Conference Paper

Towards a Multilevel Cognitive Probabilistic Representation of Space

Author(s):
Tapus, Adriana; Vasudevan, Shrihari; Siegwart, Roland

Publication Date:
2005

Permanent Link:
https://doi.org/10.3929/ethz-a-010085160 →

Rights / License:

Towards a Multilevel Cognitive Probabilistic Representation of Space

Adriana Tapus, Shrihari Vasudevan and Roland Siegwart

Ecole Polytechnique Federale de Lausanne (EPFL)
Autonomous Systems Lab
1015 Lausanne, Switzerland
{Adriana.Tapus, Shrihari.Vasudevan, Roland.Siegwart}@epfl.ch

ABSTRACT

This paper addresses the problem of perception and representation of space for a mobile agent. A probabilistic hierarchical framework is suggested as a solution to this problem. The method proposed is a combination of probabilistic belief with "Object Graph Models" (OGM). The world is viewed from a topological optic, in terms of objects and relationships between them. The hierarchical representation that we propose permits an efficient and reliable modeling of the information that the mobile agent would perceive from its environment. The integration of both navigational and interactional capabilities through efficient representation is also addressed. Experiments on a set of images taken from the real world that validate the approach are reported. This framework draws on the general understanding of human cognition and perception and contributes towards the overall efforts to build cognitive robot companions.

Keywords: cognitive mapping, hierarchical topological representation, probabilistic belief

1. INTRODUCTION

Interpreting and understanding a scene from the environment beyond single object recognition is a hard task. Humans use various sensory cues to extract crucial information from the environment. This is processed in the cortex of the brain in order to obtain a high-level representation of what has been perceived. Intuitively, it appears that humans represent knowledge in a hierarchical fashion. With a view of having robots as companion of humans, we are motivated towards developing a knowledge representation system along the lines of what we know about us. While recent research has shown interesting results, we are still far from having concepts and algorithms that interpret space, coping with the complexity of the environment.

Most of the related research on formalizing levels of abstraction in literature can be found in cognitive science (e.g. hierarchical representation and reasoning with knowledge), mobile robotics and networking. Space representation plays an important role in any cognitive and autonomous intelligent system. The idea of cognitive maps (i.e. the animal internal representation of space) was introduced for the first time by Tolman in [17]. Significant progress has been made since the seminal papers by Kuipers [7, 8] where the cognitive maps are described as a body of knowledge representing large scale space. In his work, a spatial semantic hierarchy is suggested, which represents space at different levels of abstraction and attempts navigation using such a representation (see [9]). An approach similar to the previous one can be found in [10] where a hierarchical multi-resolution space representation is addressed. Voicu uses landmarks and associations between them to construct a cognitive map of a large environment in [19]. The information from this cognitive map is then used for path planning and exploration. The authors of [14] use a hierarchical hidden Markov model (HHMM) to learn the route between two labs. The higher level states are the more abstract/distinct ones like corners and intersections. The lower level states represent intermediate positions. All the above mentioned works seem to capture a hierarchical representation of space with a navigational flavor in them.

The work done by Brezetz et al in [1] bears close resemblance to our work presented here. It assumes the presence of an even ground with objects on it. Range images are obtained and segmentation of these images yields semantic information contained in them. The information is further used to extract spatial relationships between objects. Motivation for this work has also originated from beyond the robotics community. The problem being addressed in

this paper has found great interest in the Computer Science and Geographical Information Systems communities as well. The work by Papadias and Theodoridis in [13] suggests different topological and directional relationships between objects based on the concept of minimum bounding rectangles (MBR's). They make use of "R-Trees" to represent information. A different variation of the "R-Tree" was conceived by the authors of this work as well. However graphs were found to be more general, more suited to our long term goals and more extensible. This explains the reasoning behind our work. Grini et al have done something similar in [4]. They however address the problem of deriving more complex and sound inferences and look into complexity issues regarding it. They also visualize a multi-resolution framework of various spatial relationships. The work in [5] explores similar concepts in a different context – content based image analysis. They use absolute and relative spatial relations to describe inter-relationships between objects. They also compare the spatial-relationship identification between objects, when performed by humans with the software tool they developed for it. This helped to identify what spatial relationships should be present and also points out the limitations that one would face (unable to describe many relationships) when using 2D image processing as compared to 3D image processing – something we also noticed.

