

Learning Concepts with Multi-Robot Systems

Ana Cristina Palacios-García, Angélica Muñoz-Meléndez, and Eduardo F. Morales

National Institute for Astrophysics, Optics and Electronics,
Luis Enrique Erro No. 1, Santa María Tonantzintla, Puebla, México
cpalacios@ccc.inaoep.mx, munoz@inaoep.mx, emorales@inaoep.mx

Abstract. This paper introduces a novel approach to learn representations of objects using a team of robots. Each robot extracts local and global visual features of objects and combines them to represent and recognize objects. Contrary to previous approaches the robots do not know in advance the number or nature of objects to learn. Individual representations of objects are learned on-line while the robots are traversing an environment. Robots share their individual concepts to improve their own concepts, and to acquire a new representation of an object not seen by them. For that, the robots have to detect if they are seeing a new object or an already learned one. We empirically evaluated our approach with a real world robot team with very promising results.

Keywords: Robotics and automation; mobile robots and autonomous systems; vision, recognition and reconstruction; network robotics.

1 Introduction

Collective robotics is a very active research area in the robotics community. Multi-robot systems or robot teams have effectively emerged as an alternative paradigm for the design and control of robotic systems because of the team's capability to exploit redundancy in sensing and actuation.

The research on collective robotics has focused on developing mechanisms that enable autonomous robots to perform collective tasks, such as strategies for coordination and communication [1, 2]; exploration, mapping and deployment [3]; sensing, surveillance and monitoring [4]; and decentralized decision making [5]. In these works, a robot team can reduce time to complete a complex task that is allocated among its members.

Despite constant research on the design of robot teams, very little attention has been paid so far to the development of robot teams capable of learning from their interaction with their environment. In addition to their capability for accelerated learning, learning robot teams can be used to acquire a much richer and varied information compared to the information acquired by single learning robots.

Learning is a key issue to achieve autonomy for both, single robots and robot teams. Learning capabilities can provide robots the flexibility and adaptation

needed to cope with complex situations. In the context of robot teams, the most common machine learning approach has been reinforcement learning, where the idea is to learn optimal policies using a set of robots to improve the coordination of individual actions in order to reach common goals [1, 2, 4, 6].

In this work we use visual information to learn, with a team of robots, descriptions of objects placed in a particular environment. Learning to recognize particular objects in an environment is important for robotics as it can be used for local and global localization tasks as well as for simple service tasks such as searching for objects in unknown places. Contrary to previous approaches, in our learning setting, the robots are not told the number or nature of the objects to be learned.

Vision is a primary source of perception in robotics and provides different features that can be used to classify objects. In general, using a particular set of features can be adequate for particular tasks but inadequate for other tasks. In this work, objects are characterized by two complementary features: (i) SIFT features [7] and (ii) information about the silhouettes of objects. Other features could be used as well, but the main objective in this work is to show the different cases and possible confusions that can arise in the recognition of objects and merging of concepts, and how they can be addressed.

This article deals with the individual and collective representation of objects from visual information using a team of autonomous robots.

The rest of the paper is organized as follows. Section 2 reviews related work. Sections 3 and 4 introduce, respectively, the stages of individual learning and collective learning of concepts. Section 5 describes our experimental results, and Section 6 provides conclusions and future research work.

2 Related Work

Interesting experiments where physical mobile robots learn to recognize objects from visual information have been reported. First we review significant work developed for individual learning, and then we review learning approaches developed for robot teams.

In [8] different learning techniques to acquire automatically semantic and spatial information of the environment in a service robot scenario are applied. In that work, a mobile robot autonomously navigates in a domestic environment, builds a map, localizes its position in the map, recognizes objects and locates them in the map. Background subtraction techniques are applied for foreground object segmentation. Then objects are represented by SIFT points [7] and an appearance-based method is used for detecting objects named Receptive Field Co-occurrence Histograms. The authors developed a method for active object recognition which integrates both local and global information of objects.

In [9] authors applied an instance-based method to train a robot for object recognition purposes. The objects are represented by color histograms. Different representations are learned from different views of the same object using a mediator that can ask questions and provide feedback on the robot's answers.

The recognition is performed by classifying new views of objects using the KNN algorithm [10].

In [11], a scheme for fast color invariant ball detection in the RoboCup context is presented. To ensure the color-invariance of the input images, a preprocessing stage is first applied for detecting edges using the Sobel filter, and specific thresholds for color removal. Then, windows are extracted from images and predefined spatial features, such as edges and lines, are identified in these windows. These features serve as input to an AdaBoost learning procedure that constructs a cascade of regression tree classifiers (CART). The system is capable of detecting different soccer balls in RoboCup and other environments. The resulting approach is reliable and fast enough to classify objects in real time.

