Journal of Intelligent Manufacturing, 16, 693–702, 2005

Using absolute metric maps to close cycles in a topological map

MARGARET E. JEFFERIES[1,*], WAI-KIANG YEAP[2], MICHAEL C. COSGROVE[1] and JESSE T. BAKER[1]

[1] *Department of Computer Science, University of Waikato, Hamilton, New Zealand*
E-mail: {mjeff, mcc2, jtb5}@cs.waikato.ac.nz
[2] *Institute for Information Technology Research, Auckland University of Technology, Auckland, New Zealand*
E-mail: wai.yeap@aut.ac.nz

Received May 2004 and accepted November 2004

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.

Keywords: ■

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(c)). Current approaches use some form of probability evaluation to estimate the most likely pose (x, y, θ) of the robot given its current observations and the current state of its map (Gutmann and Konolige, 1999; Hähnel *et al.*, 2003a, b; Thrun *et al.*, 2003) (x and y are the robot's location in 2D coordinates and θ is the robot's orientation). Detecting the cycle allows the map to be aligned correctly but means the error has to be corrected backwards through the map.

* Author for correspodence.

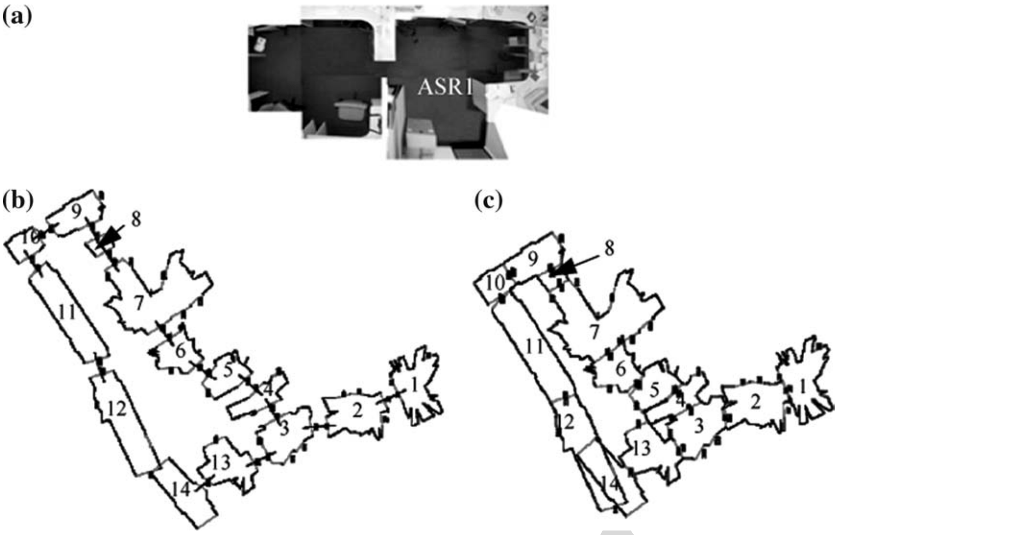

Fig. 1. The topological and metric maps. Note that ASRs 12 and 14 correspond to the same physical space and will be detected as such if they overlap sufficiently (a) a corner of the robot's environment, a large semi-open laboratory and its surrounding corridor. (b) The topological map. (c) The global 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 (Yeap and Jefferies, 1999; Kuipers, 2000; Tomatis *et al.*, 2001; Bosse *et al.*, 2003). 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 (Thrun, 1998; Tomatis *et al.,* 2002; Bosse *et al*., 2003; Thrun *et al*., 2003) and in Bosse *et al*. (2003) 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 (Kuipers and Byun, 1988; Yeap, 1988; Chown *et al*., 1995; Yeap and Jefferies, 1999) 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 (Hähnel *et al*., 2003a, b; Thrun *et al*., 2003). 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 (Jefferies and Yeap, 1988; Yeap and Jefferies, 1999). 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 (Kuipers and Byun, 1988; Yeap, 1988; Yeap and Jefferies, 1999). Recently cognitive mapping researchers have begun to adapt their theories and algorithms for the real world

Journal: **JIMS** MS.: **NO00004372** PIPS: **NO00004372** ☑TYPESET ☐DISK ☐LE ☑CP Disp.: **14/9/2005** Pages: **10**

problem robots encounter (Beeson *et al.*, 2003; Jefferies *et al.*, 2003; Kuipers *et al.*, 2004; Modayil *et al.*, 2004).

