

The Correspondence Problem in Topological Metric Mapping - Using Absolute Metric Maps to Close Cycles

Margaret E. Jefferies,¹ Michael C. Cosgrove¹, Jesse T. Baker¹,
and Wai-Kiang Yeap²

¹Department of Computer Science
University of Waikato, Hamilton, New Zealand
email: {mjjeff, mcc2, jtb5}@cs.waikato.ac.nz

²Artificial Intelligence Technology Centre
Auckland University of Technology, Auckland, New Zealand
email: wai.yeap@aut.ac.nz

Abstract.

In Simultaneous Localisation and Mapping (SLAM) the correspondence problem, specifically detecting cycles, is one of the most difficult challenges for an autonomous mobile robot. In this paper we show how significant cycles in a topological map can be identified with a companion absolute global metric map. A tight coupling of the basic unit of representation in the two maps is the key to the method. Each local space visited is represented, with its own frame of reference, as a node in the topological map. In the global absolute metric map these local space representations from the topological map are described within a single global frame of reference. The method exploits the overlap which occurs when duplicate representations are computed from different vantage points for the same local space. The representations need not be exactly aligned and can thus tolerate a limited amount of accumulated error. We show how false positive overlaps which are the result of a misaligned map, can be discounted.

1 Introduction

In this paper we describe one of the approaches we are using to solve the corresponding problem in Simultaneous Mapping and Localisation (SLAM). This is regarded as one of the hard problems in SLAM. It is often termed cycle or loop closing because the problem presents itself when the robot traverses a cycle in its environment. The challenge is how to recognise that the cycle has been closed - that parts of the environment observed from different vantage points *correspond* to the same physical space.

The problem is encountered in both topological and absolute metric maps. For absolute metric maps current localisation methods provide consistent enough local maps but residual error accumulates over large distances. By the time a large cycle is encountered the map will contain significant inconsistencies (see Fig.1). Current approaches use some form of probability evaluation to estimate the most likely pose of the robot given its current observations and the current state of its map [1-4]. Detecting the cycle allows the map to be aligned correctly but means the error has to be corrected backwards through the map.

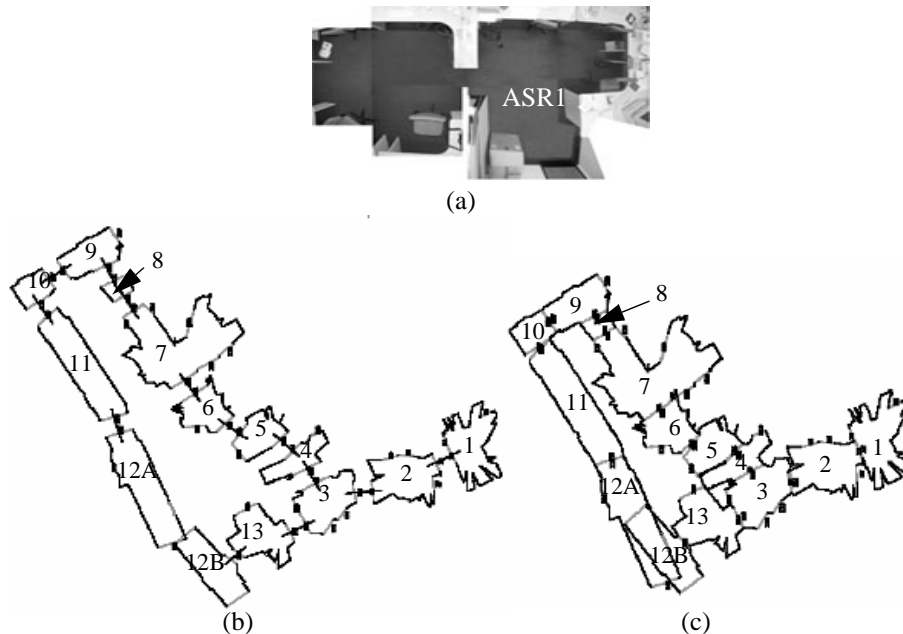


Fig.1 The topological and metric maps. (a) a corner of the robot's environment, a large semi-open laboratory and its surrounding corridor. (b) the topological map (c) the global metric map. The ASRs are numbered in the order they are encountered.

