Merging Topological Maps for Localisation in Large Environments

F. Ferreira * and J. Dias

Institute of Systems and Robotics- Coimbra, Portugal
Coimbra, 3030-290 Coimbra, Portugal
**E-mail: {cfferreira* | jorge}@isr.uc.pt*
http://paloma.isr.uc.pt

V. Santos

Department of Mechanical Engineering,
Aveiro, 3810-193 Aveiro, Portugal
E-mail: vitor@ua.pt

This article presents a method for the creation of a topological map without having to previously create a geometric map. Independently obtained topological paths are compared pairwise to search for possible overlap in the view sequence. Each individual topological path is created from a sequence of raw data views that are sampled by leading the robot along a path in the environment. The multiple topological paths are 'stitched together' into a generic Topological map by identifying overlapping segments in the individual sequences. A general topological map can be created by considering all the multiple sequences or separate runs through the environment. Results on the merger of upto 8 seperate paths indicate that this method of mapping large environments, by selective touring through the environment, can generate robust topological maps.

Keywords: topological localisation, map merging.

1. Introduction

The merging of smaller maps to create larger maps is a relevant topic in map building for Robot Navigation. There are various reasons why it might required to merge two or more smaller maps. One recent application of map merging techniques has been to integrate individual maps from cooperating robots. [1] The use of topological information to robustly merge maps is reported in a number of other recent works in the literature.

Within Geometric maps, Simultaneous Localisation and Mapping (SLAM) algorithms and their variants maintain the relation between features through the use of the relative positions and error correlation matrices. [2] SLAM is well suited to the process of incrementally mapping the environment. An extension of the map is performed by adding the position of new features and updating the error matrices with respect to the existing features (the existing map). The merger of two separate maps or sections is a difficult problem and is often referred to as the loop closure problem. In the case of SLAM, loop closure involves the conciliation of the position and error matrices for the maps/separate sections whilst simultaneously associating features from one of the maps/sections with the other.

The problem is difficult to solve since the position is not accurately known and other means, image features in the case of,[3] must be used to decide where merging will occur and where only position update of features in each map is performed. Once the position at which the maps must be merged is established, methods such as[4] can be used to locally search for and optimise the transformation that best 'fits' one map into another.

The similarity between places can be usually established by comparing changes in the sensor data over a short path. A similarity term might be used to split regions in the environment. For example in,[5] Schmidt et al. use the 'width' or smaller dimension of the environment, as measured using ultrasound sensors as the robot moves, to verify the similarity of places in the environment. In[6] the Fourier components that make up Omni-directional images are compared for an evaluation of similarity. Ho and Newman[7] and Posner, Schroeter and Newman[8] aim to use place similarity measures to improve the robustness of loop-closure and map merging algorithms that originally use geometric information. A recent article by Zivkovic et al[9] also attempts to cluster views by comparing near-by scenes to define places in the environment. In,[10] Thomas and Donikian hypothesise a hierarchical set of [topological] representations that represent the environment using similarity of places. The developers of these methods claim that such labeling of (similar-looking) places is in line with the spatial concepts that humans employ.

In[11] a geometric map is broken up into a number of smaller, 'more unique', environment layouts such as junctions and door openings and then strung back together in the form of a Dirichlet process. The method dwells on the problem of estimating the likelihood of previously unseen environment layouts and builds the prior probabilities of the different environment layouts.

Hierarchical representations of the environment have been used in the literature since the introduction of the Spatial Semantic Hierarchy, SSH, by Kuipers.[12] These methods seek to identify previously visited places within a topological representation. For example, in,[13] an extension of the SSH architecture, Kuipers and Beeson seek to merge topological information by testing out the merging hypothesis. The environment is actually sensed using very imprecise sensors and is subsequently abstracted in terms of corners and junctions. The hypothesis for merging is evaluated by actually moving the robot around the environment such that loop closure can be verified.

The selection of 'good' places at which the maps can be merged is an important problem that must be solved before merging can occur. In approaches such as the Semantic Hierarchy,[14] and again in,[15] a pre-defined set of indoor environmental characteristics such as junctions, doors and corners are utilised as possible candidates for merging sections of the map given the reliability with which these places can be detected.