The remainder of this paper is organized as follows. We present in Section 2, the way objects from the environment are modeled and the method used for performing the recognition. In Section 3, a framework to represent objects occurring in the environment and inter-relationships between them, as a topological map, is suggested. Experimental results are presented in Section 4. Section 5 concludes the paper with a discussion of the proposed approach and further research directions.

2. OBJECT MODELING AND RECOGNITION

Modeling the physical world in terms of the objects present in it and the way they relate to each other is one highly intuitive method of interpreting space. Both spatial and semantic inferences are demonstrated in this work, from the model thus created. In this context, object modeling and recognition capabilities together with methods to detect spatial and semantic relationships between objects are required. A simple probabilistic feature based object recognition system has been used here. The system has been kept simple in terms of the features used. This work demonstrates various concepts using single colored boxes in a non-cluttered setting. It must be emphasized that object modeling and recognition are not the themes of this paper. However, they are critical to realizing the overall concept presented here.

2.1. Object Modeling

Figure 1 depicts different low-level features that can be extracted from the visual information used to model the objects. Broadly, features may be classified as color, shape, edges, etc. and several methods can be employed to represent them. In our system, we used only color and shape, thereby keeping it in very simple. More details about the extraction of low-level visual features can be found in [16].
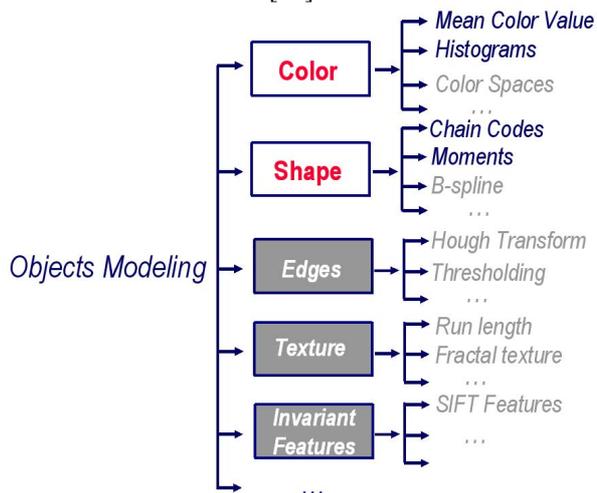

Figure 1: Objects Modeling: low-level features representation

A brief description of the features employed in our approach is given below:

A. Mean Color Value

The Hue Saturation Value (HSV) color space was used. The hue represents the color, the saturation represents the "purity" of the color and the value denotes the total power of the spectrum. All three values are normalized to the (0, 1) range for ease of use. The HSV color space is obtained by a non-linear transformation of the RGB color space. This color space was preferred due to its intuitive appeal over most other standard color spaces. The means of the hue, saturation and value components were computed.

B. RG-Chromaticity histograms

The RG Chromaticity Space is a two dimensional color space with no color intensity information. Each pixel represents the contribution of the red and green components. It is derived from the RGB color space as follows:

$$r = \frac{R}{(R + G + B)} \tag{1}$$

$$g = \frac{G}{(R + G + B)} \tag{2}$$

As each component of the RGB is normalized, the blue contribution can be easily computed, whenever required, as

$$b = 1 - (r + g) \tag{3}$$

The advantage of this space is that changing light intensities in the environment will not cause a change in the basic color of the object. A histogram with 100 bins for each component (i.e. r, g and b) of the image was formed.

C. Freeman's Codes (Chain Codes)

The shape of an object is detected by using the Freeman's code (chain code) [3] method as illustrated in [16]. The chain code models an object shape as a series of directed unit-pixel line segments. The result is a sequence of symbols that represents the inner boundary of the object/region under consideration.

D. Moments

Moment descriptions of regions describe a normalized gray level/binary image as a probability density of a 2 dimensional random variable. The moments are computed as shown in [16]. They are based on the principle of the moments of order (p+q) given by the formula:

$$m_{pq} = \sum_{i=-\infty}^{\infty} \sum_{j=-\infty}^{\infty} i^p \ j^q \ f(i, j) \tag{4}$$

This is translation, rotation and scaling dependant. Translation invariance is obtained by using centralized moments. To achieve scaling invariance, scaled central moments are introduced and rotation invariance is incorporated by choosing the coordinate axes such that the (1+1)th order central moment is zero. Essentially, a group of seven translation, rotation and scaling invariant moments as discussed in [6] were used.

