

A Bayesian Conceptualization of Space for Mobile Robots

Shrihari Vasudevan and Roland Siegwart
Autonomous Systems Laboratory (ASL)
Swiss Federal Institute of Technology Zürich (ETHZ)
8092 Zürich, Switzerland.
{shrihari.vasudevan, r.siegwart}@ieee.org

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 of making robots more spatially cognizant, 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 from the perceived objects and attempts to classify places using the concepts gleaned. The approach is based on learning from exemplars, clustering and the use of Bayesian network classifiers. 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 [1]. Robot maps can be generally classified into three categories - metric ([2], [3]), topological ([4], [5]) and hybrid ([6], [7]). 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 ([3], [5] & [8]) 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.

Typically, humans 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 [9]. Thus, a *cognitive* spatial representation, for a mobile robot, could be expected to encode similar information. The work reported in [10] attempted to create such a representation by encoding typical household objects and doors within a hierarchical probabilistic framework. It used a SIFT [11] 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) [12] that did not learn from negative exemplars. The likelihood formulation in the conceptualization was not useful for encoding the *importance* or *significance* of objects (handling multiple occurrences) - e.g. the occurrence of two tables is more significant for a concept than a single object occurrence. 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 addresses these issues in the larger context of proposing a consistent Bayesian framework for the incorporation of spatial semantics in representations for mobile robots. Further, [10] 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 [13] 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 [14]. 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 [17] and [18]. 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 relies on object occurrence cues, used in a logic or rule based framework, possibly with a predefined ontology [19]. 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 [20]. The work [21] 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 [22]. Additionally, this work attempts to be fully probabilistic and is grounded on a Bayesian Programming methodology as described in [23].

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 concept-oriented representation that results from this process, enables robots to be much more spatially cognizant of their surroundings and yet totally compatible with humans (demonstrated in [9]).

III. APPROACH

A. Overview

Figure 1 illustrates the overall approach that is being pursued. In [10], 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

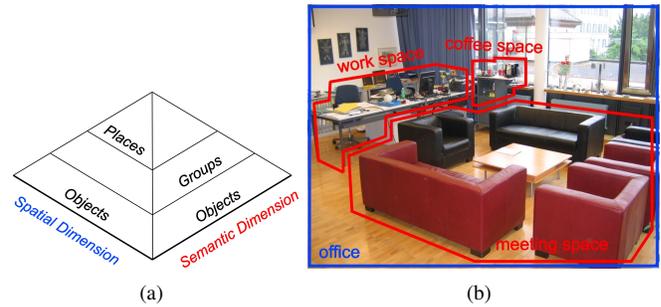


Fig. 1. (a) General approach - A robot uses the sensory information it perceives, to identify high level features such as objects, doors etc. These objects are grouped into abstractions along two dimensions - spatial and semantic. Along the semantic dimension, objects are clustered into groups so as 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 proposed approach 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.

objects and doors. 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 ? The principle idea is that adding concepts (created for instance using the functional similarity of the underlying objects) to a purely navigation oriented map would result in the incorporation / usage of spatial semantics and the formation of a concept-oriented representation of space. In this report, 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.

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.

B. On the clustering methodology

The conceptualization process to actually infer the concepts works on clusters of objects. Different clustering approaches inspired by [22] 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 was computed from the concept models learnt, these encoded the prior belief in a concept and the likelihood of observing an 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” (threshold set empirically) 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.