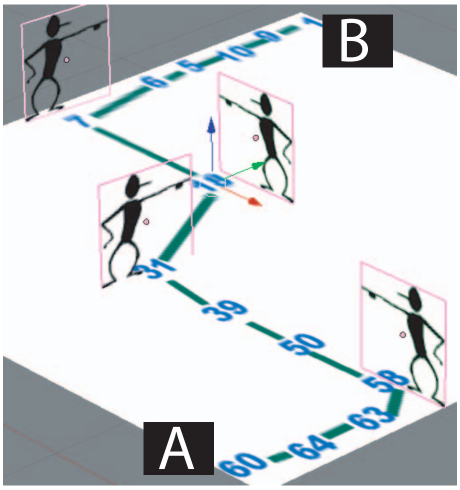

Fig. 1: The robot is led through the environment on the Environment Familiarisation run, top. The 3D point clouds must be registered in the environment, bottom.

In previous work, we have presented a method for representing a topological path using a sequence of raw data,[16] The Environment Familiarisation process results in the creation of a sequence of views referred to as the Reference Sequence. Let us suppose that the robot is led down two distinct paths within the environment, repeatedly sampling data using its sensors (laser range finder and cameras). Now, the robot has two sequences of views within which it can concurrently localise itself. To know its global position, the robot must decide along which path it is located and then localise itself within that path.

This article presents a method that can maintain a global estimate of the robot position by successively locating the robot along multiple paths. Using this method, the robot performs place recognition independently and simultaneously along each path. The complete environment within which the robot operates can be taken to be composed of multiple paths.

Two issues must be solved in order to obtain this topological map:

- a procedure to maintain a consistent probability distribution across multiple paths must be developed, and,
- a procedure for merging paths with a metric or measure that indicates the similarity between individual paths must be specified.

In the next section, a brief review of the procedure used to create topological representations of each path is presented. In section 3, a procedure to simultaneously localise the robot within multiple sequences is presented. In section 4, an algorithm is presented to merge segments of data sequences into a generalised topological map. Experiments and results of merging topological paths are presented in section 5.

2. Place recognition along a Single Sequence of Views

During an environment familiarisation phase, depicted at top in Fig. 1, the robot samples the environment according to a sampling plan, collecting features using its various sensors. A repetition of the sequence of motion performed during the place recognition will propel the robot along the original path, i.e. the graph depicted in Fig. 2. Any action other than the one taken during the environment familiarisation phase might result in the robot acquiring a previously unseen view. To accommodate these possible views, the graph that denotes any subsequent run includes some additional states, the $Lost\_Places$ as seen in Fig. 3. These $Lost\_Places$, a total of K in all, represent insertions in the original Reference Sequence. Each $Lost\_Place$ therefore takes into account the fact that the robot might be seeing views that were not seen in the environment familiarisation phase.

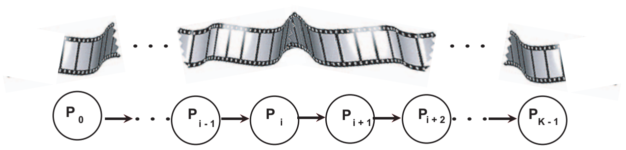

Fig. 2: This figure models the 'K' Reference Sequence as a Markov Chain according to the order in which they were sampled during the environment familiarisation phase.

The sequence begins with P_{Lost_0} which indicates the robot is completely lost or has never localised. Also, before every original place P_i there is a P_{Lost_i}. By moving forward from a lost place , the robot can transition from P_{Lost_i} to any node P_k where $k > i$. Similarly, from P_i the robot can transition to $P_k : k > i$ or to P_{Lost_i+1}. The graph does not allow a one-step transition from any P_{Lost_i} to another P_{Lost_j}.

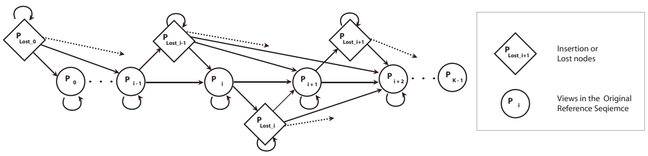

Fig. 3: During subsequent runs, '$Lost\_Places$' are inserted within the original Reference Sequence. The dotted lines indicate the transitions to each of the Places in the original Reference Sequence which have not been drawn to avoid cluttering the figure.

