Towards a Cognitive Probabilistic Representation of Space for Mobile Robots

Shrihari Vasudevan, Viet Nguyen & Roland Siegwart

Autonomous Systems Laboratory (ASL)
Swiss Federal Institute of Technology Zurich (ETHZ)
8092 Zurich, Switzerland

shrihari.vasudevan@ieee.org , viet.nguyen@epfl.ch , r.siegwart@ieee.org

Abstract - **Robots are rapidly evolving from factory work-horses to robot-companions. The future of robots, as our companions, is highly dependent on their abilities to understand, interpret and represent the environment in an efficient and consistent fashion, in a way that is comprehensible to humans.**

This paper is oriented in this direction. It suggests a hierarchical probabilistic representation of space that is based on objects. A global topological representation of places with object graphs serving as local maps is suggested. Experiments on place classification and place recognition are also reported in order to demonstrate the applicability of such a representation in the context of understanding space and thereby performing spatial cognition. Thus the theme of the work is – representation for spatial cognition.

Index Terms - Cognitive Spatial Representation, Robot Mapping, Conceptualization of spaces, Spatial Cognition

I. INTRODUCTION

Robotics today, is visibly and very rapidly moving beyond the realm of factory floors. Robots are working their way into our homes in an attempt to fulfill our needs for household servants, pets and other cognitive robot companions. If this "robotic-revolution" is to succeed, it is going to warrant a very powerful repertoire of skills on the part of the robot. Apart from navigation and manipulation, the robot will have to understand, interpret and represent the environment in an efficient and consistent fashion. It will also have to interact and communicate in human-compatible ways. Each of these is a very hard problem. These problems are made difficult by a multitude of reasons including the extensive amount of information, the huge number of types of data (multi-modality), the presence of entities in the environment which change with time, to name a few. Adding to all of these problems are the two simple facts that everything is uncertain and at any time, only partial knowledge of the environment is available.

The underlying representation of the robot is probably the single most critical component in that it constitutes the very foundation for all things we might expect the robot to do – these include the many complex tasks mentioned above. Thus, the extent to which robots will evolve from factory work-horses to robot-companions will in some ways (albeit indirectly) be decided by the way they represent their surroundings. This report is thus dedicated towards finding an appropriate representation that will make today's dream, tomorrow's reality.

II. RELATED WORK

Robot mapping is a relatively well researched problem, however, with many very interesting challenges yet to be solved. An excellent and fairly comprehensive survey of robot mapping has been presented in [1]. It introduces the problem of robot mapping, surveys the state-of-art and presents some of the open challenges in the field. Robot mapping has traditionally been classified into two broad categories – metric and topological.

Metric mapping tries to map the environment using geometric features present in it. These features could represent the occupancy of space [2], giving rise to occupancy grid maps, or other geometric features present in the environment such as lines [3], corners, landmarks etc., giving rise to what are better known as metric maps. A related concept in this context is that of the relative map, a map state with quantities invariant to rotation and translation of the robot. Ref. [4] is a recent example of this concept – it creates a relative map of landmarks, which are corners obtained as the intersection of line segments.

Topological mapping usually involves encoding place related data and information on how to get from one place to another. Different approaches have been used to form topological maps. These include the application of voronoi graphs, as shown in [5] and more recently the usage of feature based signatures to form topological maps, as demonstrated in [6]. Topological maps encode "high-level" [†] information and are thus less detailed, less voluminous and also less accurate than metric maps. More recently, a new scheme has become quite popular – the one of hybrid mapping [7, 8]. This kind of mapping typically uses both a metric map for precision navigation in a local space and a global topological map for moving between places.

The one similarity between all these representations is that all of them are navigation-oriented, i.e. all of them are built around the single application of robot-navigation. These maps are useful only in the navigation context and fail to encode the semantics of the environment. The focus of this

---

[†] Objects, doors etc. are considered "high-level" features contrasting with lines, corners etc. which are considered as "low-level" ones.

work is to address this deficiency. Several other domains inspire our approach towards addressing this challenge – these include hierarchical representations of space, high-level feature extraction, scene interpretation and the notion of a Cognitive Map.

A closely related set of works (to the work presented here) are those that suggest the notion of a hierarchical representation of space. Ref. [9] suggests one such hierarchy for environment modeling. In [10], Kuipers put forward a "Spatial Semantic Hierarchy" which models space in layers comprising respectively of sensorimotor, view-based, place-related and metric information. The work [11] probably bears the most similarity with the work presented in this paper. The authors use a naive technique to perform "object recognition" and add the detected objects to an occupancy grid map. The primary difference in the work presented here is that this approach uses objects as the functional basis in building the representation – i.e. the map is created and grown with the objects perceived.