Our approach to mapping the robot's environment extends the hybrid model of Yeap and Jefferies (1999) and adheres to the dominant cognitive mapping tenet, that the prime representation is the topological map (see Yeap and Jefferies, 1999; Kuipers, 2000 for a discussion on why this is so). Yeap and Jefferies (1999) topological map of metric local space descriptions (see Fig. 1(b)) has been implemented on a Pioneer 2DX mobile robot with minor adaptations to handle input from a forward facing laser range sensor with a 180° "viewing" angle. Yeap and Jefferies (1999) 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 (Gallistel and Cramer, 1996). 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 could only detect very small 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 (Jefferies *et al.*, 2003).

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. Thus the nodes in the topological map are metric representations of ASRs. The edges are the transitions which take the robot from one local space to another. The global metric map is computed alongside the topological map.

The basic algorithm described in Yeap and Jefferies (1999) 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 (1999) 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.

The raw laser range data from a 180° scan is converted into lines representing surfaces which block the robot's line of sight using a straightforward regression algorithm. The coordinate system for the first ASR is centred on the robot's initial pose (x, y, θ), where x, y, and θ are all set to 0. Initially an occlusion map is constructed (see Fig. 2(b)) which comprises these lines and their occlusions and it is from this map that the exits are constructed. Figure 2(b) shows the exits overlaid on the occlusion map. Exits occur where there is a gap that is large enough for the robot to pass through. In the environment depicted in Fig. 2(a), gaps that will not allow the robot passage under a table or desk often occur between chair legs. Once the exits are formed it is a 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. Parts of the ASR which require further investigation are marked as *unknown*. Figure 2(d) shows the ASR which results. Figure 2(e) shows how the ASR is extended when the unknown regions are investigated. The ASR in Fig. 2(d) indicates the regions which need exploring. The robot moves towards each of these in turn. The new laser range data is incorporated into the occlusion map (Fig. 2(c)) and the ASR recomputed. Note that once the

