A Bayesian Approach to Conceptualization and Place Classification: Using the Number of Occurrences of Objects to Infer Concepts

Shrihari Vasudevan, Ahad Harati and Roland Siegwart
Autonomous Systems Laboratory (ASL)
Swiss Federal Institute of Technology Zürich (ETHZ)
8092 Zürich, Switzerland.
{vasudevs, haratia, rsiegwart}@ethz.ch

***Abstract*— The future of robots, as our companions is dependent on their ability to understand, interpret and represent the environment in a human compatible manner. Towards this aim, the presented work is part of an attempt to create a hierarchical probabilistic concept-oriented representation of space, based on objects. Specifically, this work details efforts taken towards learning and generating concepts and attempts to classify places using the concepts gleaned. Inference is based on the number of occurrences of various objects. The approach is based on learning from exemplars, clustering and the use of Bayesian network classifiers. Such a conceptualization and the representation that results thereof would be useful for enabling robots to be more cognizant of their surroundings and yet, compatible to us. Experiments on conceptualization and place classification are reported. Thus, the theme of the work is - conceptualization and classification for representation and spatial cognition.**

I. Introduction

Robot mapping is a 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 [18]. Robot maps can be generally classified into three categories - metric ([2], [1]), topological ([4], [14]) and hybrid ([17], [19]). The one similarity between all these representations is that all of them are navigation-oriented. Thus, while these maps are certainly useful in getting robots to move around, they fail to encode much of the spatial semantics in the environment. This results in robots having a very modest level of spatial awareness. The focus of this work is to address this deficiency. Further, a robot may use such representations to perform spatial cognition to different extents. While (metric) localization and place recognition (is this *my* office ?) have been well explored ([1], [14] & [20]) in the research community, place classification (is this *an* office ?) is a more general problem and warrants the formation of a conceptual model of the place. The work reported here addresses this issue in the overall context of improving the semantic content of state-of-the-art robot representations.

Typically, humans seem to perceive space in terms of objects, states and descriptions, relationships etc. This seems both intuitive and is also validated through user studies that were conducted in [22]. Thus, a *cognitive* spatial representation, for a mobile robot, could be expected to encode similar information. The work reported in [21] attempted to create such a representation by encoding typical household objects and doors within a hierarchical probabilistic framework. It used a SIFT [10] based object recognition system and a door detection system based on lines extracted from range scans. It also proposed a first conceptualization of different places, based on the objects that were observed. Spatial cognition was demonstrated in two ways - place classification using the models learnt and place recognition using the probabilistic relative object graph representation (a graph encoding objects and 3D relative spatial information between them). The conceptualization and place-classification that was performed were preliminary steps in the direction. The classification was based on a very simplistic Naive Bayesian Classifier (NBC) [5] that did not learn from negative exemplars. The likelihood formulation in the conceptualization was not useful for handling multiple occurrences of objects. It also did not use any explicit relationship between the number of occurrences of an object and the concept. Classification was done only on the basis of the evidence that was present and did not consider that which was absent, the latter is very significant information.

The approach presented in this work attempts to address these issues in the larger context of proposing a consistent Bayesian framework for the incorporation of spatial semantics. Further, [21] represented spatial semantics through only the presence of objects. This report aims at taking this one step further - by forming meaningful semantic concepts, based on the objects. For instance, consider a kitchen that is composed of a storage-space, a cooking-space and a dining-space, each of which are in turn composed of several objects pertinent to it. This work enables a robot exploring the kitchen to actually understand (and internally represent) that there is an area to dine, to cook and to store things in the place, and that the place is a kitchen because of this.

II. Related Work

Many works either inspire or are closely related to the work presented here. In the artificial intelligence (AI) community, the problem of *generalization* has been well addressed. The work [24] provides a good overview of different generalization strategies that exist and how they relate to each other. The approach presented in this work can be likened to a data driven approach which requires a set of positive / negative exemplars (or a "teacher") to learn from. The problem of *conceptual clustering* is another closely related and well established research area. Perhaps, the best known example of this, is the COBWEB system [6]. This system attempted to perform unsupervised incremental probabilistic conceptual clustering.