Typically, humans seem to perceive space in terms of high-level information such as objects, states & descriptions, relationships etc. Thus, a human-compatible representation would have to encode similar information. The work reported here attempts to create such a representation using typical household objects and doors. Perhaps the single biggest obstacle towards constructing a representation as suggested in this work is the problem of high-level feature extraction (HLFE). Two examples of problems in this field are object recognition and door detection. Even the state of the art in the domain of HLFE can boast of very few techniques that can reliably perform in less than simple scenarios that involve challenges such as handling multiple objects, multiple views and occlusions to name a few. In the context of object recognition, a very promising approach, that has also been used in this work, is the one based on SIFT [12]. In our experience, this was found to be a very effective tool for recognizing textured objects. Several works attempt to model and detect doors. The explored techniques range from modeling/estimating door parameters [13] to those that model the door opening [14] and to methods that use more sophisticated algorithms such as [15], which is based on boosting.

More recently, scene understanding / interpretation has been gaining popularity in the robotics community. Two fine examples of this are [16] and [15]. In the former, the authors suggest a low dimensional global image representation; they use this in conjunction with hidden markov models to perform place recognition and categorization. They further use this information as a contextual prior towards object recognition and localization. The latter work uses the AdaBoost algorithm and simple low-level features (extracted from scans), vision together with hidden markov models to classify places.

This work takes inspiration from the way we believe humans represent space. Kuipers, in [17], elicits a conceptual formulation of the cognitive map. Yeap et al in their work [18] trace the theories that have been put forward to explain the

phenomenon of early cognitive mapping. They classify representations as being space based or object based. The work suggested here is based on the object based representation methodology. In summary, a single unified representation that is multi-resolution, multi-purpose, probabilistic and consistent is still a vision of the future and this is exactly what this work is intended to move towards.

III. APPROACH

A. Problem Definition

This work is aimed at developing a generic representation of space for mobile robots. Towards this aim, in this particular work, two scientific questions are addressed - (1) How can a robot form a high-level probabilistic representation of space? (2) How can a robot understand and reason about a place?

The first question directly addresses the problems of high-level feature extraction, mapping and place formation. The second question may be considered as the problem of spatial cognition. Together, when appropriately fused, they give rise to the hierarchical representation being sought. This representation must consider and treat information uncertainty in an appropriate manner. Also, in order to understand places, the robot has to be able to conceptualize space; to be able to classify its surroundings and to recognize it, when possible.

B. Overview

Figures 1 & 2 respectively show the mapping process and the method used to demonstrate spatial cognition using the created map. In an integrated system, the mapping and reasoning processes cannot be totally separated, but it is done here so as to facilitate understanding of the individual processes. Subsection C elicits the details of the perception system – this includes the object recognition and door detection processes. Subsection D specifies the details on how the representation is created – both local probabilistic object graphs and individual places. Subsection E addresses the issue of learning about place categories (kitchens, offices etc.). Subsections F and G explain how such a representation could be used for spatial cognition and finally subsection H deals with the process of updating such a representation. The remaining parts of the papers discuss the experiments conducted and the conclusions drawn thereof.
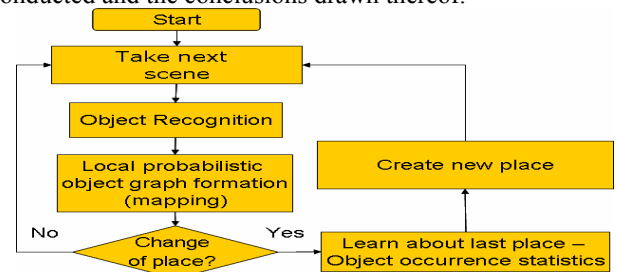

Fig. 1 The mapping process. High-level feature extraction is implemented as an object recognition system. Place formation is implemented using door detection. Beliefs are represented and appropriately treated. Together, these are encoded to form a hierarchical representation comprising of places, connected by doors and themselves represented by local probabilistic object graphs. Concepts about place categories are also learnt.
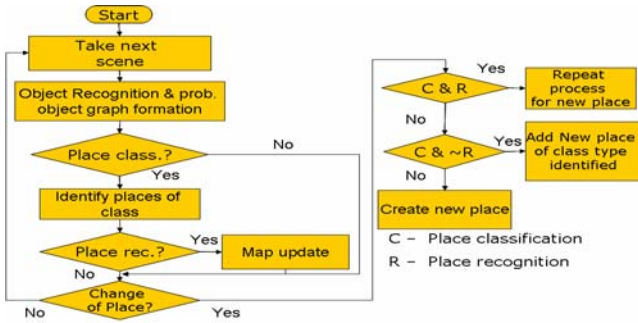
Fig. 2 The reasoning process for each place. First step is place classification – the robot uses the objects it perceives to classify the place into one of its known place categories (office, kitchen etc.). Next step is - recognizing specific instances of the place it is aware of – place recognition. Accordingly map update / adding of new place is done.