Most topological approaches to robot spatial mapping partition the environment in some way and link these partitions as they are experienced to form a topological map [5-8]. The advantage of this approach is that global consistency is not an issue because the error cannot grow unbounded as in absolute metric maps. Consistency is not a problem within the partitions as they are usually around the size of a local environment. State of the art localisation methods are good enough for local environments. In closing cycles in a topological map the problem is to match two nodes in the topological map if they represent the same physical space (the correspondence problem) and to distinguish two nodes that look the same if they represent different parts of the environment (the perceptual aliasing problem).

Recently hybrid topological/metric approaches have emerged [7, 9, 10] and in [7] the advantages of both the topological and metric mapping paradigms are exploited in closing large cycles. Hybrid approaches are popular in the cognitive mapping community [8, 11-13] however the metric and topological maps do not have equal status. The topological map is the dominant representation in their models. Cognitive maps are often regarded as being like a "map in the head" that an agent (human, animal or robot) has for its experience of its spatial environment. In absolute metric maps the need to match the local map associated with a particular pose and the need to propagate error corrections backwards through the map has seen the introduction of topologically linked local metric maps for sequences of poses [1-3]. However these are a means to an end which is more consistent absolute metric maps.

Our mapping system is based on our previous work where a computational theory of cognitive mapping has been derived from empirical evidence of how humans and animals solve similar problems.[8, 14]. An agent could be human animal or robot. Cognitive mapping researchers have been interested in the correspondence problem for some time but it was not clear from their computer simulations that their algorithms would handle all the uncertainties that a robot faces in the real world [8, 11, 12]. Recently cognitive mapping researchers have begun to adapt their theories and algorithms for the real world problem robots encounter [15-17].

Our approach to mapping the robot's environment extends the hybrid model of [8] and adheres to the dominant cognitive mapping tenet, that the prime representation is the topological map (see [5, 8] for a discussion on why this is so). Yeap and Jefferies' [8] topological map of metric local space descriptions (see Fig.1) has been implemented on a mobile robot with minor adaptations to handle input from a laser range sensor. Yeap and Jefferies [8] proposed a limited (in size) absolute metric map to close small cycles in the topological map. The restricted size of their absolute metric map accounts for the limitations in the human or animal path integration system with accumulating error [18]. The idea is that parts of the map that are distant enough from the agent's current pose will be significantly misaligned with rest of the map due to accumulating error. These would simply drop out of the map. In practice, however, without some error correction the global metric map can detect very few cycles. In the implementation we describe here, using a locally consistent global metric map, we are able to detect significant cycles. Using this method, we use the global metric map to detect and close cycles in the topological map. False positive matches are possible but using the method in conjunction with topological verification we are able to eliminate most false positive matches [17].

2 The Basic Mapping Approach

The topological map comprises a representation for each local space visited with connections to others which have been experienced as neighbours. The local space is defined as the space which "appears" to enclose the robot. The local space representation is referred to as an Absolute Space Representation (ASR) a term which emphasises the separateness and independence of each individual local space. Each ASR in the topological map has its own local coordinate frame. Note that these are *local* absolute spaces in contrast to the *global* absolute metric representations referred to in section 1.

The basic algorithm described in [8] was modified to handle input from a laser range sensor and accumulating odometric and sensor errors. However the fundamentals of the algorithm remain. Yeap and Jefferies [8] argued that the exits should be constructed first because they are the gaps in the boundary which tell the robot how it can leave the current space. An exit will occur where there is an occlusion and is formed by creating the shortest edge which covers the occlusion. Once the exits are formed it is straightforward process to connect the surfaces which lie between them to form the boundary of the ASR. At the same time surfaces which are viewed through the exits, and are thus outside the ASR, are eliminated. Fig.2 (b) shows a sequence of two ASRs so computed. See [8] for an in-depth description of the basic algorithm and [17, 19] for the




Fig.2(a) A section of the robot's environment. (b) The ASRs constructed correspond to the labelled sections of the environment in (a). E1 and E2 are exits, E1 links ASR1 and ASR2.

details of how it is implemented on an autonomous mobile robot using laser range sensing.

Rofer's [20] histogram correlation localisation method is used to provide consistency within ASRs. New ASRs are computed whenever the robot crosses an exit into an unexplored region and ASRs are linked, as they are experienced, via the exits which connect them to their neighbours in the topological map. The ASRs are the nodes of the topological map and the exits are its edges. Fig.1(b) shows the topological map constructed in our large L-shaped open plan laboratory and its surrounding corridor. ASRs 1-8 and ASR 13 comprise the laboratory and the remaining the corridor. Tables and protruding desks provide occlusions where exits are computed. In large open spaces there are fewer occlusions and thus fewer opportunities to partition the space, for example ASR 7 in Fig.1(b).