In [12] authors present an on-line method for learning objects for human-robot interaction. In this case, the robot's user is included in the learning framework as an instructor. For the learning method, authors used a lifelong incremental learning system that evolves with any new learned object based on one-class learning (OCLL). Objects are represented using normalized shape features. Additionally, the authors proposed an experimental methodology for evaluating a word learning method, and for comparing the word learning capabilities of different agents as well as to assess research progress in similar works.

Concerning the problem of collective learning of objects using robot teams there are, as far as we know, very few works. In [13] the authors address the problem of mobile object recognition based on kinematic information. The basic idea is that if the same object is being tracked by two different robots, the trajectories and therefore the kinematic information observed by each robot must be compatible. Therefore, location and velocities of moving objects are the features used for object recognition instead of features such as color, texture, shape and size, more appropriate for static object recognition. Robots build maps containing the relative position of moving objects and their velocity at a given time. A Bayesian approach is then applied to relate the multiple views of an object acquired by the robots.

In [14], objects are represented with Principal Components (PC) learned from a set of global features extracted from images of objects. An object is first segmented and its global features such as color, texture, and shape are then extracted. Successive images in a sequence are related to the same object by applying a Kalman filter. Finally, a 3D reconstructed model of an object is obtained from the multiple views acquired by robots. For that purpose, a Shape From Silhouette based technique [15] is applied.

The main drawbacks of previous works can be summarized as follows: (i) Most approaches are able to cope with a limited number of learning objects, usually 3 to 12 objects, and (ii) the number of learning objects is known *a priori*. In contrast to previous works, in our method each robot learns on-line individual representations of objects without prior knowledge on the number or nature of the objects to learn. Individual concepts are represented as a combination of global and local features extracted autonomously by the robots from the training objects. A Bayesian approach is used to combine these features and used for

classification. Individual concepts are shared among robots to improve their own concepts, combining information from other robots that saw the same object, and to acquire a new representation of an unnoticed object. We also analyzed and provide solutions to the different cases that can occur from information of two individual concepts.

3 Individual Learning of Concepts

The individual concepts are learned on-line by a robot team while traversing an environment without prior knowledge on the number or nature of the objects to learn. The individual learning of concepts consists of three parts: object detection, feature extraction, and individual training.

Individual concepts of objects are represented by Principal Components (PC) over the information about the silhouettes of objects and by Scale Invariant Features (SIFT). Learned concepts are shared among robots.

3.1 Object Detection

Robots move through an environment and learn descriptions of objects that they encountered during navigation. Objects are detected using background subtraction. In this paper we assume a uniform and static background. We performed morphological operations (closing and erode) to achieve better segmentation. Once an object is detected, it is segmented and scaled to a fixed size, to make the global PC features robust to changes in scale and position.

3.2 Feature Extraction and Individual Training

The segmented objects are grouped autonomously by the robots in sets of images containing the same object. Robots assume that they are observing the same object while it can be detected, and they finish to see it when they can not detect objects in the captured images. Only one object can be detected in an image at the same time. For each set of images, the robot obtains an individual concept that represents the object.

Training using global features: We applied Principal Component Analysis (PCA) over the average silhouettes that are automatically extracted from the set of images of a particular object. The average provides a more compact representation of objects and reduces segmentation errors. Figure 1 (a) shows an image of the object *vase* used in the training phase, Figure 1 (b) shows its silhouette, and Figure 1 (c) illustrates the average silhouette obtained from a set of images that represent the object of Figure 1 (a). Once the robot has obtained an average silhouette, this is added by the robot to a set of known average silhouettes. After that, the robot uses PCA to reduce the dimensionality of all average silhouettes learned to get the PC features that represent them.

Training using local features: Each robot extracts local SIFT features of each image of the set of images, and groups them in a final set which contains all

the different SIFT features that represent an object. In Figure 1 (d) we show an example of the SIFT points obtained from a set of images of a *vase* and the final set of SIFT points obtained. The PC features and the SIFT features represent the individual concept of the observed object.




Fig. 1. Examples of the silhouette (b) and average silhouette (c) of an object (a). Examples of the SIFT features extracted from a set of images and the final set of SIFT features (d).

3.3 Sharing Concepts

The concepts learned by robots are shared among them to achieve collective learning. This can be done off-line or on-line. In the case of collective off-line learning the robots share their individual concepts once they have learned all the training objects. On the other hand, in the collective on-line learning the robots share their individual concept as soon as a new object is learned.