The Hidden Markov Model (HMM) and the Viterbi algorithm, is utilised to match the current sequence of views with the sequence of views previously obtained during the environment familiarisation phase. A simple robot motion model generates a distribution that provides a non-zero probability of transitioning to the places further down the sequence. Details of transition probabilities can be found in.[17]

In the absence of any better information the observation probability at any $Lost\_Place$ is arbitrarily defined as an Uniform distribution over the K original views.

The observation model of the HMM is based on matching the view currently captured by the robot sensors with views gathered during the environment familiarisation phase. The features from the views in the Reference Sequence are converted into binary form as described in[16] and are represented within a *Feature Incidence Matrix (FIM)*. Due to the large dimensionality of the FIM, it is modelled as a Bernoulli Mixture Model (BMM). These parameters of the BMM are obtained by running the Expectation Maximisation(EM) algorithm. The *Maximum Likelihood Estimation* approach is used to obtain the lost match for the current observation.

3. Simultaneously localisation in Multiple Topological Paths

When more than the one Reference Sequence is available for the robot to localise itself against, the global place recognition probability distribution can be represented as a combination of two independent probability distributions (1).

$$P(k_s, s|V^{obs}) = \frac{P(s) * P(k_s|V^{obs}, s)}{\text{constant}} \tag{1}$$

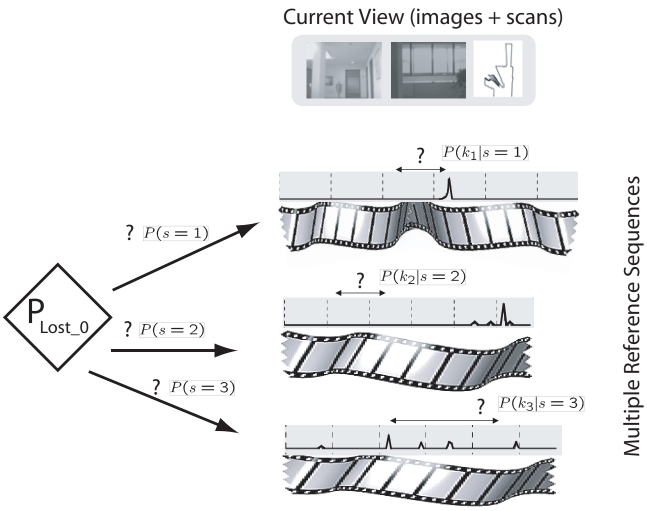

Fig. 4: A schematic for the comparison with multiple paths.

This evaluation of the position of the robot in the entire environment is depicted in Fig. 4. This method is proposed based on the corollary proposed by Montemerlo[18] which notes that the sensor observations are conditionally independent on the robot path.

The second term in the numerator in the right-hand-side of (1), is the probability of the place or sequence of places that most likely generated the current views, as summarised in section 2.

The first term in the numerator on the right side of (1) refers to the probability $P(s)$, of the robot being within a Reference Sequence, s. We have developed a measure that reflects the marginal probability, P_s of the robot being within each Reference Sequence. This measure is based on the *quality* of the place recognition within the Reference Sequence. The correct path along which the robot is moving should have the lowest uncertainty for the current observations. The uncertainty of being at a single place in a Reference Sequence can be modelled using an expression based on the uncertainty of the estimated probability distribution on the places within each of the Reference Sequence. This uncertainty can be evaluated using the appropriate expression to calculate the entropy of a random variable as in (2), where X represents a random variable, $\hat{k_s}|V^{obs}$ in the context of a single view localisation and $\hat{\text{traj}}_s|V^{obs}$ in the context of multiple view localisation.

$$H(X) = \sum P(x) \log(P(x)) \qquad (2)$$

In the case of single view place recognition, the distribution for $P(k_s|V^{obs}, s)$ provides the estimate of the robot position given the current view. An expression based on the entropy of this variable is used as a proxy for the probability of the robot being on the current Reference Sequence (3).

$$P(s) = \frac{1}{\sum_{s=1}^{S} P(s)} \Big[1 - \frac{1}{\log(K_s)} \times \sum_{k_s=1}^{K_s} P(k_s|V^{obs}, s) \log(P(k_s|V^{obs}, s))\Big]. \tag{3}$$

The above expression is based on the 'normalised entropy' and has been used to model the probability of being in a particular sequence, $P(s)$, based on a single view. Such a measure gives an idea of how uncertain a place recognition is within that Reference Sequence. The *less uncertain* a probabilistic distribution is over a Reference Sequence, the higher the probability of that Reference Sequence leading to better results for localisation. With the arrival of new sensor data, the probability of being in one sequence should begin to improve to the detriment of the other Reference Sequences. The normalising term $\frac{1}{\log(K_s)}$ allows the comparison of trajectories of different lengths, where K_s denotes the number of places in s.