3 Closing Cycles with a Global Absolute Metric Map

The main advantage of global absolute metric mapping should be that because the robot's location is measured in absolute terms, returning to a previously visited place is clearly apparent by virtue of robot's location within the absolute map. In reality, however, this is not the case - significant misalignment of the map occurs as residual errors accumulate (see Fig.1(c)). However we noted that this misalignment is often not complete, that even though there is significant misalignment in the map, the corresponding ASRs may continue to have substantial overlap. For example, in Fig.1(c) due to the misalignment along the corridor comprising ASRs 11 and 12A one cannot detect immediately from the robot's pose that the robot has re-entered ASR12A from ASR13. However it can be seen that ASR12A overlaps with the ensuing duplicate ASR12B. Note that ASR12B is smaller than ASR12A as the robot has yet to fully explore it. If we maintain the global metric map as a collection of ASRs in a single global coordinate system, we can exploit this overlap to detect that the robot is re-entering a known part of its environment.

The global metric map is discretised into the local space descriptions which correspond to the nodes in the topological map. Whenever the robot crosses an untraversed exit the robot computes a new ASR for its current local environment. It then checks its known ASRs in the global metric map for overlap. We want to detect true overlap, i.e.




Fig.3 (a) The environment (b) An example of where the overlap would not be detected. The centers of each of the overlapping ASRs are not inside the corresponding ASR.

the overlap which is probably not as good as it should be due to the misaligned map rather than the false overlap which results from the map misalignment. To minimise the effect of the latter we match ASR centres. The robot's position is firstly projected to the centre of the current ASR and this location is checked for inclusion in the ASRs in the global map. For example, in Fig.1(c) the robot's position is projected to the centre of ASR12B. This position is checked for inclusion in ASRs 1-12A. This is true for ASR12A. To minimise the effect of the spurious overlaps which are the result of the misalignment we perform a crosscheck of the matching ASRs' centers. In Fig.1(c) we take the centre of ASR12A and check it for inclusion in ASR12B. This eliminates many of the false positive matches at very little cost. The trade-off is that some positive matches will be missed. The method tolerates a significant but limited amount of accumulated error - each of the centers of the duplicate ASRs must lie inside the other. Fig.3(b) shows an example of an overlap which would fail the centers crosscheck.

While the above check discounts many false positive matches, if the accumulated error is significantly large then some false matches may pass this test. The next step in the process is to "close the loop" in the topological map. In the example of Fig.1(c), this means that ASR12A is linked to ASR13. To achieve this "linking of ASRs" the corresponding exits need to be found, in particular the pair belonging to the edge which closes the cycle (see Fig.4). Fortuitously this provides another mechanism for eliminating false positive matches. If the pair of corresponding exits cannot be found the match is rejected.

We do not attempt to combine ASR12A and ASR12B into a single integrated representation. The problem is that even accounting for the fact ASR12B has not been fully explored, there are significant differences in the boundary of ASR12A and ASR12B. Some of this is due to sensing and odometry errors but it is also attributed to the fact that the ASRs are viewed from different vantage points. The same physical space does not look the same when viewed from different locations. Combining the ASRs would provide a neater map. However, from whichever viewpoint the robot encountered the ASR, the map would be a compromise. This is problematic in dynamic environments where discrepancies in the representation of the current view as compared with a previous representation need to be attributed to either map errors or real changes in the environment.

Thus we maintain duplicate representations for the same physical space which correspond to the different vantage points from which they were initially computed. The




Fig.4. The topological map with its cycle closed, i.e. ASR12A is linked to ASR13

links in the topological map which correspond to duplicate ASRs are unidirectional. For example, in Fig.4 when traversing ASR11 to ASR13, ASR12A is used. When traversing ASR3 to ASR11, ASR12B is used.

The main purpose of our approach is to close cycles in the topological map. However with the cycle closed there is the opportunity to realign the global map, correcting the error backwards through the map and develop a model of the residual error to assist future cycle detection. We are currently investigating this aspect of our approach and are comparing it with Yeap and Jefferies [8] limited in size global metric map where the misaligned parts of the map would simply drop off.

We also employ landmark matching to identify and cycles in the topological map [17]. Cycles detected in the topological map provide supporting evidence for cycles detected in the global metric map and vice versa.