C. Perception

This work deals with representing space using high-level features. In particular, two kinds of high-level features are used here – typical household objects and doors. Reliable and robust methods for high-level feature extraction are yet unavailable. It must be emphasized that the perception component of this work, is not the thrust of this work. Thus, simplified and/or established algorithms have been used.

For this work, a SIFT based object recognition system was developed (Fig. 3) along the lines of [12]. Very briefly, the SIFT approach to object recognition is a "local-features" method. It does not learn any general properties of objects – in order to categorize and classify objects. It does however, transform a set of features, obtained from the object-of-interest using a naive technique, into a robust feature set that incorporates invariance to scale and rotation changes; to a considerable extent deals with illumination changes and changes in viewing direction as well. In our experience, this method was found to be very effective for recognizing most textured objects. More details on SIFT based object recognition can be obtained from [12]. The objects detected are used to represent places as explained in sub-section D.

Doors are used in this work in the context of place formation. A method of door detection based on line extraction and the application of certain heuristics, was used. The sensor of choice was the laser range finder. The door detection process involved the following steps:
  (1) Detect lines in 2D space (split & merge method, [19])
  (2) Apply heuristics to identify doorways.
  (3) Track door hypotheses. Use a likelihood-based tracking process to eliminate false positives.

References for places are placed at the end of the door that occurs first in the anticlockwise direction (generally the left) when the door is crossed. Occasionally, due to the coordinates of the door with respect to the robot position during detection, the same end of the door may be chosen as the reference both when entering the place and when leaving it. This algorithm was applied on a dataset with over 150 scans taken over several rooms of our laboratory. The robot was rigorously tested by moving it into corners, between

tables, walls etc. A promising performance was observed – this is indicated in table 1. and fig. 4.
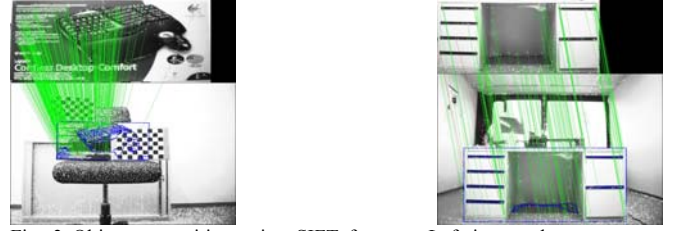

Fig. 3 Object recognition using SIFT features. Left image shows a carton being recognized, right image shows a table being recognized. Objects used in this work include cartons of different kinds, a table, a chair, a shelf & a mug.

TABLE 1
Door detection performance

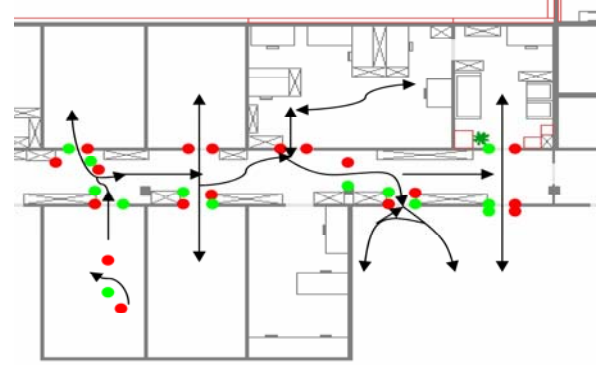| Number of good detections (door detections – going into a room + leaving a room)	14 (of 15 expected => 93.3 %)
Number of false positives (additional "noise" detections)	3

Fig. 4 Door detection algorithm tested over a part of our lab. A pair of dots represents a door. The red ones are the references of the place that are explored on crossing the door. The few false positives that were obtained were observed and removed using simple techniques.

D. Representation

The representation put forward here is a hierarchical one that is composed of places which are connected to each other through doors and are themselves represented by local probabilistic object graphs (probabilistic graph encoding the objects and relationships between them). Objects detected in a place are used to form a relative map for that local space. Doors are incorporated into the representation when they are crossed and link the different places together.

Object graphs were used by the authors in [20]. The problem with this work is that the information encoded in the representation was purely semantic and not persistent i.e. not invariant and not re-computable based on current viewpoint. This work addresses this drawback by drawing on the relative mapping approach in robotics. It suggests the use of a probabilistic relative object graph as a means of local metric map representation of places. The metric information encoded between objects includes distance and angle measures in 3D space. These measures are invariant to robot translation and rotation in the local space. Such a representation not only

encodes the inter-object semantics but also provides for a representation that could be used in the context of robot navigation.