4. Merging Topological Paths

In the context of this article, at a coarser or 'higher' level, the environment is viewed as being composed of multiple paths. Each path in turn is made up of a sequence of views. This representation is depicted in Fig. 5, by using the topological map described in.[19] The global topological map that we propose to build from a merger of paths can be contrasted with the method described in,[19] where the 'global' topological map is constructed by abstracting out certain properties of the 'global' geometric map.

The merging of paths to represent an environment in the form of distinct sequences of views includes the identification and removal of overlapping segments of the Reference Sequence from multiple paths.

The merging of topological maps is performed in a two-step procedure: 1) finding tentative sections for path overlap and 2) Identifying real overlapping sections. This two-step process is described in the following sub-sections.

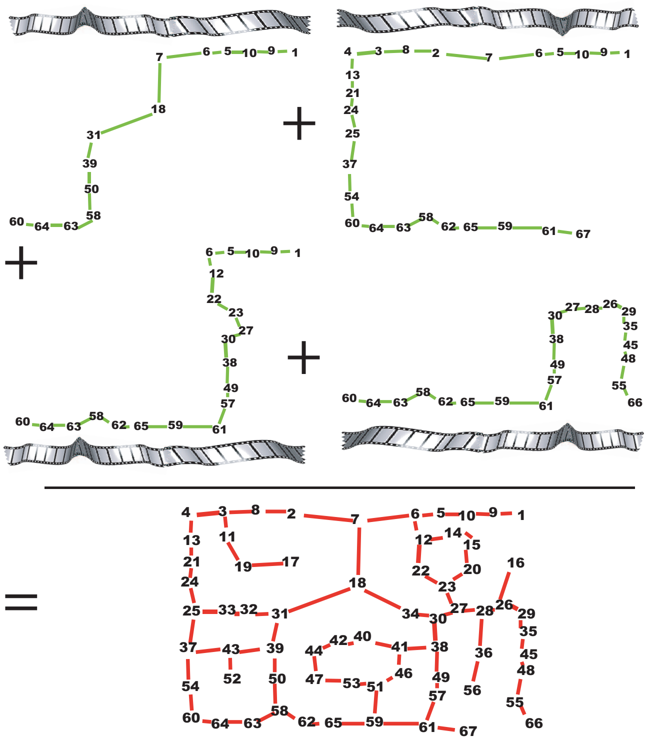

Fig. 5: The global topological map can be viewed as a collection of multiple, non-overlapping paths through the environment. This process of creating a topological map by merging the topological representation along multiple paths can be contrasted with methods such as,[19] where the 'global' topological map is constructed by abstracting out the 'global' geometric map.

4.1. *Finding tentative sections for place recognition*

An algorithm that measures view-similairty is used to identify individual instances of such overlaps. The proposed method compares a pair of sequences with each other such that a *distance* is calculated between each view in one sequence and every view in a second sequence. The result is a similarity matrix that indicates how similar a particular view is with each view inside the Reference Sequence. An example of such a comparison is seen in Fig. 6 for two sequences taken along Department of Electronics and Computer Engineering, DEEC, at the University of Coimbra. Places that are very similar in this matrix are utilised to detect possible overlaps between parts of the sequences. The overlap between the sequences is verified according to procedure described in the next sub section.

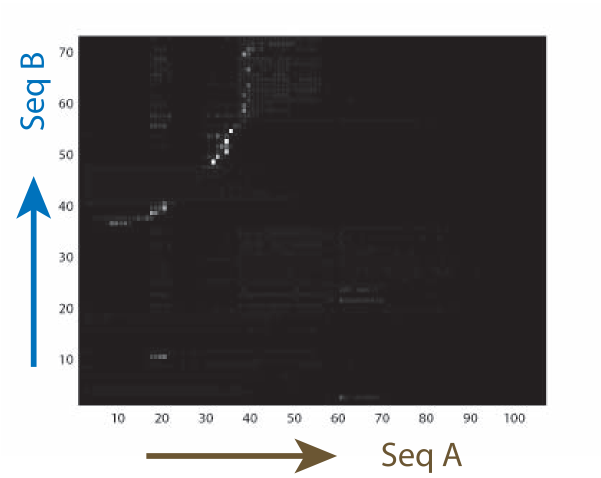

Fig. 6: View similarity matrices for a comparison between two paths. The similarity (lighter indicates greater similarity) between views is calculated using a metric based on the number of features that the views have in common after having reduced the dimensionality of the feature space using the Bernoulli Mixture Model.