2.2. Object Recognition

As mentioned previously, object recognition is realized as a probabilistic feature matching process. The recognition methodology can be understood as an essentially color-based one. Simple polyhedral and single colored objects in uncluttered environments are used to demonstrate the concept proposed here. Prior to the recognition process, a database of object (the model objects) features is constructed. The color and the shape features are extracted and stored using the representations described above. The recognition process proceeds in the following manner:

- **Color Thresholding**
- **Feature Extraction**

- **Feature Matching** – The extracted features are compared with those of the object models in the database. Dissimilarity measures are computed. These are used to arrive at a belief measure of how closely the region resembles the object with which it is compared, on the basis of the particular feature.
- **Belief Computation** – Probabilities obtained from the feature matching process are combined to produce a single "belief" value that represents the degree of similarity between the region and the object (model) under consideration. Similar measures for other object models are found. The region is then inferred to be an object of the type represented by the object model for which the highest belief measure is obtained.

3. OBJECT GRAPH MODELS: CONSTRUCTION, UPDATING AND INFERENCE

The main significance of this work is to establish a correspondence between animals (humans) and robots. By understanding how animals (humans) navigate and build their own spatial representation, the observed phenomena are applied in robotics. In order to have a robust and reliable framework for navigation (i.e. in order to move within an environment, manipulate objects in it, avoid undesirable mishaps (collisions), etc.) space representation, perception, localization and mapping are all needed. The work of Tolman [17] showed that space is modeled as mental or cognitive maps. Tolman's model advocates that animals (rats) don't learn space as a sequence of movements; instead, the animal's spatial capabilities rest on the construction of maps which represent spatial relationships between various objects encountered in the environment. This has led to the concept of topological representation of space. Later, O'Keefe and Dostrovsky [11] discovered the hippocampal place cells (i.e. cells whose firing pattern is dependent on the location of the animal in the environment), which led to the idea that the hippocampus works as a cognitive map of space [12].

Next, we discuss a framework for describing an approach similar to the phenomena encountered in the hippocampal place cells and perirhinal cortex (i.e. part of the brain where the recognition of objects is performed).

3.1. OGM Construction and Update

Physical space is viewed in terms of objects and relationships between them. Thus, the environment is represented as a graph structure – which is referred to as an Object Graph Model (OGM). Knowledge representation using OGM involves: object recognition (described in section 2.2) and graph representation. After performing the object recognition step, a set of objects and object descriptions are obtained. These are used in order to model/describe a scene from the environment by using a graph representation. The existence of the object has an associated belief, which is computed from the beliefs of its features (i.e. from object recognition step). Object descriptions include color, shape, centroid coordinates, bounding boxes, and so on. Each object is represented as a node in the graph. The object descriptions are used to check for the existence of every possible relationship between each pair of objects (using binary relationships). Identified relationships are represented as directed arrows between two nodes. Each relationship also has an associated belief which depends on how well (unambiguously) the two objects satisfy the criteria that define the existence of a particular relationship apart from the beliefs in the existence of the objects involved. The resulting graph structure, so formed, is what we refer to as the Object Graph Model (OGM).

The OGM's previously formed must be updated with change in scene description. Three kinds of update operations are requisite:

A. New OGM

Each new scene observed is compared with all previously modeled scenes which have been modeled using the OGM. There is a data association problem involved here, in identifying two corresponding scenes. Two metrics are used to approach this, (1) the number of objects the two scenes share in common and (2) the number of these objects obeying the same relationships (on the vertical plane only, as horizontal plane relationships may change). Each of these metrics produces a belief measure, which are in turn combined. The resulting belief is a measure of how similar the current scene is, with respect to a previously modeled scene. Depending on this measure and a threshold that is preset, either a new OGM is constructed for the current scene or the previously modeled scene is updated as explained below.

B. Belief Update

When objects are repeatedly viewed across scenes (i.e. different views which correspond to the same modeled scene), the belief in their existence should increase. On the other hand, the belief in the existence of objects, not observed across scenes, should decrease. This is similar to the approach adopted in [18]. The same trend should also be applicable for relationships between objects. A data association problem involved in identifying

"corresponding" objects between a previously modeled scene and the scene under consideration is present here. This is overcome by taking a simple experimental scenario wherein two objects do not look exactly alike (e.g. they differ in either shape / color).

C. Add/Delete objects

When a new scene corresponds to a previously modeled scene (i.e. thus already existing as an OGM), a provision is required by which the new objects in the current scene are incorporated in the OGM while un-observed objects have their beliefs reduced. If their belief goes below a certain threshold – they are removed from the OGM as the belief in their existence (i.e. in that scene) has become negligibly small.