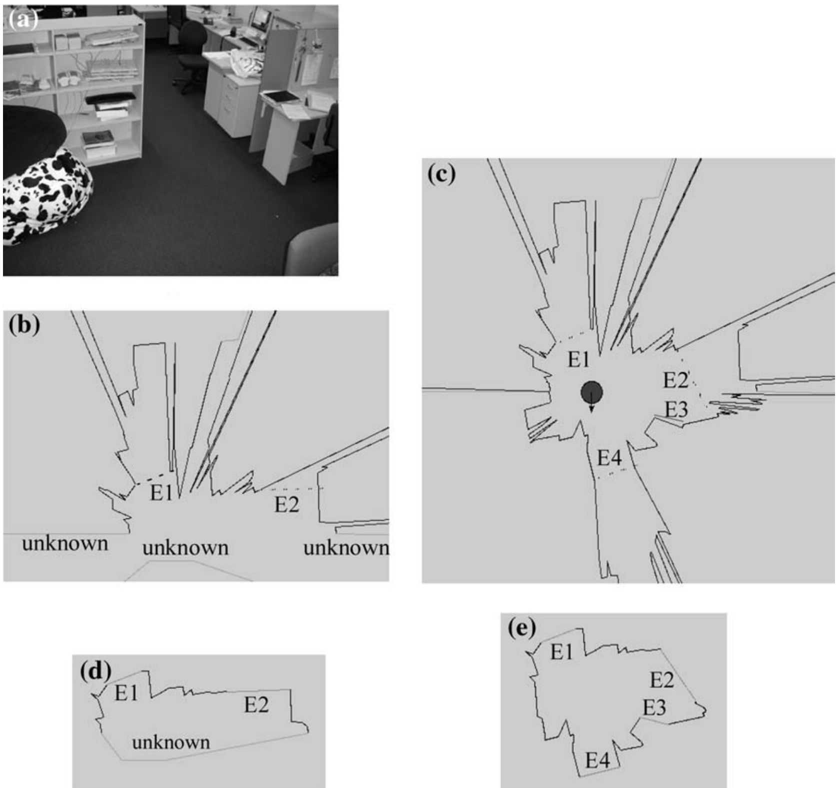

Fig. 2. (a) A section of the robot's environment from which the ASR in (d) and (e) was constructed. (b) The occlusion map extracted from raw laser range data overlaid with the exits E1 and E2. The regions behind the robots line of "sight" are marked as *unknown*. (c) The updated occlusion map when the robot investigates the unknown regions. (d) The temporary ASR constructed from the occlusion map in (b). (e) The final ASR constructed from the occlusion map in (c).

unknown regions are incorporated the ASR could be structured differently. This can be seen in the resulting ASR in Fig. 2(e). In particular, exploring the peripheries of the robot's view often results in ASRs with are structurally different from an initial ASR. The initial ASR merely provides a reasonable guide as to the overall shape of the local space and indicates where the robot should explore to obtain a complete enclosure. Further exploration refines the ASR to a better fitting representation of the robot's local space.

With its first ASR complete, the robot chooses an exit by which to explore the rest of its environment. In our exploration strategy, the robot investigates in a depth first manner, choosing the next largest exit to explore at each step. For the ASR in Fig. 2(e) this is exit, E2. When this exit has been crossed, the robot is in a new local space, and a new local space ASR2 is constructed. The process proceeds as for ASR1 (see Fig. 3(b) and (c)), the coordinate system being centred on the robot's initial pose in the new local space. ASR2 is then connected to ASR1, in the topological map (Fig. 3(d)). The edge connecting the two ASRs indicates that while in ASR1, the transition to ASR2 is via ASR1's exit, E2. From ASR2 the transition to ASR1 is via ASR2's exit, E3. The global map (Fig. 3(e)) comprises both ASR1 and ASR2 in a single frame from of reference centred on the coordinate system of the current ASR, ASR2. See Yeap and Jefferies (1999) for an indepth description of the basic algorithm and (Jefferies *et al*., 2002, 2003) for the details of how it is implemented on an autonomous mobile robot using laser range sensing.

Rofer's (2002) 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

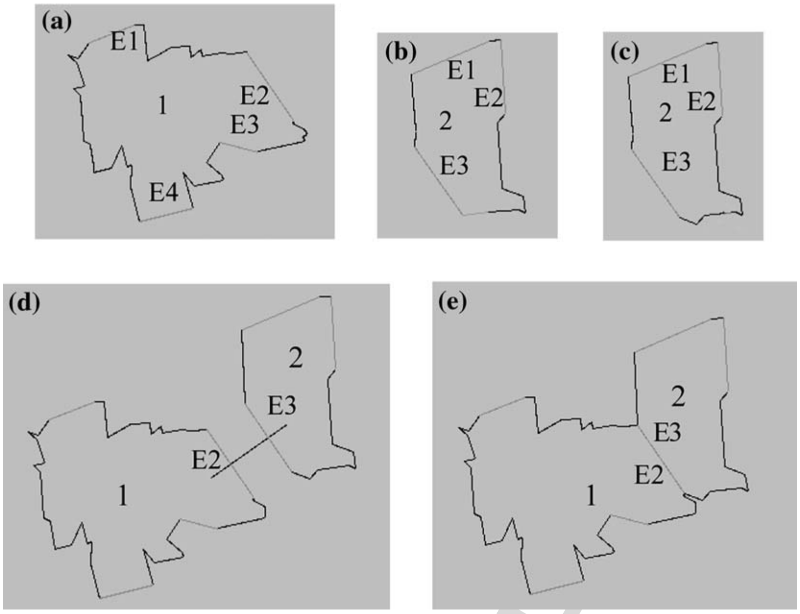

Fig. 3. (a) The first ASR. (b) The initial enclosure for ASR2, computed when the robot crossed E2 in ASR1. (c) The final enclosure for ASR2. (d) ASR1 and ASR2 are connected via the traversed exits to form a topological map. (e) The global metric map.

and ASRs are linked, as they are experienced, via the exits which connect them to their neighbours in the topological map. Figure 1 shows the topological and global maps 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. Doorways, tables and 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.

