithms for big maps: The algorithms currently in the BMT are general purpose algorithms and are not designed for multi-level mapping specifically. Techniques that warrant investigtation include integrated local and global optimisation (eg using the simplex algorithm inside an MCMC). This would enable many adjacent nodes to be sampled simultaneously (which cannot be achieved with (non-blocky) Gibbs sampling).
- BMT development: the API to BMT needs to be extended to, for example, enable the functional representation to be sufficiently flexible that it can define multi-level mapping problems succinctly: the BMT manipulates very general functional trees to define non-linearities and probability densities.
- Complex models: The type of the transformations relating the maps and the atlas needs to be refined such that they can model uncertainties that occur in real-world situations. The hierarchical models for the atlas also need to be refined (eg to cater for 3D features). Finally, the sensor models need to reflect real sensor sub-systems' performance.
- Demonstration. The approach needs to be demonstrated on the exchange of multiple 3D maps.

It is recommended that these be researched as part of a future SEN project in the SEAS DTC.

During the IF study, as a result of interaction with Oxford University, the scope of a potential Algorithms and Architectures related research project has been identified. This relates to posing consistent large-scale metric SLAM as sequential non-linear non-Gaussian parameter estimation and is the subject of current discussion with the SEAS DTC.

The study also identified, through discussion with Waterfall Solutions, the potential to use improved data association techniques 0 within the 3D visualisation tools being developed by the DTC.

References

[1] Cowell, R.G., Dawid, A P., Lauritzen, S.L. and Spiegelhalter, D.J., *Probabilistic Networks and Expert Systems*, 1999, Springer.

[2] Jensen, F.V., *Bayesian Networks*, 1996. UCL Press.

[3] Maskell, S., *Signal Processing with Reduced Combinatorial Complexity*. July 2003. Patent reference: 0315349.1.

Acknowledgements

The work reported in this paper was funded by the Systems Engineering for Autonomous Systems (SEAS) Defence Technology Centre established by the UK Ministry of Defence. The authors gratefully acknowledge the guidance and support provided by the theme leaders, Paul Kimber and Neil Short. The discussion with Paul Newman and Mark Bernhardt was also of significant benefit. The authors are also grateful to Chris Booth and Jeremy Baxter for initial discussions regarding this IF study and their previous work for the SEAS DTC. Finally, the BMT team's contribution to the project is very much appreciated. Thanks go to: Kevin Weekes, Richard Everitt, Will Addison and Andrew Edwards.