C. The Concept Model and Conceptualization

$$P(c, o_1, \dots, o_n) = P(c) * \prod_{i=1}^n P(o_i | c) \quad (1)$$

Equation 1 shows the joint probability distribution (JPD) of the proposed model. It computes the belief in a concept given the objects perceived. This is done using Bayes rule, which interprets this in terms of the prior belief in the concept and the likelihoods of the occurrences of the objects given the concept in consideration. These priors and likelihoods are encoded in the learnt (training) concept models. Given that a NBC is the underlying model, the objects 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 rule-of-succession, as used in previous work [10] to estimate likelihoods, was un-conducive towards modeling multiple object occurrences. This is necessary in order to model the *importance* or *significance* of various objects towards the formation of more abstract concepts. Thus the likelihoods were generated by choosing an appropriate function so as to be able to effectively handle multiple object occurrences. The aim was to use a likelihood function that would generate probabilities, that would encode the fact that multiple occurrences of an object may imply a greater significance of the object for the occurrence of the concept and that would exhibit a monotonically increasing behavior. Given these constraints, the exponential function

$$f(x) = 1 - \exp(-\lambda * x) \text{ where } x = \frac{n + \delta}{N + (2 * \delta)}$$

was chosen as the likelihood function. Here, n and N are respectively, for instance, the number of occurrences of a particular object in positive exemplars of a concept and the total number of positive 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 and not something that may never

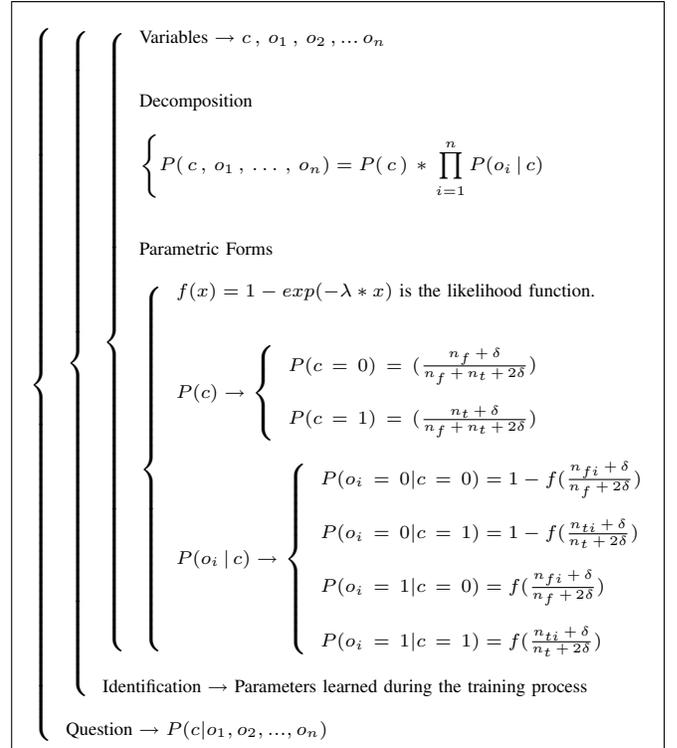


Fig. 2. The Bayesian program that summarizes the conceptualization and classification processes. It is characterized by the specification of the variables of the system, the decomposition of the joint probability distribution (JPD), the parametric forms of each of the components of the JPD, a specification of how the parameters of the distributions are learnt and finally, the question that is to be answered by the system. ‘c’ denotes the concept and the various o_i are the objects observed. ‘0’ denotes a ‘false’ and ‘1’ denotes a ‘true’. $\lambda = \ln(4)$ and $\delta = 0.001$. n_{f_i} and n_{t_i} are the number of occurrences of the i^{th} object in negative and positive exemplars (concept) respectively. n_f and n_t are the number of negative and positive exemplars (concept) respectively. The same process can be applied to infer about places given the concepts observed.

occur. The parameter λ is chosen to be $\ln(4)$ to achieve this effect - $f(0.5) = 0.5$. Also, the behavior of the likelihood function is data-driven in that it closely follows the relative frequency and tapers off beyond a certain point in order to deal with multiple object occurrences. The value of δ decides (inversely) 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.

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 [23]. The complete Bayesian program used to do the learning and inference processes in this work, is summarized as shown in figure 2.

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 [9]. The approach however is not ontology-specific.