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 even when 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 12 one cannot detect immediately from the robot's pose that the robot has re-entered ASR12 from ASR13. However, it can be seen that ASR12 overlaps with the ensuing duplicate ASR14. Note that ASR14 is smaller than ASR12 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, matching 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 ASR14. This position is checked for inclusion in ASRs 1–12. This is true for ASR12. 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 ASR12 and check it for inclusion in ASR14. 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 potentially duplicate ASRs must lie inside the other. Figure 5(b) shows an example of an overlap which would fail the centres 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 ASR12 is linked to ASR13. In this example the exits can be aligned and the link made via the corresponding exits (see Fig. 4). We do not attempt to combine ASR12 and ASR14 into a single integrated representation. The problem is that even accounting for the fact ASR14 has not been fully explored, there are significant differences in the boundary of ASR12 and ASR14. Some of this is due to sensing and odometry errors but it can also be 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. If multiple representations are recorded real changes can be tracked over time; the most appropriate ASR can be selected and out of date representations can disappear once it is certain they are no longer relevant.

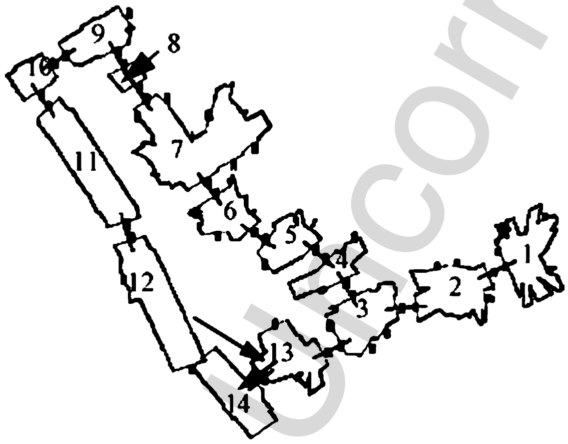

Fig. 4. The topological map with its cycle closed, i.e. ASR12A is linked to ASR13.

Thus we maintain duplicate representations for the same physical space which correspond to the different vantage points from which they were initially computed. The links in the topological map which correspond to duplicate ASRs are currently unidirectional. For example, in Fig. 4 when traversing ASR11 to ASR13, ASR12 is used. When traversing ASR3 to ASR11, ASR14 is used.

Figures 5 and 6 show the mapping of the cycle around the group of tables in our large laboratory. This cycle raises some interesting issues which we are currently investigating. In Fig. 5(a) and (b), the topological and global maps, respectively, the robot is currently in ASR5. Note that the corner of ASR5 overlaps ASR1 in the global map but appropriately this does not signify a match. In Fig. 5 (c) and (d) the robot has moved into ASR6. It can be seen in Fig. 5 (d) that ASR6 is almost entirely contained within a corner of ASR1. However, this match will fail the centre cross match check, i.e. the centre of ASR6 is within ASR1, but not vice versa. This demonstrates the circumspect nature of our approach. In this case a match is appropriate, however, matches such as these are often the result of spurious overlaps due to misalignment errors. Currently we err on the side of caution and reject all such matches. However, inadvertently rejecting a true positive such as in Fig. 5 (d) often means that detecting the cycle is delayed rather than being missed altogether. In Fig. 6 (a) and (b) the robot has entered ASR7. It can be seen clearly in Fig. 6 that ASR7 covers the greater part of ASR1. Cross matching the centres of these two ASRs does indicate that they are representations for the same physical space. Armed with the knowledge we have of this

Journal: JIMS MS.: NO00004372 PIPS: NO00004372 ☑TYPESET ☐DISK ☐LE ☑CP Disp.: 14/9/2005 Pages: 10