The problem, approach and the methodology of generating and using probabilities is different from that presented here. Among more recent works, the aspects dealt with in this work, bear similarities with [15]. This work presented a generative probabilistic model for classification and clustering of relational data. The model is based on previous work by the authors on Probabilistic Relational Models. The model incorporates a large set of dependencies between the latent variables representing the entities of the data; it used an approximate Expectation-Maximization algorithm to learn the parameters of the underlying model and the inference was based on Belief Propagation. Another closely related work, to that presented here, is reported in [16]. It provides a Bayesian approach to learning concepts from a few positive exemplars. The specific example demonstrated is that of learning axis-parallel rectangles in multi-dimensional space.

Recent works in robotics that are relevant in the context of this work include [7] and [12]. The former used an AI based reasoning engine that specified rules for each concept based on an ontology. The latter used the object occurrences to differentiate between similar structured rooms - this was done by integrating the object cues within an AdaBoost framework. The state-of-the-art in robot place classification typically relies on object occurrence cues, used in a logic or rule based framework, possibly with a predefined ontology. A recent contribution that works along these lines is detailed in [11]. The objective of this work is to formulate a principled Bayesian approach in order to incorporate semantic concepts in robot spatial representations and enable robots to reason about their surroundings. The scenario envisioned is that of a robot being taught different concepts by its human user.

A concept that provides for the basis of the approach presented here is that of the Bayesian network classifiers- in particular, the Naive Bayesian Classifier (NBC). It is well known that NBC's (generative classifiers) although being unarguably simplistic models that make strong assumptions, are able to successfully compete with any of the other state-of-the-art (discriminative) classifiers [13]. The work [3] gives a nice overview on the different kinds of Bayesian network classifiers that exist and also elicits on ways to learn them. The approach presented in this report also draws on the vast amount of work done in the area of *clustering*, a good survey of which is presented in [8]. Additionally, this work attempts to be fully probabilistic and is grounded on a Bayesian Programming methodology, as described in [9].

The contribution of this work is the formulation of a sound Bayesian methodology to enable a robot to conceptualize and classify its environment as it explores it. The representation that is formed as a result of the conceptualization, encodes a greater level of semantic information (concepts) than before and enables a robot to be more spatially aware of its surroundings. Also, the representation would be totally compatible with humans (demonstrated in [22]).

III. Approach

A. Overview

Figure 1 illustrates the overall approach that is being pursued. In [21], a key idea was to enhance robots spatial representation by changing the feature set from the now common lines, corners etc. to higher level features such as objects and doors. It established the link between the robots sensors, the objects and the places. This work attempts to build on that idea by asking the question - given a set of objects, how can a robot be made to gain a deeper understanding of its surroundings ? It attempts to form groups in accordance with the hierarchy shown. The objective is a greater incorporation and usage of spatial semantics, thereby producing a concept oriented (thus more semantic) representation of space. In this report specifically, two questions are addressed - (1) How can a robot build a conceptual model of a place ? and (2) How can a robot understand that it is in a particular type of place. The former refers to the problem of conceptualization and the latter, the problem of place classification.

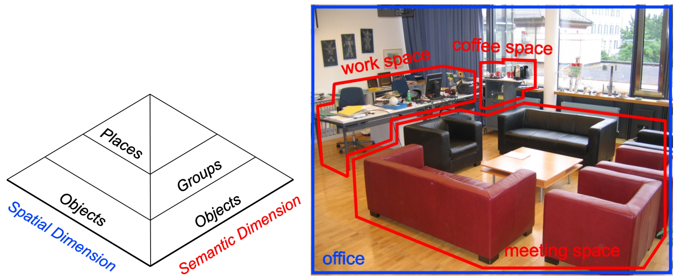

(a) (b)

Fig. 1. **(a)** General approach - A robot uses the sensory information it perceives to identify high level features such as objects, doors etc. These objects are grouped into abstractions along two dimensions - spatial and semantic. Along the semantic dimension, objects are clustered into groups so as to capture the spatial semantics. Along the spatial dimension, places are formed as a collection of groups of objects. Spatial abstractions are primarily perceptual formations (occurrence of walls, doors etc.) whereas semantic or functional abstractions are primarily conceptual formations (similarity of purpose / functionality ; spatial arrangement). The representation is a single hierarchy composed of sensory information being mapped to increasingly abstract concepts. **(b)** An example scenario - The figure depicts a typical office setting. The approach proposed in this work would would enable a robot to recognize various objects, cluster the respective objects into meaningful semantic entities such as a meeting-space and a work-space and even understand that the place is an office because of the presence of a place to work and one to conduct meetings.