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 fig. 2, for each concept. Each concept was trained with its set of positive exemplars and against all other exemplars as negative exemplars. Testing and evaluation involved the comparison of each of the 991 objects with the corresponding training input in the context of both clustering and conceptualization. Conceptualization resulted in a total of 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 have been conceptualized incorrectly or it may belong to a group that has not been classified at all. The latter-most outcome could itself happen in two situations - an object may be a part of a conceptual group that has been identified as two or more correct concepts (e.g. two mugs in a work-space also being identified as a mug-group), or alternatively, it may belong to a conceptual group that requires more evidence (objects) to be classified. If multiple correct concepts are identified, the cluster is not classified but the outcome is not considered incorrect as it requires a representation of relationships between concepts - something which will be addressed in future. Figures 3 and 4 respectively depict the outcome of conceptualization process applied to an office and a kitchen.

B. Evaluation of the clustering algorithm

TABLE I
EVALUATION OF THE CLUSTERING ALGORITHM

Case	Number (of 991)	Percentage (%)
Singleton	18	1.8163
Fused or Broken	274	27.6488
Correct	699	70.5348

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.

TABLE II
EVALUATION OF THE CONCEPTUALIZATION ALGORITHM

Case	Number	% of N(classified)	% of 991
Incorrect	3	0.6085	0.3027
Not classified	480	-	48.4359
Multiple acceptable	18	-	1.8163
Correct	490	99.3915	49.4450

C. Evaluating the Conceptualization algorithm

The concept models indirectly encoded the extent to which each object contributed/did-not-contribute to the concept. While every concept model had many significant entries, there were also numerous entries that did not contribute one way or the other. Such entries from the model were pruned off in order to minimize the size of the model to its most distinctive form. To this effect, an empirically determined distinctiveness threshold was applied to each concept model, so as to avoid the use of entries that were not distinctive enough. Distinctiveness was computed as the absolute difference between $P(o_i = 1|c = 1)$ and $P(o_i = 1|c = 0)$. Entries with very similar values in the concept model were thus pruned off. The outcome of the conceptualization experiment is described in table II. Its worth noting that although the number of unclassified outcomes is very high, the the actual number of erroneous outcomes is extremely low.

D. Improving the classification rate - constraining the model

TABLE III
EVALUATION OF THE CONCEPTUALIZATION ALGORITHM

Outcome	Number	% of N(classified)	% of 991
Incorrect	51	8.5427	5.1463
Not classified	329	-	33.1988
Multiple acceptable	65	-	6.5590
Correct	546	91.4573	55.0959

The dataset used had 71 objects. It was observed that some concepts were quite sparse in that they were typically defined by the presence of only a small number of objects (and mostly the absence of all the other objects). For instance, a garbage-space is typically defined by the occurrence of a few objects and the non-occurrence of every other object. In a limiting case, such conceptualization would seem counter-intuitive. Further, during inference, the non-occurrence of many objects could signal the occurrence of one or more “wrong” concepts. Even if this happened alongside the signaling of the “right” concept, the conceptualization would fail to occur as no clear outcome would be possible. Since at any time, the robot observed only a small fraction of the set of all known objects, this issue was a major problem in the calling and most likely explained the numerous unclassified outcomes in table II. In order to deal with this issue, a constraint was imposed on the model. The constraint enforced that the concept model always had at least as much information about objects contributing positively to it (i.e. for it) as objects contributing otherwise. This was

implemented as $\Sigma P(o_i = 1|c = 1) \geq \Sigma P(o_i = 1|c = 0)$. The model now only comprised of those elements that contribute positively to the occurrence of the concept and the most distinctive of those which contribute negatively. The outcome of the conceptualization process is shown in table III. In relation to the previous case, a clear decrease in the number of unclassified clusters was observed. Although, the number of correct outcomes (of those classified) decreases, the total number of correctly classified outcomes increases very significantly. The incorrect outcomes, although not as insignificant as in II, is not high. The constrained model thus tends to infer more about the same data as compared to the unconstrained one.