3.2. Example – OGM Representation of a scene

Let us illustrate the construction of an OGM with an example (see Figure 2). Figure 2 a) illustrates a possible scene in the environment and it is composed of seven objects. These are: one table, three mugs represented with the orange color and by the symbol m_i (i =1,2,3) and three green boxes of different sizes, denoted by b_j (j =1,2,3). The relative positions (e.g. above/below etc.) between the objects has been used to represent the relationships between them. In the topological framework, the objects are the nodes and the relationships between the objects are the edges. Thus the OGM of the current scene can be constructed as shown in Figure 2 b).
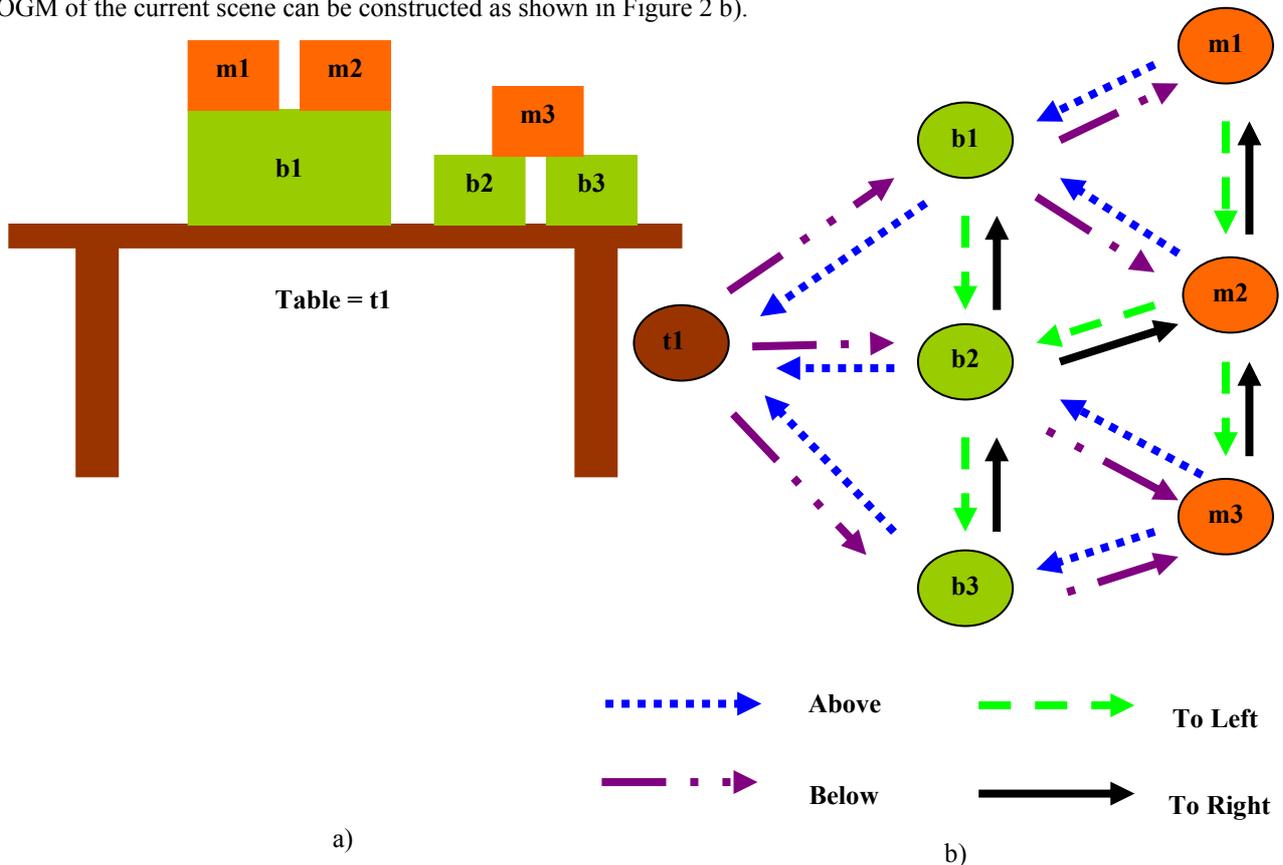

a)

b)

Figure 2: Scenario Representation. a) Table (t1) with Boxes (bi) and Mugs (mi) on it; b) Object Graph Model (OGM) for the example shown in Figure a).