In accordance with figure 1(a), objects are incrementally grouped into clusters which are conceptualized as functional groupings (concepts or groups in this report). These groups provide for meaningful semantics that the robot can glean as it explores a place. The robot can then use the groups to infer about or classify the place. Inference is based on the Naive Bayes Classifier (NBC). The key improvement lies in the creation of an intermediate level of semantic understanding, which certainly increases semantic content in the representation but may also improve understanding at higher levels of abstraction. The following sub-sections detail the clustering methodology used, the conceptual model, and the conceptualization process.

B. On the clustering methodology

The conceptualization process to actually infer the concepts works on clusters of objects. Different clustering approaches inspired by [8] were attempted. Most were based on nearest neighbor approach as distance between objects was a reasonable metric to cluster them. The objective, however, was to also make use of the semantic information captured in the concept models learnt by the robot. Thus, a nearest neighbor approach in conjunction with a Maximum-a-posteriori (MAP) estimate of the best case concept (for the incoming object) was the basis of the clustering method that has finally been used in this work. The former used the distance to the center of the cluster as the metric whereas the latter was the concept that had the maximum posterior belief given the occurrence of the single object (it is computed by learning, from the training data, the likelihood of observing the object, given the occurrence of the concept). The behavior of the algorithm can be briefly summarized in three steps in the same order of precedence - (1) choose the nearest cluster that has the same concept as the best case concept suggested for the incoming object, (2) choose the nearest cluster that is conceptually dissimilar but "acceptably likely" with respect to the best case concept and (3) create a new cluster with the incoming object of type suggested by the best case concept. More details on this can be obtained from [23].

C. The Concept Model and Conceptualization

$$P(\,c\,,\,X_0\,,\,X_1\,,\,\dots\,,\,X_n\,) = P(\,c\,)\,*\,\prod_{i=0}^{n} P(X_i\,|\,c) \tag{1}$$

Equation 1 shows the joint probability distribution (JPD) of the model used in this work. Given a set of objects (o_i), the equation computes the belief in a concept given number of occurrences (m_i) respectively, of each of the objects. Every X_i in the equation denotes an $o_i = m_i$ for the corresponding object; c denotes the concept that is to be inferred. The inference is principled on the Bayes rule that interprets this in terms of the prior belief in the concept and the likelihoods of observing the specific number of occurrences of the respective objects, given the concept. Given that a NBC is the underlying model, the object occurrences are assumed to be independent of each other, given the concept. The same method is also used to infer about the place given the occurrence of one or more concepts.

The conceptual model that is used for the inference, encodes the likelihood of the occurrence of a specific number (over a range, in the cluster under consideration) of a certain object towards the formation of a particular concept. It is worth noting that encoding and using the number of occurrences of various objects rather than just their individual occurrences is a more informative method of distinguishing between concepts. For instance, chairs and tables are common to both a work-space and a dining-space, however, the number of occurrences of each of them is one distinguishing factor. For a set of concepts c_i, a set of objects o_i and a range of possible number of occurrences m_i, the training process uses a collection of positive and negative concept exemplars to compute the likelihoods as shown in equation 2.

$$P(o_i\,=\,mi\,|\,c_i) = \frac{N_{o_i=m_i}\,+\,\delta}{N_{(+/-)\,exemplars}\,+\,(2*\delta)} \tag{2}$$

Variables

$c\,,\,X_0\,,\,X_1\,,\dots X_n\;where\;X_i \equiv o_i = m_i$
i.e. $X_i \equiv m_i$ occurrences of object o_i and c $\equiv$ the concept.