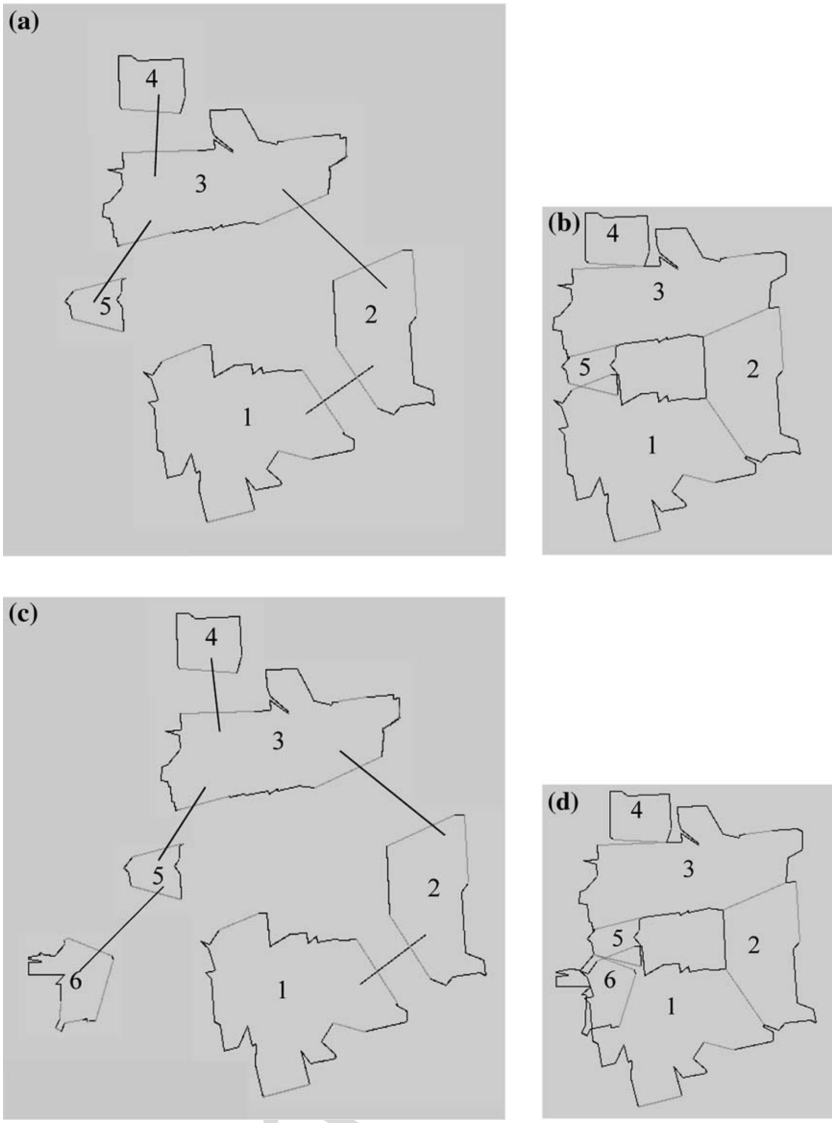

Fig. 5. (a) and (b) The topological and global metric maps, respectively. The robot is currently in ASR5. (c) and (d) The topological and metric maps, respectively. The robot is in ASR6. ASR6 overlaps ASR1 in (d) but fails the centre crosscheck.

match it should be possible to backtrack to the previously rejected match, accepting it in hindsight. We have not implemented this yet. Once a match is indicated it can be verified using our topological verification approach (Jefferies *et al.*, 2003).

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, to correct 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 (1999) 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 close cycles in the topological map (Jefferies *et al.*, 2003, 2004). Cycles detected in the topological map provide supporting evidence for cycles detected in the global metric map and vice versa.

Journal: **JIMS** MS.: **NO00004372** PIPS: **NO00004372** ☑TYPESET ☐DISK ☐LE ☑CP Disp.: **14/9/2005** Pages: **10**

Fig. 6. (a) and (b) The topological and global maps, respectively. The robot is in ASR7. ASR7 clearly covers most of ASR1 in (b). It passes the centre crosscheck and a match is thus indicated.