4.2. *Identifying Overlap in View Sequences*

In the second step, a simple procedure, Algorithm 4.1, tests each of the hypothesis for overlap that are provided by the previous step. The algorithm entails a node-by-node alignment in the neighbourhood of the nodes which were deemed to be similar from the view-similarity matrix. If the similar nodes are relatively close to each other suggesting that the non-similar intermediate nodes results from sensor noise, that part of the paths are considered to be overlapping.

A degree of robustness is imparted to the merging technique by searching for a minimum amount of overlap between a pair of sequences around the location of the potential cross-overs.

Without loss of generality, Fig. 8a shows the effect of matching a segment of two Reference Sequences in for the experiment described in section 5.

Algorithm 4.1 Merging Hypothesis For Topological Paths

$H_{new} = NULL$ // the new sequence is completely separated
${}^{min}t_{new}$ // the minimum trajectory length for declaring overlap
Require: $V_{new} \geq^{min} t_b$ // the new sequence has min length
$T_{new} = V_{new}$ // all new views are potential match
while $(T_{new} > 0)$ **do**
  $i \leftarrow V^1_{new}$
  $j \leftarrow V^{1+^{min}t_{new}}_{new}$
  if TEST$(t_{i,j})$) **then**
    // test for no overlap
    $H \leftarrow t^b_{i,j}$ //no overlap, current trajectory in B is added to Hypothesis
    remove $V^1_{new}, \ldots V^{1+^{min}t_{new}}_{new}$ from T^b //current trajectory removed from further tests
  else
    // overlap confirmed
    remove $V^1_{new}, \ldots V^{1+^{min}t_{new}}_{new}$ from T^b //current trajectory removed from further tests
  end if
end while
return H // the non-overlapped sequence is returned

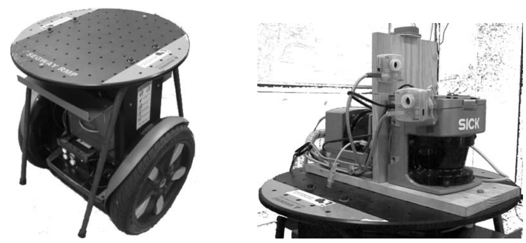

(a) The Segway RMP 200.

(b) Sensor Platform consisting of Camera 1 and Camera 2.

Fig. 7: The Segway Robotic Mobile Platform (RMP) and the sensor platform.

5. Experiments and Results.

In the above sections we have described a method that could be used to merge overlapping sections of multiple paths and a procedure to localize along these multiple paths.

Two data sequences were obtained by leading the robot, depicted in Fig. 7a, through two paths having some overlap. Each sequence consisted of 640x480 gray scale images taken by the two cameras seen in Fig. 7b. Using the two-step procedure identified in section 4, the overlapping segment is identified in between Places 16 and 32 in sequence #1 and Places 22 and 46 in sequence #2, see Fig. 8a.

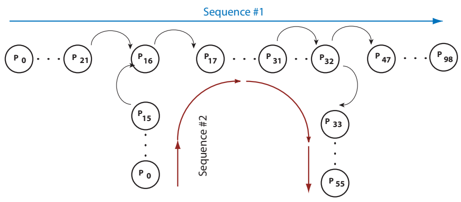

(a) Path #2, 98 views, is merged to path #1, 55 views.

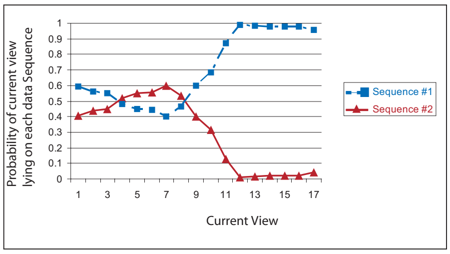

(b) The figure plots the probability of the robot being on each of the original paths #1 and #2.

Fig. 8: Experiment 1: Merged Topological graph and probability distribution.