Decomposition

$$\left\{ P(\,c\,,\,X_0\,,\,\dots\,,\,X_n) = P(\,c\,)\,*\,\prod_{i=0}^{n} P(X_i\,|\,c) \right.$$

Parametric Forms

$$P(c) \rightarrow \begin{cases} P([c\,=\,0]) = (\frac{n_f+\delta}{n_f+n_t+2\delta}) \\ P([c\,=\,1]) = (\frac{n_t+\delta}{n_f+n_t+2\delta}) \end{cases}$$

$$P(X_i\,|\,c) \rightarrow \begin{cases} P([X_i\,=\,0]|[c\,=\,0]) = 1 - \frac{n_{fi}+\delta}{n_f+2\delta} \\ P([X_i\,=\,0]|[c\,=\,1]) = 1 - \frac{n_{ti}+\delta}{n_t+2\delta} \\ P([X_i\,=\,1]|[c\,=\,0]) = \frac{n_{fi}+\delta}{n_f+2\delta} \\ P([X_i\,=\,1]|[c\,=\,1]) = \frac{n_{ti}+\delta}{n_t+2\delta} \end{cases}$$

Identification $\rightarrow$ Parameters learned during the training process

Question $\rightarrow P(c|X_0, X_1, ..., X_n)$

Fig. 2. The Bayesian program that summarizes the conceptualization and classification processes. It is characterized by the specification of the variables of the system, the decomposition of the joint probability distribution (JPD), the parametric forms of each of the components of the JPD, a specification of how the parameters of the distributions are learnt and finally, the question that is to be answered by the system. '0' denotes a 'false' and '1' denotes a 'true'. δ = 0.001 to accurately reflect on the training data, it ensures that a previously unseen event is an unknown event (belief = 0.5) and not one that never occurs. n_{fi} and n_{ti} are the number of occurrences of the case $o_i = m_i$ in negative and positive exemplars respectively. n_f and n_t are the number of negative and positive exemplars respectively. The same process can be applied to infer about places given the concepts observed.

where the numerator encodes the number of occurrences of the particular case $o_i = m_i$, for every object over a range of occurrences, and the denominator encodes the number of positive or negative exemplars of the concept. The terms δ and $2 * \delta$ ensure that an event that has not been encountered during prior training, is only something that the robot has no prior information about (belief = 0.50) and not something that may never occur. The value of δ decides the reliance on the training data. In the experiments reported in this work, δ takes a very low value of 0.001 so as to reflect the training data accurately. The likelihoods were also limited to taking values between upper and lower bounds in order to avoid 'uninteresting' inferences that could be produced in the limiting cases.

A Bayesian program is a systematic formulation for the creation and usage of Bayesian models such as the one used in this work. Elaborate details on the concept, its structure and its semantics are available in [9]. The Bayesian program used to do the learning and inference process is summarized as shown in figure 2. The complete probabilistic model for the system, including the parameters, likelihoods and the question to be answered, are depicted in it.

IV. EXPERIMENTS

A. Overview

Experiments were conducted on a dataset that included physically measured object and coordinate information from 11 offices and 8 kitchens. The office data was represented in terms of three concepts (apart from some free-standing objects). These were work-space, storage-space and meeting-space. The kitchen data was described in terms of ten concepts, namely cooking-space, garbage-space, dining-space, bottle-group, glass-group, box-group, mug-group, bag-group, poster-group and book-group. Concepts used in this work represent the manner in which the places were understood by the authors; they are similar to those observed in [22]. The approach however is not ontology-specific. Developing a standardized ontology that could perhaps enable high-level communication between robots is beyond the scope of this work.

Two instances each, of office and kitchen data were used only for testing and the others for both training and testing. Training was performed to learn the unknown parameters shown in figure 2, for each concept. Each concept was trained with its set of positive exemplars and against all other exemplars as negative ones. Testing and evaluation involved the comparison of each of the 991 objects (total number in 19 places) with the corresponding objects in the training input. Conceptualization resulted in four outcomes. An object may have been conceptualized correctly, i.e. it belongs to the correct conceptual group with respect to the training data; it may belong to a group that has not been classified (due to insufficient evidence or multiple competing hypotheses inhibiting a clear inference); it could be a free-standing object in training, that has been assigned a label; finally, the object may belong to a group that has been incorrectly classified. Figures 3 and 4 respectively depict the outcome of clustering, conceptualization and place classification of an office and a kitchen.