3.3. Spatial Relations

Spatial relations are required in order to make meaningful and useful inferences about the objects in the scene. This work uses both directional and topological relationships between objects. A more comprehensive repertoire of relationships is possible if the image processing and object recognition are done in 3D. Directional relationships include those like the North (N), South (S), East (E) and West (W). Egenhofer, in [2], points out the existence of as many as nine meaningful topological relationships between objects. The use of 2D image processing limited us to use

only those like overlap and disjoint. It must be noted that horizontal plane relations are relative to the line of sight and hence cannot be relied on, to hold good, for more than the scene in which they are observed. Vertical plane relations are more persistent in that the robot (which models the environment) is assumed not move in the vertical plane. Relationships between objects also have associated beliefs. These beliefs are computed by taking into account the probability of the relationship holding between the two objects given their features and the probabilities of the existence of each of the objects. It may be computed using the following expression:

$$P(R_{A,B}) = P(R \mid A, B) \ P(A) \ P(B) \hspace{2cm} (5)$$

where R is a binary relationship and A and B are the objects under consideration. The belief value is intuitive in that the belief in the presence of a relationship between two objects also takes into account the "clarity of the evidence" of the relationship existing between two objects.

3.4. Inference

The inference methods used here are based on binary relationships between objects. When objects are recognized, their bounding boxes are used to compute the existence and belief of these relationships. The following types of inference have been conceived in the system: purely spatial, purely semantic, spatio-semantic and multi OGM. Purely spatial inference tests for specific/all possible spatial relationships between objects (e.g. find all objects related to object *A* or how are objects *A* and *B* related). Semantic inference is realized by making queries on properties of the objects like shape and color (e.g. find all objects of a given color / shape). A hybrid inference, the spatio-semantic one, is obtained by combining the previous two types queries (e.g. find all existing relationships between all *red* colored objects in a scene). Finally, multi OGM inference makes queries across scenes (i.e. each scene being modeled by a separate OGM) (e.g. find all *red* colored objects between two different scenes). In our system, inference is realized by looking up an already computed exhaustive set of relationships for a particular one – this is in effect a graph search. More complex inference may be obtained by "chaining" inference results to obtain higher level complex results of greater use in practical situations.

4. EXPERIMENTAL RESULTS

For the experiments, a digital firewire color camera SONY DFW-VL500 has been used. The camera has a resolution of 640x480 pixels at 25Hz. The camera is mounted horizontally on the BIBA robot, a fully autonomous mobile robot.

4.1. Object Recognition

The test setup was the following: The training set included a single image of each of the possible objects used for the experiments. The robot extracted the four features from a test set of 15 images. The images have been taken from different angles and the distance from the robot to the objects was almost the same throughout the experiment. This simplification could be omitted by using scale invariant features. The recognition output was classified into the following categories: object identified correctly, data association problem between objects of the same color, false positives, merging of objects and completely missing the presence of an object. As pointed out in section 2, we extracted only low-level visual features to represent an object just for simplicity of the model, so that the OGM and related concepts presented in this paper can be quickly demonstrated.

Before giving a detailed description of the results and their implications, we will use an example to explain the different situations that can be obtained (see Figure 3):
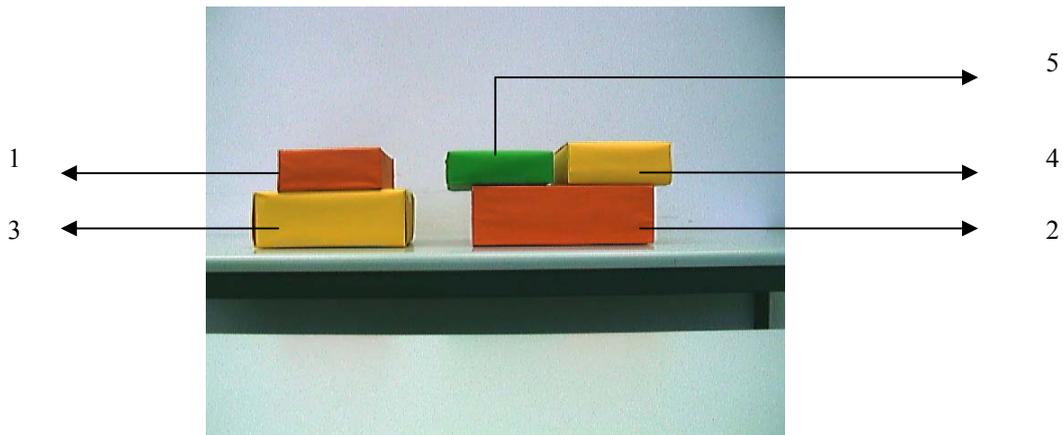
Figure 3: Test image composed of 5 different objects (annotated from 1 to 5). The different boxes can be associated to real world objects. For instance, the green box could correspond to a can of coke, the yellow box to a bottle of orange juice, etc.