E. Place Classification

TABLE IV
EVALUATION OF PLACE CLASSIFICATION

Model	Office (11)	Kitchen (8)	total (19)
Distinctiveness only	11	3	14
Distinctiveness + constraint	11	8	19

Table IV compares both models on their ability to classify 19 places (11 offices and 8 kitchens). The distinctiveness-only model performed poorly in classifying kitchens, in relation to the one that included the constraint. The reason for this being the non-occurrence of many kitchen concepts made the system infer the presence of an office (absence of a kitchen related concept contributes negatively to a kitchen outcome and positively to an office outcome). This problem was eliminated in the concept-model that uses the constraint.

V. CONCLUSION

A Bayesian approach to conceptualization 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 algorithm is supervised in that it is “taught” different concepts but it is not ontology-specific and would rely on a robot’s inherent object recognition capability. The algorithm incrementally formed conceptual groups of objects - these represented semantic (functional) groupings that were aimed at capturing the spatial semantics. The results obtained, provide the foundation for future work. In order to incorporate more semantic information in the conceptualization process and to attempt to improve overall performance, explicit encoding of number of occurrences of various objects in various concepts and that of spatial relationships would be attempted. A comparison with other learning methodologies is also planned.

ACKNOWLEDGMENTS

The authors thank Ahad Harati, Rudolph Triebel and Ralf Kästner for their suggestions, 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] S. Thrun, *Exploring Artificial Intelligence in the New Millennium*. Morgan Kaufmann, 2002, ch. Robotic mapping: A survey.
- [2] R. Chatila and J. P. Laumond, “Position referencing and consistent world modeling for mobile robots,” in *IEEE International Conference on Robotics and Automation*, 1985.
- [3] K. O. Arras, “Feature-based robot navigation in known and unknown environments,” Ph.D. dissertation, Swiss Federal Institute of Technology Lausanne (EPFL), 2003, Thesis number 2765.
- [4] H. Choset and K. Nagatani, “Topological Simultaneous Localization And Mapping (SLAM): toward exact localization without explicit localization,” *IEEE Transactions on Robotics and Automation*, vol. 17, pp. 125–137, April 2001.
- [5] A. Tapus, “Topological SLAM - Simultaneous Localization And Mapping with fingerprints of places,” Ph.D. dissertation, Swiss Federal Institute of Technology Lausanne (EPFL), 2005, Thesis Number 3357.
- [6] S. Thrun, “Learning metric-topological maps for indoor mobile robot navigation,” *Artificial Intelligence*, vol. 99 (Issue-1), pp. 21–71, 1998.
- [7] N. Tomatis, I. Nourbakhsh, and R. Siegwart, “Hybrid Simultaneous Localization And Map building: A natural integration of topological and metric,” *Robotics and Autonomous Systems*, vol. 44, pp. 3–14, July 2003.
- [8] I. Ulrich and I. Nourbakhsh, “Appearance-based place recognition for topological localization,” in *IEEE International Conference on Robotics and Automation (ICRA)*, San Francisco, CA, USA, April 2000, pp. 1023–1029.
- [9] S. Vasudevan, S. Gächter, and R. 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.
- [10] S. Vasudevan, S. Gächter, V. T. Nguyen, and R. Siegwart, “Cognitive Maps for Mobile Robots - An object based approach,” *Robotics and Autonomous Systems*, vol. 55, no. 5, pp. 359–371, 2007.
- [11] D. G. Lowe, “Distinctive image features from scale-invariant keypoints,” *International Journal of Computer Vision*, vol. 60, no. 2, pp. 91–110, 2004.
- [12] R. O. Duda, P. E. Hart, and D. G. Stork, *Pattern Classification*. Wiley Interscience, 2000.
- [13] B. W. Wah, “Generalization and generalizability measures,” *IEEE Transactions on Knowledge and Data Engineering*, vol. 11, no. 1, pp. 175 – 186, January - February 1999.
- [14] D. H. Fisher, “Knowledge acquisition via incremental conceptual clustering,” *Machine Learning*, vol. 2, pp. 139–172, 1987.
- [15] B. Taskar, E. Segal, and D. Koller, “Probabilistic classification and clustering in relational data,” in *Seventeenth International Joint Conference on Artificial Intelligence (IJCAI)*, Seattle, Washington, USA, August 2001.
- [16] J. B. Tenenbaum, “Bayesian modeling of human concept learning,” in *Advances in Neural Information Processing Systems (NIPS) 11*, S. S. M.S. Kearns and D. Cohn, Eds. Cambridge, MA, USA: MIT Press, 1999.
- [17] C. Galindo, A. Saffiotti, S. Coradeschi, P. Buschka, J.-A. Fernandez-Madrigo, and J. Gonzalez, “Multi-hierarchical semantic maps for mobile robotics,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Edmonton, Canada, 2005, pp. 3492–3497.
- [18] Ó. M. Mozos, R. Triebel, P. Jensfelt, A. Rottmann, and W. Burgard, “Supervised semantic labeling of places using information extracted from sensor data,” *Robotics and Autonomous Systems*, vol. 55, no. 5, pp. 391–402, 2007.
- [19] Ó. M. Mozos, P. Jensfelt, H. Zender, G.-J. M. Kruijff, and W. 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.
- [20] A. Y. Ng and M. 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.
- [21] J. Cheng and R. Greiner, “Learning bayesian belief network classifiers: Algorithms and system,” in *Canadian Conference on Artificial Intelligence (CSCSI)*, Ottawa, Canada, May 2001.
- [22] A. K. Jain, M. N. Murty, and P. J. Flynn, “Data clustering: a review,” *ACM Computing Surveys*, vol. 31, no. 3, pp. 264–323, 1999.
- [23] O. Leibel, P. Bessière, J. Diard, and E. Mazer, “Bayesian robots programming,” *Autonomous Robots*, vol. 16, pp. 49–79, 2004.