B. Evaluating the clustering algorithm

TABLE I
EVALUATION OF THE CLUSTERING ALGORITHM

Outcome	Cases	Percentage (%)
Singleton	10	1.0091
Fused or Broken	304	30.6761
Correct	677	68.3148

Table I summarizes the evaluation of the clustering process. *Correct* cases correspond to objects which belonged to the respective clusters, in comparison with the training data. A significant number of clusters were either fused or broken with others. In most cases, this resulted in for instance, the fusion of two adjacent work-spaces or the inclusion of one or more objects of one cluster in another one. Few cases did occur, where objects characteristic of one concept were clustered with those of another. A clear conceptualization would be unlikely in these cases. A few objects were separated from the rest and formed clusters by themselves - these were regarded as being inaccurate with respect to the training input (where only large objects such as cupboards were treated as singleton clusters). The number of such cases however, was quite low.

C. Evaluating the conceptualization algorithm

TABLE II
EVALUATION OF THE CONCEPTUALIZATION ALGORITHM

Outcome	Cases	% (of classified)	% (of total)
Incorrect	168	19.3772	16.9526
Not classified	124	-.	12.5126
Free Object	7	0.8074	0.7064
Correct	692	79.8155	69.8285

The outcome of the conceptualization experiment is described in table II. This experiment used the algorithm exactly as described in the approach. The outcomes correspond respectively to those detailed in IV-A. Clearly, the number of incorrect and unclassified cases, although not insignificant, is quite small in comparison to the number of correct cases.

The algorithm was also evaluated in the context of place classification. All of the eleven offices and eight kitchens were tested against the models learnt. The concepts that were inferred in the previous step were used to infer about the place. The algorithm performed perfectly, in that the place classification accuracy was 100%.

D. Improving the generalization capability

As observed till now, the training phase resulted in the successful learning of conceptual models for both abstract concepts (groups) and places. However, generalization was limited to the exemplars observed and assumed that the exemplars were void of any uncertainty. The former aspect is quite significant as, for instance, if in training, four chairs were always observed in a dining-space, the occurrence of three chairs (a very plausible scenario) would probably render the algorithm being unable to comprehend the group. In such a scenario, the algorithm should infer the possible existence of a dining-space, albeit with a greater uncertainty. Thus, the algorithm should be able to generalize in a manner such that it is able to handle at least conceptually "adjacent" cases to what it has observed before. Also, given that a user is expected to teach the robot in an on-line learning scenario (and not use some predefined data-base of models), it would only be appropriate to consider the training input as being uncertain.

Two techniques were attempted towards these aims - (1) a smoothing of the concept models, based on a moving-average algorithm and (2) the use of a Gaussian uncertainty in the training process - so that every training input affects not only $P(o_i = m_i|c)$ but also its neighbors. Only the results of the latter approach are reported here. The Gaussian uncertainty incorporates a fixed uncertainty in training input; the choice of the Gaussian noise to be used would depend on the local circumstances and the aspects that need to be modeled. In the experiments presented here, $N(0.0, 0.4472)$ was used in order to consider only the number of occurrences $o_i = m_i$, $o_i = m_i - 1$ and $o_i = m_i + 1$ respectively (i.e. only the immediate neighbors). Results of the conceptualization process with the new models are given in table III - they were very encouraging. Even in the context of place classification - they produced perfect results, identifying all 11 offices and 8 kitchens correctly.

TABLE III
EVALUATION OF THE MODIFIED CONCEPTUALIZATION ALGORITHM

Outcome	Cases	% (of classified)	% (of total)
Incorrect	175	18.5381	17.6589
Not classified	47	-.	4.7427
Free Object	9	0.9534	0.9082
Correct	760	80.5085	76.6902

E. Discussion

While the results from both models are quite promising, it is worth noting the effect of adding the Gaussian uncertainty in training process. From the earlier results, it could be inferred that the modified (Gaussian) conceptual models have managed to successfully classify several of the previously unclassified cases while maintaining or marginally improving the previously obtained classification accuracies - both for concepts (groups) and for places. This indicated a clear improvement in generalization capability of the algorithm. Further, the incorporation of the Gaussian uncertainty enables the algorithm to be more realistic, by accounting for possible sources of uncertainty in the training phase.