The robot uses odometry to know the robot pose which is in turn used towards the creation of the relative object graph. The odometry model employed was that of a standard differential drive model as suggested in [3]. A stereo camera is used to know the positions of various objects in 3D space. The stereo model suggested by Jung in his work on SLAM (Simultaneous Localization and Mappping) using stereo vision [21] is used here.

As mentioned before, the representation is probabilistic. "Existential" beliefs (discrete probability values) are obtained from the perception system for each object that is observed. Simultaneously, "precision" beliefs are maintained in the form of covariance matrices. These beliefs are based on detailed mathematical formulations – the end result of which is shown below in figure 5. By representing both kinds of beliefs, such a representation will serve in the context of both, high level reasoning and scene interpretation and yet be useful for lower level navigation related tasks. As mentioned earlier, the relative spatial information encoded, include distance and angle measures in 3D space. These also have associated existence and precision beliefs – the computation of which takes the form shown in figure 6.

$$X_o = f(X_C) \; where \; f = M_{RO} * M_{CR}$$

(object in local place reference; O = local place reference frame,

C = camera reference frame & R = robot reference frame

$M_{RO} = $ Transformation between robot frame & absolute reference

$M_{CR} = $ Transformation between camera frame & robot frame)

$$P_o = F_1 P_1 F_1' + F_2 P_2 F_2' \;, \; where \; F_1 = J_{X_1}(f) \;\; \& \;\; F_2 = J_{X_2}(f)$$

(covariance matrix – uncertainty in object position

$X_1 \rightarrow (X_R, Y_R, \theta_R) \rightarrow$ robot pose

$X_2 \rightarrow (X_C, Y_C, \theta_C) \rightarrow$ object position in camera frame

$P_1 = $ covariance matrix $\rightarrow$ uncertainty in robot position.

$P_2 = $ covariance matrix $\rightarrow$ uncertainty in the object position.

$F_1 = $ jacobian of " f " with respect to X_1

$F_2 = $ jacobian of " f " with respect to X_2)

Fig. 5. Belief representation for individual objects in a place. The objects in 3D camera space are transformed into local place reference coordinates. The uncertainty in its position in the place reference frame is dependent on the uncertainty in the robots pose, the uncertainty in the object pose and the dependence of the transformation on them. This is computed as shown.

E. *Learning about place categories*

Concepts are learnt when creating the representation of various places. These encode the occurrence statistics of different objects in different place categories (office, kitchen, etc.). These statistics serve as likelihood values for a place classification procedure discussed in the next section. Thus, in a future exploration task, a robot could actually understand its environment and thereby classify its surroundings based on the objects it perceives.

Let $X_1 (x_1, y_1, z_1)$ & $X_2 (x_2, y_2, z_2)$ represent two objects

$f(X_1, X_2) \rightarrow$ relative spatial information between the two objects. (distance & angle measures)

P_1 & P_2 $\rightarrow$ uncertainty in object positions (covariance matrices)

Precision belief $\rightarrow Bel_1(f) = F_1 P_1 F_1' + F_2 P_2 F_2'$

where $F_1 = J_{X_1}(f)$ & $F_2 = J_{X_2}(f)$ are the jacobians of ' f '

with respect to X_1 and X_2 respectively and

Existential belief $\rightarrow Bel_2(f) = \min (belief \; in \; existence \; of \; objects)$

Fig. 6 Belief representation for relationships between objects. Given any two objects (in 3D space) in the place reference, the distance and angle measures are computed. The existence of a relationship is subject to the belief in the existence of the objects themselves. The precision of the relationship is subject to the relationships dependence on the objects coordinates and the uncertainty in the object positions. These are computed as shown.

F. *Spatial Cognition 1 – Place Classification*

Place classification is done in an online incremental fashion, with every perceived object contributing to one or more hypotheses of previously learnt place concepts. An informal specification of the algorithm is given below in figure 9. Hypothesis selection uses both, distinctiveness metric and a threshold metric to classify a place. The distinctiveness metric is used as a means of distinguishing between multiple competing hypotheses. A hypothesis would only be selected if it was sufficiently different from the nearest competing one. The threshold metric simply identifies potential hypothesis based on preset thresholds (in this case the prior belief – uniform value for all hypotheses in the beginning) for deeming that a classification has indeed occurred. In the case of place classification, the distinctiveness metric is given more importance as this process is about disambiguating between multiple competing hypotheses. The Laplace succession law is used for computing the likelihood of being in a place on observing a particular object, as shown in figs. 7 and 8.