The image depicted in Figure 3 contains five objects annotated from 1 to 5. As illustrated in Table I, all 5 objects were recognized as expected. One object was identified by its color correctly, but it was confused with another object of the same color (i.e. the color association was perfect but the correct object association was inaccurate in this case). Thus the recognition rate was 90%. There were no false positives (i.e. existence of a object when it isn't actually the case) observed, no merging problems between objects and no objects were completely missed.

TABLE I. OBJECTS CLASSIFICATION FOR THE EXAMPLE SHOWN IN FIGURE 3

Results	ID Objects				
Expected recognition output	1	2	3	4	5
Observed recognition output	2	2	3	4	5

Many statistics have been collected for the test set of 15 images. Table II shows how many cases of false positives, merged objects and missed objects were encountered. A good result that must be underlined is that no false positives were detected. Another outcome that was checked for was the case of merged objects. From the entire experiment, only three occurrences of this were found. This was due to an inaccuracy in the color thresholds. With regards to completely missing objects, our system did well in that only one object in one case was overlooked. This was caused by an illumination change, too pronounced for our system to handle. Table II thus draws out a lot of positive aspects about our simple object recognition mechanism.

TABLE II. MEASURE ON FALSE POSITIVE, MERGED OBJECTS AND MISSED OBJECTS CASES FOR THE SET OF 15 IMAGES.

Cases of False Positives	Cases of Merged objects	Cases of overlooking of objects
None	3	1

The table below (Table III) shows the results obtained from the object recognition module for 15 different images. The image annotation (i.e. the id corresponding to each object type) is exactly the same as shown in Figure 3. For the first few images, two objects of type 1 were used. The sequences (both observed and expected ones) are formed by listing the objects in the order of its color – first, the orange ones, followed by the yellow and finally the green ones. The table points out several interesting facts. Firstly, most of the objects were recognized in each image and most of the identified objects were in turn identified correctly. A correct identification corresponds to correct object association. There are however cases wherein object color is identified correctly but it is confused with other objects of the same color. In Table III, all cases wherein the number of observed objects is less than the number of expected objects, one of the three errors detailed in Table II occurred.

TABLE III. STATISTICS PERFORMED ON THE SET OF 15 IMAGES.

Image ID	Output Observed	Expected Output	Number of Objects expected	Number of objects observed	Number of correctly identified objects	Number of objects : color identified correctly but object confused with another of same color
0	122335	112345	6	6	4	2
1	135	245	3	3	1	2
2	1435	12345	5	4	1	3
3	12345	112345	6	5	4	0
4	1235	11235	5	4	3	0
5	222345	112345	6	6	4	2
6	1345	112345	6	4	3	0
7	212335	211435	6	6	4	2
8	14	24	2	2	1	1
9	24	24	2	2	2	0
10	22345	12345	5	5	4	1
11	1435	2345	4	4	1	3
12	23	13	2	2	1	1
13	235	235	3	3	3	0
14	2235	1235	4	4	3	1

The positive feature about the method is that in most cases, the number of objects correctly identified exceeds the number of objects which were identified correctly in terms of color.

4.2. Belief Update

The belief update method has also been tested. Its performance however depends heavily on the object recognition module. Figure 4 shows three images (i.e. image id = 10, 11, 12) that represent three different scenes in the environment. These images have used to demonstrate the updating concept. In order to correctly demonstrate the belief update procedure, the object association errors in the output of the recognition module have been corrected (i.e. as shown in Table III, in all three cases, there are objects whose color has been correctly identified, but they have been confused with other objects of the same color. These errors have been corrected to the expected outputs in order to facilitate the demonstration of the concept). Next, we will show the overall mechanism used for the belief update process.