V. CONCLUSION

A Bayesian approach to conceptualization and classification of space for mobile robots was presented. The suggested algorithm was based on the Naive Bayes Classifier (NBC) and was implemented using a clustering mechanism and a sound Bayesian Programming methodology. The results vindicated the use of the number of occurrences of various objects towards concept formation; the addition of the Gaussian uncertainty clearly improved the generalization capability of the algorithm and made it more realistic in being able to account for uncertain training data. The algorithm incrementally formed conceptual groups of objects - these represented semantic (functional) groupings that were aimed at capturing spatial semantics; further, they were used for classifying places. The generated concepts increase the amount of semantic information contained in a robot's spatial representation. They also endow the robot with the capability of being more spatially aware machines, capable of reasoning about spatial semantics. Future work will focus on adding spatial semantics and understanding the system's response to dynamic situations.

ACKNOWLEDGMENTS

The authors thank all people who allowed the use of their offices and kitchens for data collection. This work has been supported by the EC under FP6-IST-002020-COGNIRON and FP6-IST-027140-BACS.

REFERENCES

[1] Kai Oliver Arras. *Feature-Based Robot Navigation in Known and Unknown Environments*. PhD thesis, Swiss Federal Institute of Technology Lausanne (EPFL), 2003. Thesis number 2765.

[2] Raja Chatila and Jean Paul Laumond. Position referencing and consistent world modeling for mobile robots. In *IEEE International Conference on Robotics and Automation*, 1985.

[3] Jie Cheng and Russell Greiner. Learning Bayesian Belief Network Classifiers: Algorithms and System. In *Canadian Conference on Artificial Intelligence (CSCSI)*, Ottawa, Canada, May 2001.

[4] Howie Choset and Keiji Nagatani. Topological Simultaneous Localization And Mapping (SLAM): toward exact localization without explicit localization. *IEEE Transactions on Robotics and Automation*, 17:125–137, April 2001.

[5] Richard O. Duda, Peter E. Hart, and David G. Stork. *Pattern Classification*. Wiley Interscience, 2000.

[6] Douglas H. Fisher. Knowledge Acquisition Via Incremental Conceptual Clustering. *Machine Learning*, 2:139–172, 1987.

[7] Cipriano Galindo, Alessandro Saffiotti, Silvia Coradeschi, Pär Buschka, Juan-Antonio Fernandez-Madrigal, and Javier Gonzalez. Multi-Hierarchical Semantic Maps for Mobile Robotics. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3492–3497, Edmonton, Canada, 2005.

[8] A. K. Jain, M. N. Murty, and P. J. Flynn. Data Clustering: A Review. *ACM Computing Surveys*, 31(3):264–323, 1999.

[9] O. Lebeltel, P. Bessière, J. Diard, and E. Mazer. Bayesian Robots Programming. *Autonomous Robots*, 16:49–79, 2004.

[10] David G. Lowe. Distinctive Image Features from Scale-Invariant Keypoints. *International Journal of Computer Vision*, 60(2):91–110, 2004.

[11] Óscar Martínez Mozos, Patric Jensfelt, Hendrik Zender, Geert-Jan M. Kruijff, and Wolfram Burgard. From Labels to Semantics: An Integrated System for Conceptual Spatial Representations of Indoor Environments for Mobile Robots. In *IEEE Int. Conf. on Robotics and Automation (ICRA) Workshop on Semantic Information in Robotics*, 2007.

[12] Óscar Martínez Mozos, Rudolph Triebel, Patric Jensfelt, Axel Rottmann, and Wolfram Burgard. Supervised semantic labeling of places using information extracted from sensor data. *Robotics and Autonomous Systems*, 55(5):391–402, 2007.

[13] Andrew Y. Ng and Michael I. Jordan. On Discriminative vs. Generative classifiers: A comparison of logistic regression and naive Bayes. In *Advances in Neural Information Processing Systems (NIPS) 14*. MIT Press, 2002.

[14] Adriana Tapus. *Topological SLAM - Simultaneous Localization And Mapping with fingerprints of places*. PhD thesis, Swiss Federal Institute of Technology Lausanne (EPFL), 2005. Thesis Number 3357.