For demonstration purposes, the overlapping segment is not removed from either Reference Sequence and the robot is subsequently, driven along a portion of path #1. The robot localises itself along each path and globally. The probability of being on each of the original sequences is plotted and shown in Fig. 8b. While the robot is on the overlapping region, either sequence is probable. Upon reaching the place where the paths separate, the probability of the robot being on the incorrect path rapidly reduces. This behaviour indicates why the overlapping segments must be removed.

The robot maintains an estimate of its position in the topological map by comparing its current view with each trajectory in map: locating the robot in the global map at the place that is most similar to the currently viewed data sequence.

To demonstrate the effect of the removal of overlapping segments, a second experiment was carried out on the a longer environment consisting of a set of

8 Reference Sequences, collected by driving the robot forwards and backwards along 4 paths in the environment. The 4 paths cover a distance of approximately 300 meters and overlap to various extents. The plan view of the stretch of the environment is seen in Fig. 9a. Sample images from six of the eight Reference Sequences are shown in Fig. 10 for the *Camera #1*. The 8 paths are named thus:

- Reference Sequence A Forward: moving forward along path A.
- Reference Sequence A Reverse: Retracing in reverse along path A
- Reference Sequence B Forward: moving forward along path B.
- Reference Sequence B Reverse: Retracing in reverse along path B.
- Reference Sequence C Forward: moving forward along path C.
- Reference Sequence C Reverse: Retracing in reverse along path C.
- Reference Sequence D Forward: moving forward along path D.
- Reference Sequence D Reverse: Retracing in reverse along path D.

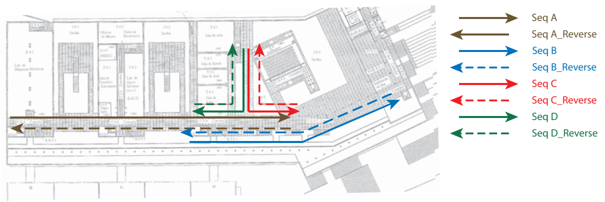

(a) Layout of the 6 Reference Sequences that were recorded on the fourth floor of the DEEC Building.

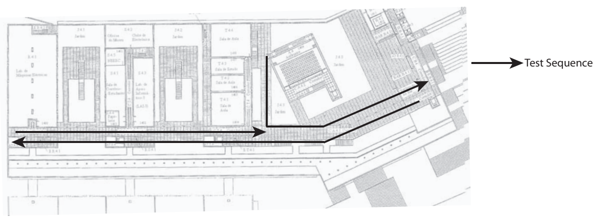

(b) Layout of the test path on the fourth floor of the DEEC building over which localisation was performed.

Fig. 9: Experiment 2: Eight sequences with overlapping segments were gathered from the environment, as shown at top. Subsequently, the robot was driven over part of the region covered by the 8 Reference Sequences, as shown at bottom.

The procedure outlined in the previous sections are followed to merge the Reference Sequences. The candidates for map-merging in the view similarity are then verified for actual overlap according to the procedure outlined in section 4 and the overlapping segments are removed from the *longer* Reference Sequence. This is a simple criteria for the removal of overlapping sequences and was adopted on account of its simplicity.

Subsequently, the robot is driven along part of the environment covered by the 8 Reference Sequences. The path covered by the robot is shown in Fig. 9b. Localisation was performed simultaneously along all 8 Reference Sequences with the overlapping segments removed. The global position of the robot is obtained by combining the probability $P(s)$, of the Robot being along a Reference Sequence s and the probability distribution $P(k_s|V_{obs}, s)$ over the views in that Reference Sequence. The probability distribution $P(s)$ is plotted in Fig. 11 over the entire test path. The robot is positioned on the Reference Sequence with the greatest value of $P(s)$.

It can be seen in Fig. 11, that when the robot is travelling along the part of the environment covered by Reference Sequences A Reverse, B Reverse, C Reverse, the values of $P(s)$ are quite stable and a single Reference Sequence is consistently selected, leading to correct global localisation of the robot over the topological map.

At other times, as can be seen in Fig. 11, there is a large variation in the probability distribution $P(s)$. This occurs partly because, even when there was no overlap, some of the paths are quite similar to each other, for example in Reference Sequence A and B both cameras were facing regions of the environment that were very similar and partly because at other places, very few features could be extracted, for example in Reference Sequence C and D the sidewards facing cameras were looking at a textureless wall.