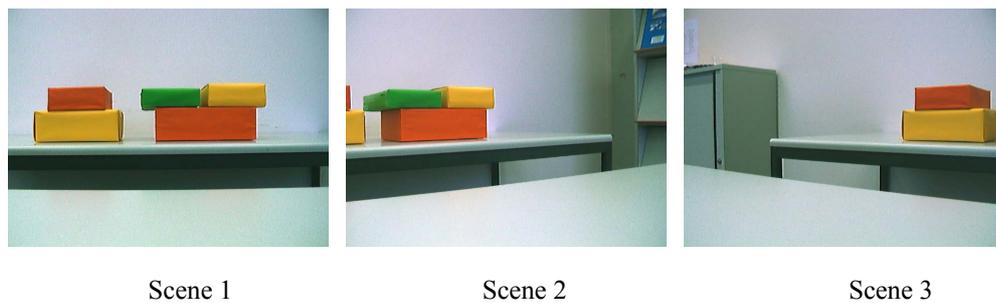Scene 1   Scene 2   Scene 3

Figure 4: Three scenes from the environment used to show the updating process. The ids of the objects are the same as the ones depicted in Figure 3.

In the beginning, scene 1 is processed by the object recognition module. All objects and their interrelationships are identified. This information is then used to construct the OGM for scene 1. Identified objects, their properties, interrelationships between them, and the beliefs in each of these were stored in the OGM. Figure 5 shows the partial OGM formed with respect to object 2.
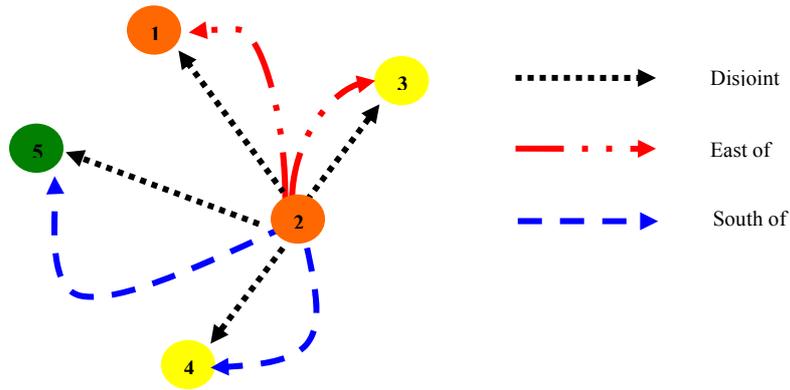
Figure 5: Partial OGM for the object 2 in scene 1

The other possible inter-relationships between objects are not shown in order to preserve clarity in the figure. After processing scene 2, its OGM is formed in a similar fashion. The two OGM's are compared to determine if they represent the same scene. The resulting comparison factor was smaller than the preset threshold and thus the OGM of scene 2 is now used to update the OGM of scene 1 in order to form a single, consistent and up-to-date representation. All objects and relationships that are also observed in scene 2 (in addition to being observed in scene 1) have their beliefs in the OGM representation of scene 1 updated with the current belief. Conversely, un-observed objects and relationships have their beliefs reduced in the OGM of scene 1. For e.g., the belief in object 2 and its relationship with object 4 are both updated with their current beliefs in scene 2. On the other hand, the belief in object 2's relationship with object 1 is reduced as object 1 is not anymore observed in scene 2. Another important OGM update operation is the removal of objects and relationships when their beliefs fall below a certain threshold. This is intuitive in that the belief is a measure of the existence of an entity. In the example considered here, object 3's relationships with object 4 are removed when scene 2 updates the OGM for scene 1, as the belief in their existence becomes negligibly small. Scene 3 is processed in a similar way. It is compared with the previously modeled scene (scene 1 updated with scene 2) and deemed to match it. Thus its OGM is used to update the previously accumulated representation. As object 2 is not observed in scene 3, the belief in object 2 and that of all of its relationships is reduced after scene 3 updates the accumulated representation. Thus the belief update procedure proceeds exactly as per expectations.

This example clearly points out that the belief update mechanism is very easy to apply, very intuitive, permits modeling the dynamics of the environment and yet produces a single consistent and up-to-date representation of the environment as perceived by the robot.