4 Conclusion

We have shown that significant cycles in a topological map can be detected from the corresponding cycles in a global metric map. The key to the approach is to ensure that the global metric map is made up of the ASRs in the topological map. The approach is conservative but combined with landmark cycle detection [17] we are able to close many cycles in large-scale environments. However our approach is conservative; we sacrifice some true positive matches so that we can reject most false positive matches. Missing the opportunity to close a cycle in a topological map is not catastrophic as in absolute metric mapping. The outcome is that the robot will take a longer route than it needs to.

References

1. Hahnel, D., Burgard, W., Fox, D., and Thrun, S. A efficient fastSLAM algorithm for generating maps of large-scale cyclic environments from raw laser range measurements. in *Proceedings Intelligent Robots and Systems*. 2003.

2. Thrun, S., Hahnel, D., Ferguson, D., Montemerlo, M., Triebel, R., Burgard, W., Baker, C., Omohundro, Z., Thayer, S., and Whittaker, W. A system for volumetric robotic mapping of abandoned mines. in *Proceedings International Conference on Robotics and Automation*. 2003.
3. Hahnel, D., Thrun, S., Wegbreit, b., and Burgard, W. Towards lazy data association in SLAM. in *Proceedings 10th International Symposium of Robotics Research*. 2003.
4. Gutmann, S. and Konolige, K. Incremental mapping of large cyclic environments. in *Proceedings International Symposium on Computational Intelligence in Robotics and Automation*. 1999.
5. Kuipers, B., The spatial semantic hierarchy. *Artificial Intelligence*, 2000. 119. 191-233.
6. Tomatis, N., Nourbakhsh, I., and Siegwart, R. Simultaneous localization and map building: A global topological model with local metric maps. in *Proceedings International Conference on Intelligent Robots and Systems*. 2001.
7. Bosse, M., Newman, P., Leonard, J., Soika, M., Feiten, W., and Teller, S. An Atlas framework for scalable mapping. in *Proceedings International conference on Robotics and Automation*. 2003.
8. Yeap, W.K. and Jefferies, M.E., Computing a representation of the local environment. *Artificial Intelligence*, 1999. 107. 265-301.
9. Tomatis, N., Nourbakhsh, I., and Siegwart, R. Hybrid simultaneous localization and map building: closing the loop with multi-hypotheses tracking. in *Proceedings IEEE International Conference on Robotics and Automation*. 2002. Washington DC, USA.
10. Thrun, S., Learning metric-topological maps for indoor mobile robot navigation. *Artificial Intelligence*, 1998. 99(1). 21-71.
11. Kuipers, B.J. and Byun, Y.-T. A robust, qualitative method for robot spatial learning. in *Proceedings of the National Conference on Artificial Intelligence (AAAI-88)*. 1988.
12. Yeap, W.K., Towards a computational theory of cognitive maps. *Artificial Intelligence*, 1988. 34. 297-360.
13. Chown, E., Chaplain, S., and Kortenkamp, D., Prototypes, location, and associative networks (PLAN): Towards a unified theory of cognitive mapping. *Cognitive Science*, 1995. 19. 1-51.
14. Jefferies, M.E. and Yeap, W.K. Representing the local space qualitatively in a cognitive map. in *Proceedings Twentieth annual conference of the Cognitive Society*. 1998.
15. Kuipers, B. and Beeson, P. Bootstrap learning for place recognition. in *Proceedings 18th International Conference On Artificial Intelligence*. 2002.
16. Beeson, P., MacMahon, M., Modayil, J., Provost, J., Savelli, F., and Kuipers, B. Exploiting local perceptual models for topological map building. in *Proceedings IJCAI-2003 Workshop on Reasoning with Uncertainty in Robotics*. 2003.
17. Jefferies, M.E., Weng, W., Baker, J.T., Cosgrove, M.C., and Mayo, M. A hybrid approach to finding cycles in hybrid maps. in *Proceedings Australian Conference on Robotics and Automation*. 2003.
18. Gallistel, C.R. and Cramer, A.E., Computations on metric maps in mammals: getting oriented and choosing a multi-destination route. *The Journal of Experimental Biology*, 1996. 199. 211-217.
19. Jefferies, M.E., Yeap, W.K., and Baker, J., Robot mapping with a topological map of local space representations, in *Advances on Simulation, Systems Theory and Systems Engineering*, N.E. Mastorakis, V.V. Kluev, and D. K., Editors. 2002, WSEAS Press. 287-294.
20. Rofer, T. Using histogram correlation to create consistent laser scan maps. in *Proceedings IEEE International Conference on Intelligent Robotics Systems*. 2002.