$$P(place \mid object) = \frac{N\_place + 1}{N\_object + 2}$$

P(place | object) = Likelihood of place category given the object observed

N_object = number of occurrences of an object in a place category

N_place = number of occurrences of the place category

Fig. 7 Likelihood of being in a place of a particular place category given the occurrence of an object. Laplace succession law is used for the computation.
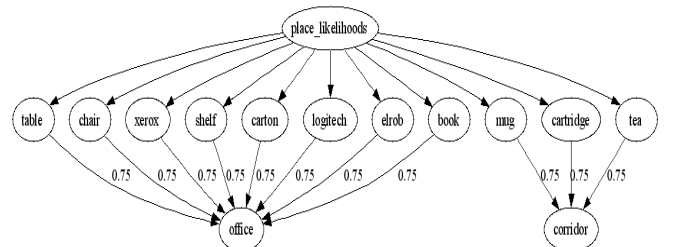

Fig. 8 Likelihoods of different place categories given the occurrence of different objects (after one mapping run). This encodes the number of times the robot was in a place (e.g. my office) of a particular place category (e.g. office) when it encountered a particular object (e.g. mug).

```
Algorithm:
(1) Initialize each place with a prior belief
(2) For each object that is observed
    a.    For each place that could have this particular object.
    b.    Posterior (place) = Prior (place) * Likelihood (place | object).
(3) Normalize beliefs of places.
(4) Hypothesis selection (based on distinctiveness & threshold)
```

Fig. 9 Algorithm for place classification. Starting with a uniform prior for all place hypotheses, the algorithm accumulates evidence (observation of objects) incrementally as the robot explores its surroundings.

G.   Spatial Cognition 2 – Place recognition

Place recognition is done by a graph matching procedure which matches both the nodes and its relationships to identify a node match. The number of node matches serves as a measure of recognition. Again, a hypothesis is finally selected on the basis of both distinctiveness and threshold metrics as explained in the context of place classification. However, in this case, the threshold is given much more importance than the other metric. This is because, the aim is to find the maximal common set of identically configured objects between a place the robot knows (previously mapped) and the one it currently perceives.

From a computational perspective, performing place classification prior to actual place recognition makes the graph matching procedure much more scalable by using the semantics of a place to zero down on a select set of places that could possibly match the one in consideration. Such a methodology would be visibly useful when the representation grows in complexity, size and number of instances of places in various place categories. Also data association at the level of objects is done by not only comparing the objects under consideration but also the relationships that the particular object obeys with its neighbors. This facilitates a distinctive representation of space.

H.   Map update

A map update operation (internal graph representation is updated) is required both for handling the revisiting of places and the re-observation of objects while mapping a place. It involves the following steps:-
  (1) Remove un-observed nodes and relations.
  (2) Increase belief for re-observed nodes and relations.
  (3) Add new nodes and relations.
  (4) Remove nodes & relationships with too low belief.
For step (1), two options exist: (a) to remove un-observed nodes and relations, (b) to reduce the beliefs and implement a gradual "forgetting" process. Both have been tried, but for the experiments in this work, the former was used.

IV. EXPERIMENTS

A.   System Overview and Scenario

The robot platform shown in fig. 10 was used for this work. The robot is equipped with several sensors including encoders, stereo and two back-to-back laser range scanners. The robot was driven across 5 rooms covering about 20m in

distance. The objects used for the representation (*and their respective annotations as used in the following figures*) comprised of different cartons (*xerox, carton, logitech, elrob, tea*), a chair (*chair*), a mug (*mug*), a shelf (*shelf*), a table (*table*) & a book (*book*).

Fig. 10 The robot platform that was used for the experiments. The encoders, stereo vision system and laser scanners were used for this work.
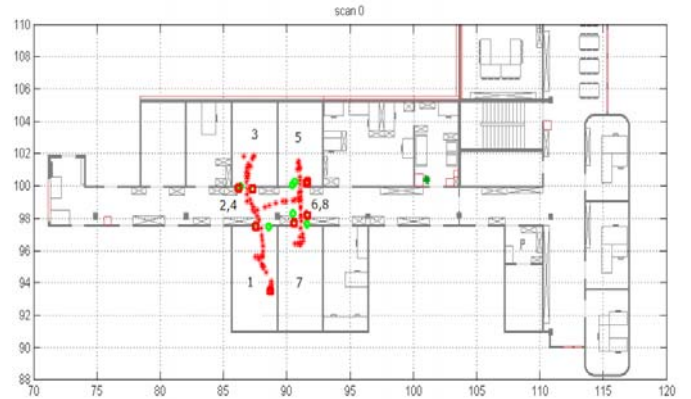