5. CONCLUSIONS AND FUTURE WORK

This paper has presented a method for topological space representation using a combination of probabilistic belief with object graph models (OGM). This kind of representation permits efficient modeling of the information that a robot would perceive from its environment. The larger goal of the work presented here is to address the integration of navigational and interactional capabilities of robots through an efficient unified multi-level representation of space. This work puts forward a general, graph-based, probabilistic representation of the environment and is meant to contribute towards the larger goal. This directly relates to the research area of Human Perception and Cognition. The concept was validated through experiments that showed very encouraging results. Some future works will concentrate on developing a method for color constancy, based on a PID controller that will stabilize the mean luminance, and on the improvement of the object recognition and modeling mechanism. Use of invariant features like the SIFT and color invariant features seems to be a promising option towards robust object recognition. 3D object modeling will also be required in order to have effective models for interaction. In addition, efforts will also be directed towards extending this work into a multi-level cognitive probabilistic representation with multi-modal data suited for both navigation and interaction. Incorporating more geometric, semantic and functional information about objects into the OGM's will be

required for higher level inference, better navigational and interactional capabilities. While adding multimodal data, sound uncertainty models for each of the inputs will also be required in order to effectively handle the uncertainty in the environment. In summary, to have real cognitive agents using this very promising concept, the implementation must be scaled up in all facets, it has to be subjected to an exhaustive set of test scenarios and complexity issues must be dealt with to ensure scalability.

ACKNOWLEDGMENTS

REFERENCES

1. Brezetz, S. B., Chatila R. Devy M., (1994), Natural scene understanding for mobile robot navigation, In the Proceedings of the IEEE International Conference on Robotics and Automation, San Diego, USA

2. Egenhofer, M.,J., Reasoning about binary topological relations. In Gunther, O. and Schek, H.J. (eds.), Advances in Spatial Databases, SSD'91 Proceedings, Springer Verlag 143-160, 1991.

3. Freeman H., (1961), On the encoding of arbitrary geometric configuration, IRE Transactions on Electronic Computers, EC-10(2):260 – 268

4. Grigni, M., Papadias, D., Papadimitriou, C. Topological Inference. Proceedings of the International Joint Conference of Artificial Intelligence (IJCAI), Montreal, Canada, AAAI Press, 1995.

5. Hollink L., Nguyen G., Schreiber G., Wielemaker J., Wielinga B. and Worring M.., Adding Spatial Semantics to Image Annotations. 4th International Workshop on Knowledge Markup and Semantic Annotation at ISWC'04

6. Hu M.K.,(1962), Visual Pattern Recongnition by moment invariants, IRE Transactions on Information Theory, 8(2):179-187

7. Kuipers, B. J. (1978), Modeling Spatial Knowledge, Cognitive Science, 2: 129-153, 1978.

8. Kuipers, B. J. (1983), The Cognitive Map: Could it have been any other way?, In Spatial Orientation: Theory, Research and Application. Picks H.L. and Acredolo L.P. (eds.), New York. Plenum Press.

9. Kuipers, B. J. (1996), A Hierarchy of qualitative representations for space, In Proceedings of the 10th International Workshop on Qualitative Reasoning (QR-96), Fallen Leaf Lake, California, USA.

10. Martinelli, A, Tapus A, Arras, K.O., and Siegwart. R. (2003), Multi-resolution SLAM for Real World Navigation, In Proceedings of the 11th International Symposium on Research Robotics, Siena, Italy.

11. O'Keefe, J., and Dostrovsky, J., The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. Brain Res. 34, 171-175, 1971.

12. O'Keefe, J., and Nadel, L., The hippocampus as a cognitive map, Clarendon, Oxford, 1978.

13. Papadias, D., Theodoridis, Y. Spatial Relations, Minimum Bounding Rectangles, and Spatial Data Structures. International Journal of Geographic Information Science, vol l.11(2), pp. 111-138, 1997

14. Rohanimanesh K., Theocharous G. and Mahadevan S. (2000), Hierarchical Map learning for Robot Navigation, In AIPS Workshop on Decision-Theoretic Planning, Breckenridge, Colorado.

15. Shih H.C. and Huang C.L. (2003), A Semantic Network modeling for understanding baseball video, In Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing, Hong Kong, China.

16. Sonka M., Hlavac V. & Boyle R., (1998), Image Processing, Analysis and Machine Vision, Brooks/Cole Publishing Company

17. Tolman, E. C. (1948), Cognitive maps in rats and men, Psychological Review, 55:189-208.

18. Tomatis, N., I. Nourbakhsh, and R. Siegwart (2003). Hybrid simultaneous localization and map building: a natural integrationof topological and metric. Robotics and Autonomous Systems, 44:3-14.

19. Voicu H. (2003), Building and Using a Hierarchical Representation of Space, In Proceedings of the International Joint Conference on Neural Networks (IJCNN), Portland, Oregon, USA.