4. Related Work

We are not aware of any approaches which combine topological and metric mapping in the way that we do. Two approaches which combine topological and global metric mapping and which have some similarity to our work are those of Bosse *et al*. (2003) and Modayil *et al*. (2004). In Bosse *et al*., ATLAS the global metric map seems incidental as a by product of topological mapping. The topological map comprises interconnected local maps, each of the same fixed size, and each with its own local coordinate frame. Restricting the local maps to a certain size has the advantage that their complexity is limited and known. However, partitioning the environment in this arbitrary way rather than exploiting the natural structure inherent in the environment to identify each local space adds complexity to the transitions from one local map to another. In our system, exits determine the boundary of the local space, and are then the transition points between adjacent local maps. These exits carry an expectation that crossing a particular exit will take the robot into a particular neighbouring ASR. ATLAS constructs a signature for its local maps which comprise non-repetitive features from within the local frame. Cycles are detected by matching the local map signatures. The idea of using a subset of distinctive features within the local map to recognise places that the robot is revisiting is similar to the topological matching approach that we employ in Jefferies *et al*. (2003, 2004). ATLAS does not use global map matching; it uses local map matching and from the consistent local maps builds the global metric map. ATLAS constructs a signature for its local maps, as we do in Jefferies *et al*. (2003, 2004), but in Bosse's local map these comprise non-repetitive features from within the local frame. Our signatures are constructed from the features in the ASR which distinguish it from other ASRs. ATLAS's non-repetitive features could easily be those features that are common to other local maps, giving a higher likelihood of false positive matches. The map-matching process is a search for a coordinate transformation, based on the signatures, that brings overlapping frames into alignment.

Journal: **JIMS** MS.: **NO00004372** PIPS: **NO00004372** ☑TYPESET ☐DISK ☐LE ☑CP Disp.: **14/9/2005** Pages: **10**

Kuipers (2000) has long argued for a layered approach to mapping with the topological map preceding the global metric map in the hierarchy. Thus like ATLAS, Kuipers combines local metric maps in a topological map to construct the global metric map (Modayil *et al*., 2004). A set of likely topological maps is maintained rather than a single map hypothesis as in Bosse *et al*. (2003) and Jefferies *et al*. (2003). Closing a cycle in the global metric map involves selecting the correct topological map. However, this assumes that the cycle has been found in the topological map. This approach is appealing as it avoids the problem of having to propagate an error correction factor back through the global metric map when a cycle is found. The nodes in Kuipers and Beeson's (2002) topological map, are distinctive states. The edges connecting them are a description of the actions required to travel between adjacent distinctive states. A k-means clustering algorithm is used to place different images of the same distinctive state in the same cluster thus reducing image variability due to noise. However, this means that similar images belonging to different states will also be placed in the same cluster (the perceptual aliasing problem). If the image variability problem is addressed, one would assume that sufficient images have been captured at each distinctive state.

Most of the occupancy grid based mapping approaches use global metric maps but rarely do they exploit advantages that a topological map can offer. One of the main problems with most grid based approaches is that because they maintain a single map hypothesis, choosing the winning hypothesis at each step, there is no way back should that hypothesis eventually fail. Hähnel *et al*. (2003b) use a "lazy data association" approach which can "repair" poor choices once it is discovered that they are wrong. The mapping method is global metric but pivotal to the approach is the linking of the occupancy grid sub map for each pose in a topological map, explicitly representing the path information. A tree of alternative path hypotheses is maintained. When the current hypothesis no longer provides the best explanation of the data, the tree is searched for an alternative best hypothesis. Thus finding the correct correspondences when a cycle is encountered in a global metric map involves discovering that the current hypothesis is incorrect, and choosing a better alternative path. This could be some time after the cycle was encountered. If the map has diverged significantly from the correct path it may not be possible to find a suitable alternative hypothesis.

5. 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 (Jefferies *et al*., 2003) we are able to close many cycles in large-scale environments. 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

Beeson, P., MacMahon, M., Modayil, J., Provost, J., Savelli, F. and Kuipers, B. (2003) Exploiting local perceptual models for topological map building. in *Proceedings IJCAI-2003 Workshop on Reasoning with Uncertainty in Robotics*, 2003.

Bosse, M., Newman, P., Leonard, J., Soika, M., Feiten, W. and Teller, S. (2003) An Atlas framework for scalable mapping. in *Proceedings International Conference on Robotics and Automation*, 2003.

Chown, E., Kaplan, S. and Kortenkamp, D. (1995) Prototypes, location, and associative networks (PLAN): Towards a unified theory of cognitive mapping. *Cognitive Science*, **19**, 1–51.