Fig. 11 Map displaying the robot path. The robot traverses through 4 rooms crossing a corridor each time it moves from one room to another. Green/red circles indicate the doors detected. The red circles also serve as the place references for the place explored on crossing the door. The numbers indicate the sequence in which the places were visited.
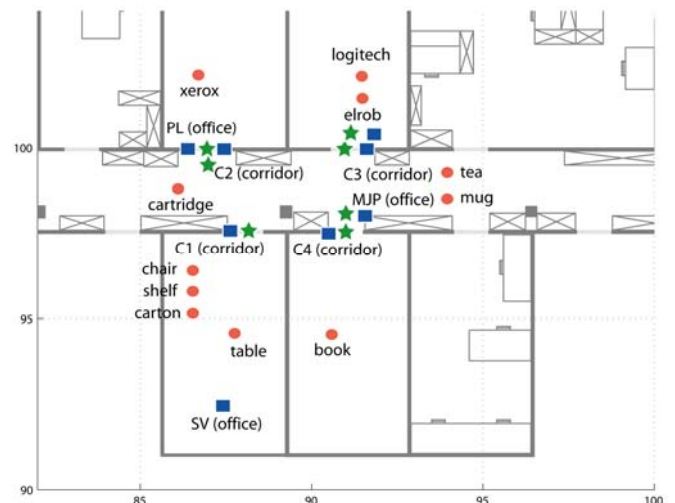

Fig. 12 Object based map produced as a result of exploring the test environment. Zoomed-in view of the above map. Blue squares are the place references, red circles are the objects and the green stars are the doors.
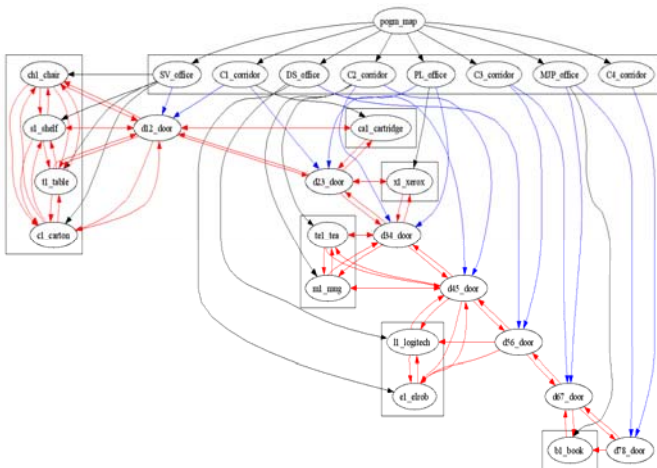
Fig.13 Probabilistic object graph representation created as a result of exploring the path shown in fig. 11
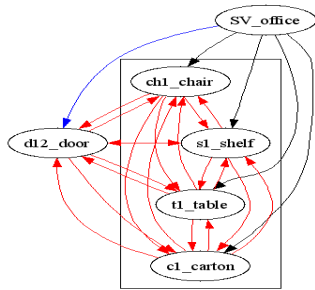

Fig. 14 Small section of the complete representation shown in Figure 13 for ease of comprehension.

B. Mapping

Fig. 11 shows the path of the robot. The objects and doors recognized are shown in the object based map depicted in fig. 12. Finally, fig. 13 illustrates the complete probabilistic object-graph representation formed as a result of the process. For the sake of clarity, a small portion of the same figure is depicted in figure 14.

The robot performed the mapping process as per expectations. Objects and doors were recognized and the representation was formed as per the methods described in the previous sections. However, the robot often observed multiple doors at the same place (due to the presence of large cupboards) on either side of the door. Further, the robot created multiple occurrences of the corridor as, the topological information between places that is encoded was not used in the experiments in this work. Also, it did not see an identical set of objects through the corridor so as to be able to recognize the previously visited corridor. These two issues (fusing of doors and loop closing) would be addressed in subsequent works.

C. Spatial Cognition – Place classification/recognition

The robot was made to traverse two previously visited places along the path shown in figure 15. The locations of movable objects (all but the table, shelf and the door) were changed so that a significant configuration change of both places was observed. The robot was then made to interpret these places.

For the first place, the robot perceived the objects in the sequence *shelf – xerox – carton – table – logitech – cartridge*. Fig. 16 displays the object map for the "unknown" place. On seeing the first two objects, the robot successfully classified the place as an office. Subsequently the robot attempted to match this place with its knowledge of prior offices it has visited. When finally crossing the door, the robot found enough objects (including the door) that are located in a matching spatial configuration to a place that it has visited before. Thus, at this point, the "unknown" place was recognized as the place SV (office) and the internal map representation of the robot is updated to reflect the changes to the place that the robot had perceived. Figure 17 displays the updated internal representation of the robot.