[15] Ben Taskar, Eran Segal, and Daphne Koller. Probabilistic Classification and Clustering in relational data. In *Seventeenth International Joint Conference on Artificial Intelligence (IJCAI)*, Seattle, Washington, USA, August 2001.

[16] Joshua B. Tenenbaum. Bayesian Modeling of human concept learning. In S.A. Solla M.S. Kearns and D.A. Cohn, editors, *Advances in Neural Information Processing Systems (NIPS) 11*, Canbridge, MA, USA, 1999. MIT Press.

[17] Sebastian Thrun. Learning Metric-Topological Maps for Indoor Mobile Robot Navigation. *Artificial Intelligence*, 99 (Issue-1):21–71, 1998.

[18] Sebastian Thrun. *Exploring Artificial Intelligence in the New Millenium*, chapter Robotic mapping: A survey. Morgan Kaufmann, 2002.

[19] Nicola Tomatis, Illah Nourbakhsh, and Roland Siegwart. Hybrid Simultaneous Localization And Map building: A natural integration of topological and metric. *Robotics and Autonomous Systems*, 44:3–14, July 2003.

[20] Iwan Ulrich and Illah Nourbakhsh. Appearance-Based Place Recognition for Topological Localization. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 1023–1029, San Francisco, CA, USA, April 2000.

[21] Shrihari Vasudevan, Stefan Gächter, Viet Tuan Nguyen, and Roland Siegwart. Cognitive Maps for Mobile Robots - An object based approach. *Robotics and Autonomous Systems*, 55(5):359–371, 2007.

[22] Shrihari Vasudevan, Stefan Gächter, and Roland Siegwart. Cognitive Spatial Representations for Mobile Robots - Perspectives from a user study. In *IEEE Int. Conf. on Robotics and Automation (ICRA) Workshop on Semantic Information in Robotics*, Rome, Italy, April 2007.

[23] Shrihari Vasudevan and Roland Siegwart. A Bayesian Conceptualization of Space for Mobile Robots. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, San Diego, USA, October 2007.

[24] Benjamin W. Wah. Generalization and Generalizability Measures. *IEEE Transactions on Knowledge and Data Engineering*, 11(1):175 – 186, January - February 1999.

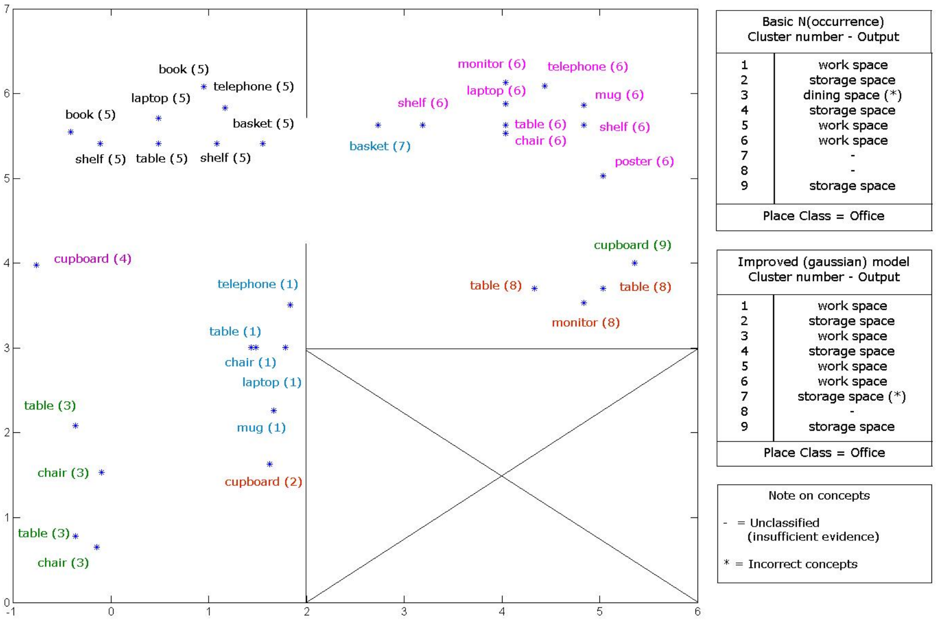

Fig. 3. Outcome of the clustering, conceptualization and place classification processes for an office. The depiction is a top-down view. Each cluster of objects is identified by a color and a number in parenthesis. On the right are the outcomes as obtained using the two cases - with only the basic model that uses the number of occurrences and with the improved (Gaussian) model. Each cluster is classified as being one of 13 concepts used in this work. Note that the basket in cluster 7 is classified as a storage space. This is primarily due to the non-occurrence of all other known objects in that cluster and the prior probability of the occurrence of the storage-space concept in relation to that of other concepts.