6. Future work

This article presented a method that could be used to create topological representations of large environments by merging multiple topological paths. The results are still preliminary but the method offers a powerful and, in our opinion, new approach to the creation of topological maps. A targeted area of application is in the mapping of large environments in which the robot can be taken on selective tours through the environment, potentially speeding up the mapping process.

From the experiments described in the earlier section, such observations lead us to believe that the merging of Topological maps will be dependable only at those places where the views in the Reference Sequences are quite distinct, i.e. wherever localisation is *less uncertain*.

We have also seen how the removal of the overlapping segments allows the identification of those places where the localisation is successful and others where it fails. The situations which result in similar values of the maximum $P(s)$ for mul-

(a) Reference Sequence A Forward

(b) Reference Sequence A Reverse

(c) Reference Sequence B Forward

(d) Reference Sequence B Reverse

(e) Reference Sequence C Forward

(f) Reference Sequence C Reverse

Fig. 10: Experiment 2: Some of the Images from *Camera #1* that are included in 6 of the 8 Reference Sequences captured prior to merging.

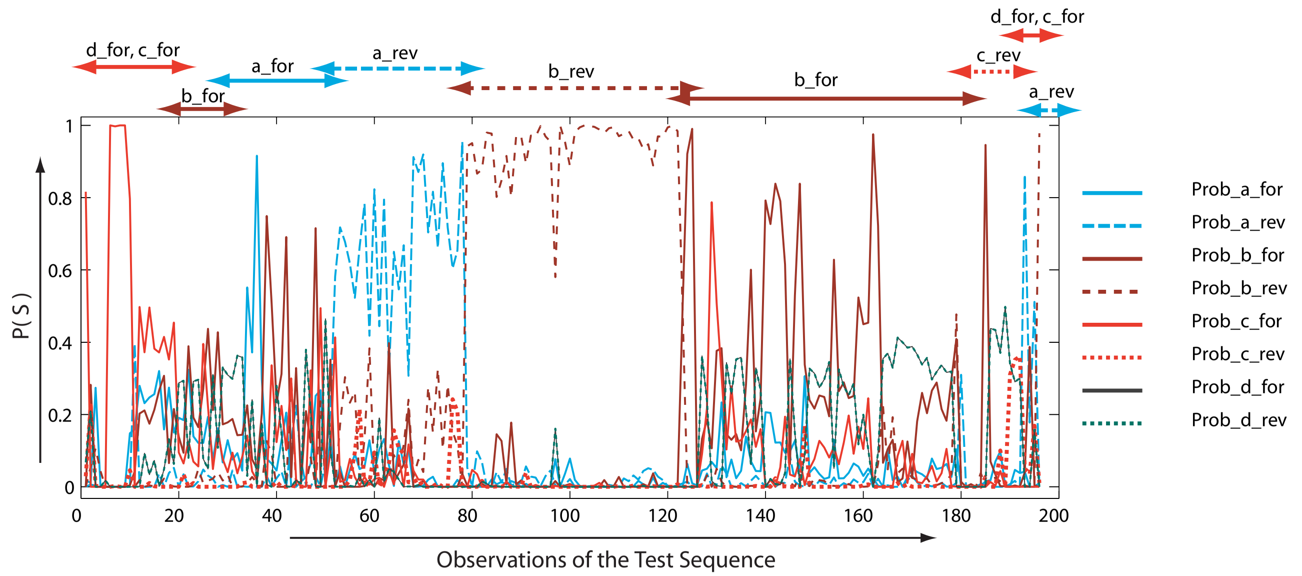

Fig. 11: Experiment 2: The figure includes a plot of the variation in $P(s)$ for each of the 8 Reference Sequences that make up the merged map over the entire path taken by the robot during the localisation test.

tiple Reference Sequences are an indication that the robot cannot localise itself.

It is important to emphasize what the 'merging' of maps does not do. Since the view sequences are devoid of spatial or geometric information, it is not possible to build up the global map of the environment where the environment paths intersect at the overlap in the environment. Therefore, it will not be normally possible to arrange the paths within a larger graph and to represent individual sequences with relation to each other.

Instead, the method that is proposed in this article seeks to create a consistent means of keeping the robot localised within the set of paths, according to the method described in section 3.