For the second place, the robot perceived the objects in the following sequence: *door – book – cartridge – elrob - mug – tea*. The robot managed to correctly classify the place as being a corridor but it could not recognize it (as the basic map that was created already has a fault that multiple instances of the corridor are represented and also due to insufficient data). Under these circumstances however, the robot would be expected to continue exploration until it crosses over to a new place, after which it would add a new node of the particular place category to its internal representation. The unknown place in this case is shown in figure 18.

Fig. 15 The robot path for the experiment on reasoning about places. The robot traversed the path shown above. The configuration of the objects within the place was significantly changed from that in Fig. 12. The numbers indicate the sequence in which the places were visited.
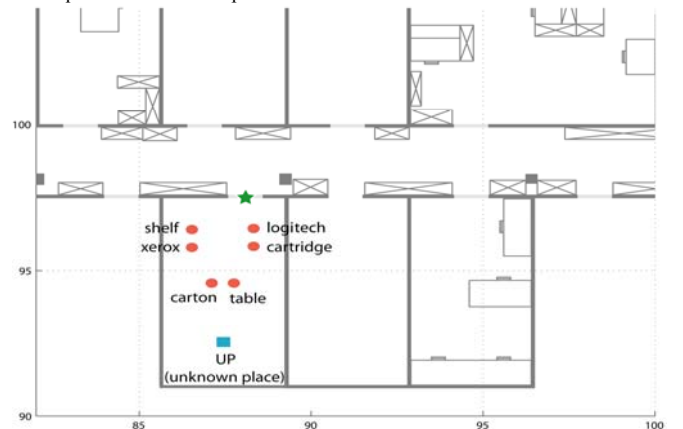

Fig. 16 First 'unknown' place at the time of place recognition. The configuration of the objects is different from that of the same place in Fig. 12. Note - The *carton* is above the table & *xerox* is above the shelf.
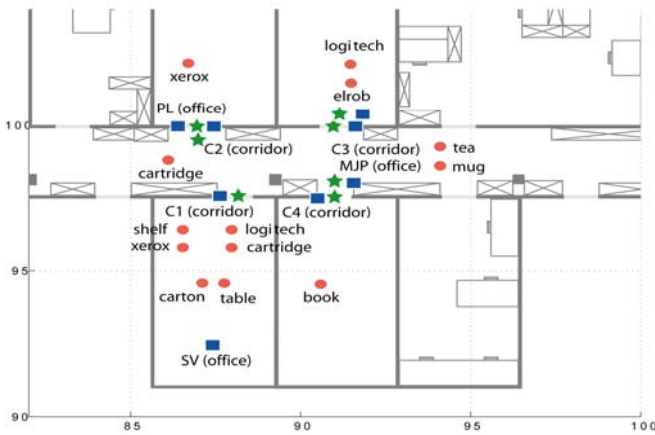
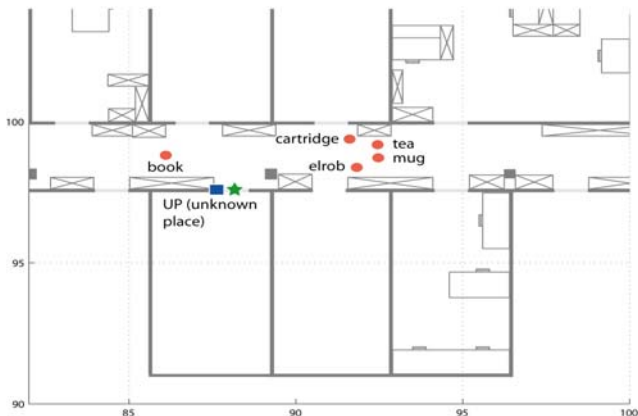Fig. 17 Updated internal representation of the robot after place recognition

Fig. 18 Second "unknown" place (configuration of objects changed). The place is correctly classified as a corridor.

V. CONCLUSIONS & FUTURE WORK

A cognitive probabilistic representation of space based on high level features was proposed. Two scientific questions were addressed – (1) How can a robot form a high level probabilistic representation of space? and (2) How can a robot understand and reason about a place? The uncertainty for all required aspects of such a representation were appropriately represented and treated. Spatial cognition using such a representation was shown through experiments on place classification and place recognition.

Fusing of doors and merging of places are both required to get a more appropriate representation of space. On the conceptual front, the suggested representation needs to be made richer but yet lighter and computationally efficient in applications.