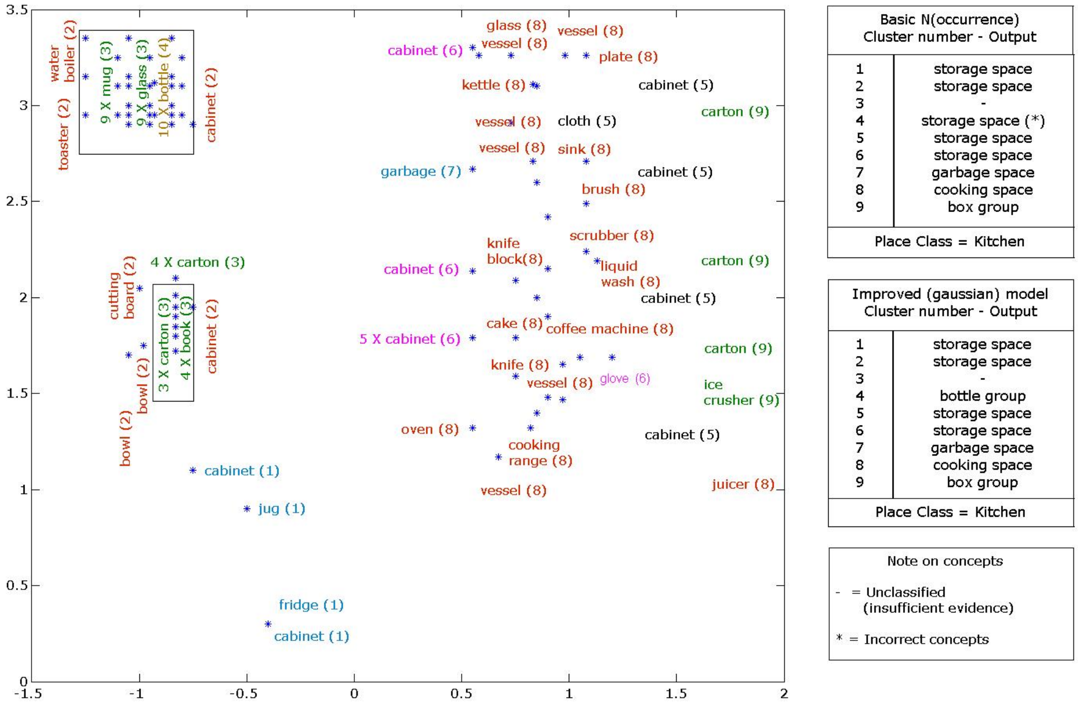

Fig. 4. Outcome of the conceptualization and classification processes for a kitchen. The depiction is a top-down view. In general, the 3D map would have the objects more to the left / right parts of the image, at higher-levels / more-on-the-inside as compared to those that are located approximately at the center of the figure (which was the walking space in the kitchen). The rectangles depict the cabinets containing various objects within it in different rows. Each cluster of objects is identified by a color and a number in parenthesis. Note that *'n X object'* is used to denote n occurrences of an object. On the right are the outcomes as obtained using the two cases - with only the basic model that uses the number of occurrences and with the improved (Gaussian) model. Each cluster is classified as being one of 13 concepts used in this work. Note that the glasses and mugs in cluster 3 are fused with the storage space and are not individual clusters or not clustered with say, the bottle-group. This is because the clustering depends on the distance threshold, the prior probability of the occurrence of the concept and the likelihood of observing the object in an instance of the concept. In the trained model, storage-spaces occurred more frequently than bottle-groups. More details on the clustering mechanism can be obtained from [23].