Gallistel, C. R. and Cramer, A. E. (1996) Computations on metric maps in mammals: Getting oriented and choosing a multi-destination route. *The Journal of Experimental Biology*, **199**, 211–217.

Gutmann, S. and Konolige, K. (1999). Incremental mapping of large cyclic environments. in *Proceedings International Symposium on Computational Intelligence in Robotics and Automation*, 1999.

Journal: **JIMS** MS.: **NO00004372** PIPS: **NO00004372** ☑TYPESET ☐DISK ☐LE ☑CP Disp.: **14/9/2005** Pages: **10**

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

Hähnel, D., Thrun, S., Wegbreit, B. and Burgard, W. (2003b) Towards lazy data association in SLAM. in *Proceedings 10th International Symposium of Robotics Research*, 2003.

Jefferies, M. E., Cree, M., Mayo, M. and Baker, J. T. (2004) Using 2D and 3D landmarks to solve the correspondence problem in cognitive robot mapping. in *Proceedings Spatial Cognition 2004*.

Jefferies, M. E. and Yeap, W. K. (1998) Representing the local space qualitatively in a cognitive map. in *Proceedings Twentieth Annual Conference of the Cognitive Society*, 1998.

Jefferies, M. E., Yeap, W. K. and Baker, J. (2002) 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 K., D., (eds), WSEAS Press, pp. 287–294.

Jefferies, M. E., Weng, W., Baker, J. T., Cosgrove, M. C. and Mayo, M. (2003) A hybrid approach to finding cycles in hybrid maps. in *Proceedings Australian Conference on Robotics and Automation*, 2003.

Kuipers, B. (2000) The spatial semantic hierarchy. *Artificial Intelligence*, **119**, 191–233.

Kuipers, B. and Beeson, P. (2002) Bootstrap learning for place recognition. in *Proceedings 18th International Conference on Artificial Intelligence*, 2002.

Kuipers, B. J. and Byun, Y.-T. (1998) A robust, qualitative method for robot spatial learning. in *Proceedings of the National Conference on Artificial Intelligence (AAAI-88)*, 1988.

Kuipers, B., Modayil, J., Beeson, P., McMahon, M. and Savelli, F. (2004) Local metrical and global topological maps in a spatial semantic hierarchy. in *Proceedings IEEE International Conference on Robotics and Automation*, 2004.

Modayil, J., Beeson, P. and Kuipers, B. (2004) Using the topological skeleton for scalable global metric map-building. in *Proceedings IEEE/RSJ International Conference on Intelligent Robotics and Systems*, 2004.

Rofer, T. (2002) Using histogram correlation to create consistent laser scan maps. in *Proceedings IEEE International Conference on Intelligent Robotics Systems*, 2002.

Thrun, S. (1998) Learning metric-topological maps for indoor mobile robot navigation. *Artificial Intelligence*, **99**(1), 21–71.

Thrun, S., Hahnel, D., Ferguson, D., Montemerlo, M., Triebel, R., Burgard, W., Baker, C., Omohundro, Z., Thayer, S. and Whittaker, W. (2003) A system for volumetric robotic mapping of abandoned mines. in *Proceedings International Conference on Robotics and Automation*, 2003.

Tomatis, N., Nourbakhsh, I. and Siegwart, R. (2001) Simultaneous localization and map building: A global topological model with local metric maps. in *Proceedings International Conference on Intelligent Robots and Systems*, 2001.

Tomatis, N., Nourbakhsh, I. and Siegwart, R. (2002) Hybrid simultaneous localization and map building: Closing the loop with multi-hypotheses tracking. in *Proceedings International Conference on Robotics and Automation*, 2002.

Yeap, W. K. (1998) Towards a computational theory of cognitive maps. *Artificial Intelligence*, **34** 297–360.

Yeap, W. K. and Jefferies, M. E. (1999) Computing a representation of the local environment. *Artificial Intelligence*, **107**, 265–301.

Journal: **JIMS** MS.: **NO00004372** PIPS: **NO00004372** ☑TYPESET ☐DISK ☐LE ☑CP Disp.: **14/9/2005** Pages: **10**