Future work is under way to allow the intervention of a user to improve/correct the process of creation of the topological map. User intervention could be directed to correct adjacency relationships on different sequences and to reduce the search space for overlapping segments during the creation of the topological map.

7. Acknowledgements

This work was partially supported by the IRPS, EU-FP6-IST-045048 project. Initial work on the procedure described in this work was orally presented at the 8th Conference on Autonomous Robot Systems and Competitions at Robotica 2008 in Aveiro, Portugal.

References

1. Kurt Konolige, Dieter Fox, Charlie Ortiz, Andrew Agno, Michael Eriksen, Benson Limketkai, Jonathan Ko, Benoit Morisset, Dirk Schulz, Benjamin Stewart, and Regis Vincent. Centibots: Very large scale distributed robotic teams. In *Proceedings ISER2004*, 2004.
2. Sebastian Thrun, Wolfram Burgard, and Dieter Fox. *Probabilistic Robotics*. MIT Press, 2005.
3. P. Newman, D. Cole, and K. Ho. Outdoor slam using visual appearance and laser ranging. In *ICRA 06*, 2006.
4. Stefano Carpin, Andreas Birl, and Viktoras Jucikas. On map merging. *Robotics and Autonomous Systems*, 53(1):1–14, 2007.
5. Jochen Schmidt, Chee K. Wong, and Wai K. Yeap. A split & merge approach to metric-topological map-building. *ICPR*, 3:1069–1072, 2006.
6. H. Ishiguro and S. Tsuji. Image-based memory of environment. In *International Conference on Intelligent Robots and Systems '96, IROS 96, Proceedings of the 1996 IEEE/RSJ*, volume 2, Osaka, Japan, November 1996.
7. Kin Leong Ho and Paul Newman. Multiple map intersection detection using visual appearance. In *Ciras 05*, 2005.
8. I. Posner, D. Schroeter, and P. Newman. Using scene similarity for place labelling. In *ISER 06*, 2006.

9. Zoran Zivkovic, Olaf Booij, and Ben Króse. From images to rooms. *Robotics and Autonomous Systems*, 55(5):411–418, 2007.
10. Romain Thomas and Stéphane Donikian. A spatial cognitive map and a human-like memory model dedicated to pedestrian navigation in virtual urban environments. In *Spatial Cognition*, pages 421–438, 2006.
11. Benjamin Stewart, Jonathan Ko, Dieter Fox, and Kurt Konolige. The revisiting problem in mobile robot map building: A hierarchical bayesian approach. In *Conference on Uncertainty in Artificial Intelligence*, 2003.
12. Benjamin Kuipers. *Representing Knowledge of Large-Scale Space*. PhD thesis, Mathematics Department, Massachusetts Institute of Technology, Cambridge, Massachusetts, June 1977. Also MIT-AI TR-418, 1977, Available via ftp://ftp.cs.utexas.edu/pub/qsim/papers/Kuipers-PhD-77.ps.gz.
13. Benjamin Kuipers and Patrick Beeson. Bootstrap learning for place recognition. In *Eighteenth National Conference on Artificial Intelligence(AAAI-02)*, 2002.
14. B. J. Kuipers. The Spatial Semantic Hierarchy. *Artificial Intelligence, Elsevier*, 119:191–233, 2000. http://www.cs.utexas.edu/users/qr/papers/Kuipers-aij-00.html.
15. David Kortenkamp and Terry Weymouth. Topological mapping for mobile robots using a combination of sonar and vision sensing. In *Twelfth national conference on Artificial intelligence (vol. 2)*, pages 979 – 984, Seattle, Washington, United States, 1994.
16. F. Ferreira, V. Santos, and J. Dias. A topological path layout for autonomous navigation of multi-sensor robots. *International Journal of Factory Automation, Robotics and Soft Computing*, 1:203–215, 2007.
17. Filipe Ferreira, Vitor Santos, and Jorge Dias. *Emerging Technologies, Robotics and Control Systems*, chapter A Topological Path Layout for Autonomous Navigation of Multi-sensor Robots, pages –. InternationalSAR, 2007.
18. M. Montemerlo, S. Thrun, D. Koller, and B. Wegbreit. FastSLAM: A factored solution to the simultaneous localization and mapping problem. In *AAAI*, 2002.
19. Sebastian Thrun. Learning metric-topological maps for indoor mobile robot localization. *Artificial Intelligence*, 99(1):21–71, 1998.