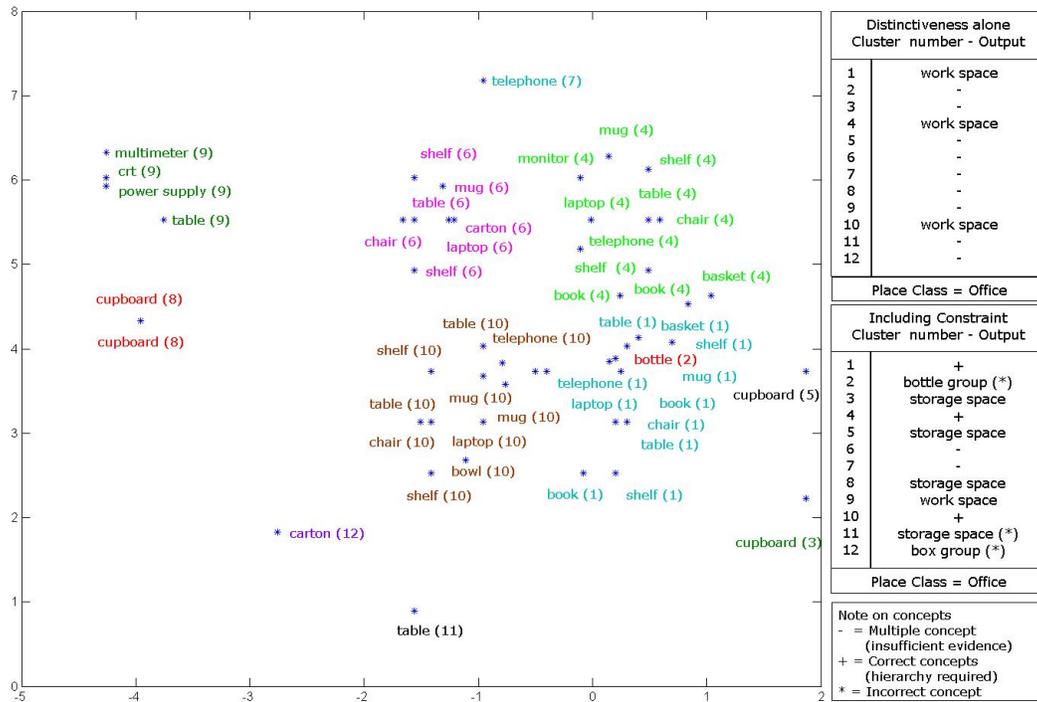


Fig. 3. Outcome of the conceptualization and classification processes for an office. The depiction is a top-down view. Each cluster of objects is identified by a color and a number in parenthesis. On the right are the outcomes as obtained using the two cases - with only a distinctiveness threshold and including a constraint in the model. Each cluster is classified as being one of 13 concepts used in this work. The output from the place classification algorithm in each case is also shown. Clearly the constrained model performs much better than the model with only a distinctiveness measure, but it also generates more erroneous outcomes.

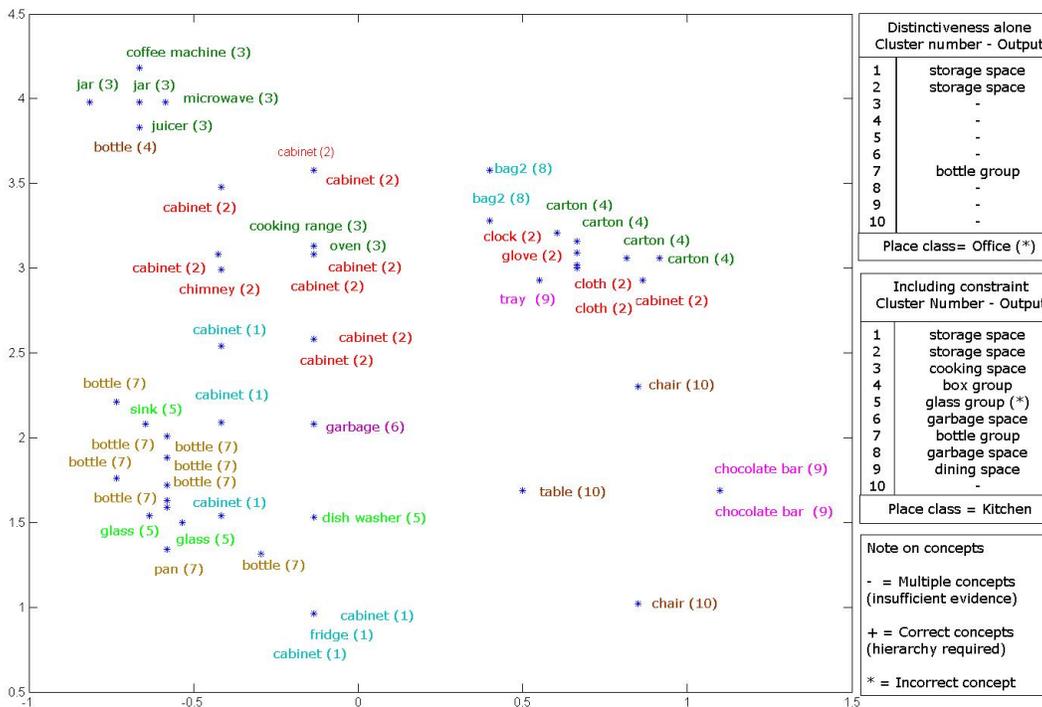


Fig. 4. Outcome of the conceptualization and classification processes for an kitchen. The depiction is a top-down view. In general, the 3D map would have the objects more to the left / right at higher-levels / more-on-the-inside as compared to those that are located approximately at the center of the figure (which was the walking space in the kitchen). 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 a distinctiveness threshold and including a constraint in the model. Each cluster is classified as being one of 13 concepts used in this work. The output from the place classification algorithm in each case is also shown. Clearly the constrained model performs much better than the model with only a distinctiveness measure.