REFERENCES

[1] S. Thrun, "Robotic mapping: A survey", In G. Lakemeyer and B. Nebel, editors, Exploring Artificial Intelligence in the New Millenium. Morgan Kaufmann, 2002.
[2] W. Burgard, D. Fox, H. Jans, C. Matenar, and S. Thrun, "Sonar-Based Mapping of Large-Scale Mobile Robot Environments using EM", In the proceedings of the International Conference on Machine Learning (ICML) 1999.
[3] K.O. Arras, Ph.D. Thesis, "Feature-Based Robot Navigation in Known and Unknown Environments", Thèse N° 2765 (2003), Swiss Federal Institute of Technology, Lausanne.
[4] A. Martinelli, A. Svensson, N. Tomatis and R. Siegwart, "SLAM Based on Quantities Invariant of the Robot's Configuration", IFAC Symposium on Intelligent Autonomous Vehicles, 2004.
[5] H. Choset and K. Nagatani, "Topological simultaneous localization and mapping (SLAM): toward exact localization without explicit localization", IEEE Transactions on Robotics and Automation, Volume: 17 Issue: 2, Apr 2001.
[6] A. Tapus, Ph.D. Thesis, "Topological SLAM : simultaneous localization and mapping with fingerprints of places", Thèse N°3357 (2005), Swiss Federal Institute of Technology, Lausanne.
[7] S. Thrun, "Learning metric-topological maps for indoor mobile robot navigation", Artificial Intelligence, Volume 99, Issue 1, February 1998, Pages 21-71.
[8] N. Tomatis, I. Nourbakhsh and R. Siegwart, "Hybrid Simultaneous Localization and Map Building: A Natural Integration of Topological and Metric", Robotics and Autonomous Systems, 44, 3-14., July 2003.
[9] A. Martinelli, A. Tapus, K.O. Arras and R. Siegwart, "Multi-resolution SLAM for Real World Navigation", In Proceedings of the 11th International Symposium of Robotics Research, Siena, Italy, 2003.
[10] B. Kuipers, "The Spatial Semantic Hierarchy", Artificial Intelligence, 119: 191-233, May 2000.
[11] C. Galindo, A. Saffiotti, S. Coradeschi, P. Buschka, J.A. Fernández-Madrigal and J. González, "Multi-Hierarchical Semantic Maps for Mobile Robotics", Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS), pp. 3492-3497. Edmonton, CA, 2005.
[12] D. G. Lowe, "Distinctive image features from scale-invariant key-points", International Journal of Computer Vision, vol. 60, no. 2, 2004.
[13] D. Anguelov, D. Koller, E. Parker, S. Thrun, "Detecting and Modeling Doors with Mobile Robots", Proceedings of the International Conference on Robotics and Automation (ICRA), 2004.
[14] D. Kortenkamp, L. D. Baker, and T. Weymouth, "Using Gateways to Build a Route Map," Proceedings IEEE/RSJ International Conference on Intelligent Robots and Systems, 1992.
[15] C. Stachniss, Ó. Martínez-Mozos, A. Rottmann, and W. Burgard, "Semantic Labeling of Places", In Proc. of the Int. Symposium of Robotics Research (ISRR).San Francisco, CA, USA, 2005.
[16] A. Torralba, K. P. Murphy, W. T. Freeman and M. A. Rubin, "Context-based vision system for place and object recognition", Proceedings of the IEEE International Conference on Computer Vision, ICCV 2003, vol.1, p.273. Nice, France.
[17] B. J. Kuipers, "The cognitive map: Could it have been any other way?", Spatial Orientation: Theory, Research, and Application. New York: Plenum Press, 1983, pages 345-359.
[18] W.K. Yeap and M.E. Jefferies, "On early cognitive mapping", Spatial Cognition and Computation 2(2) 85-116, 2001.
[19] V.T. Nguyen, A. Martinelli, N. Tomatis and R. Siegwart, "A Comparison of Line Extraction Algorithms using 2D Laser Rangefinder for Indoor Mobile Robotics". In Proceedings of the IEEE/RSJ Intenational Conference on Intelligent Robots and Systems, IROS, Edmonton, Canada.
[20] A. Tapus, S. Vasudevan, and R. Siegwart, "Towards a Multilevel Cognitive Probabilistic Representation of Space", In Proceedings of the International Conference on Human Vision and Electronic Imaging X, part of the IS&T/SPIE Symposium on Electronic Imaging 2005, 16-20 January 2005, CA, USA.
[21] I.K. Jung, Ph.D. Thesis, "SLAM in 3D environments using stereo vision", Thèse No 04132 (2004), Institut National Polytechnique, Toulouse, France.