4 Collective Learning of Concepts

Collective learning of concepts enables robots to improve individual concepts combining information from other robots that saw the same object, and to acquire a new representation of an object not seen by them. Therefore, a robot can learn to recognize more objects of what it saw and can improve their own concepts with additional evidence from other robots.

A robot has to decide whether the concept shared by another robot is of a new object or of a previously learned concept. A robot can face three possibilities: coincident, complementary or confused information. The shared concepts are fused depending on the kind of information detected, as described below.

4.1 Pre-analysis of Individual Concepts

The concept learned by a robot is defined as follows:

$$C_k^i = \{Sil_k^i, SIFT_k^i\} \quad (1)$$

where C_k^i is the concept k learned by robot i , Sil_k^i is the average silhouette, and $SIFT_k^i$ is the set of SIFT features that form the concept k .

In order to determine if a shared concept is previously known or not to a robot, it evaluates the probabilities that the PC features and SIFT features are previously known by the robot. The probability vectors of PC features calculated by robot i , v_P^i , indicate the probability that a concept shared by robot j , C_k^j , is similar to the concepts known by robot i , $C_1^i, \dots, C_{numObjs}^i$, given the global features. $numObjs$ is the number of concepts of objects known by robot i . The process to obtain the probability vector PCA is described as follows:

- A temporal training set of silhouettes is formed by adding the average silhouettes of concepts known by robot i or actual robot, $Sil_1^i, \dots, Sil_{numObjs}^i$, and the average silhouette of the new or shared concept Sil_k^j .
- The PCA is trained using the temporal set of average silhouettes. The projection of the average silhouettes known by robot i is obtained as a matrix of projections, $matProys$. The projection of the average silhouette Sil_k^j is obtained in a vector, $vectProys$.
- The Euclidean distance (dE) is calculated between each vector of the matrix $matProys$ and the vector $vectProys$ as shown in formula 2, i.e. we obtain the distance between all the projections already computed and the projection of the new silhouette.

$$dE_l^i = \sqrt{\sum_{r=1}^{nEigens} (matProys_{(l,r)} - vectProys_{(1,r)})^2} \quad (2)$$

where $nEigens$ is the number of eigenvectors used during the PCA training ($nEigens = numObjs^i - 1$), and l is the index of the distance vector, where the maximum size of the vector dE^i is $numObjs^i$.

- The distance value dE^i is divided by a maximum distance value, $ThresholdMax$, determined experimentally to obtain a similarity metric also called the probability vector PCA, v_P^i , as shown in formula 3.

$$v_{P_i}^i = 1 - \frac{dE_l^i}{ThresholdMax} \quad (3)$$

If dE_l^i is bigger than the $ThresholdMax$ value, then the probability will be fixed as shown in formula 4, which indicates that the projections of the object k and the one of the object l are completely different.

$$v_{P_i}^i = \frac{1}{numObjs} \quad (4)$$

The value of the SIFT similarity metric also called the probability vector SIFT at the position $v_{S_i}^i$, is obtained calculating the number of coincident SIFT points, n_{coin} , between the individual SIFT concept $SIFT_l^i$ learned by robot i ,

and the individual SIFT concept $SIFT_k^j$ shared by robot j . If the number n_{coin} is bigger than an average of coincidences determined experimentally, $AverageCoin$, then the probability will be fixed to $v_{S_i}^i = 1.0$, which means that both concepts contain the same local SIFT features. Otherwise, the probability will be calculated using formula 5.

$$v_{S_i}^i = \frac{n_{coin}}{AverageCoin} \quad (5)$$

The constant $AverageCoin$ represents the average of coincidences between two sets of SIFT points of the same object from different perspectives.

4.2 Analysis and Fusion of Individual Concepts

This section describes how detection is achieved if the new or shared concept is one of the following: coincident, complementary or confused, and how the individual concepts are fused to form collective concepts depending on the kind of detected concept.

Coincident Concepts: A coincident concept is detected when one, two or more robots of the robot team learned individual concepts from similar views of the same object. A new or shared concept is classified as coincident if $v_{P_i}^i \geq \alpha$ and $v_{S_i}^i \geq \alpha$. That is, if both probabilities (PCA and SIFT) of a previously learned concept are greater than a predefined threshold value (α). If a new or shared concept is determined as coincident it is merged with the most similar known concept as follows:

PCA fusion: It is obtained by evaluating a new average silhouette from the average of the known Sil_i^i and new Sil_k^j silhouettes. After that, it is necessary to re-train the PCA substituting the concept Sil_i^i with the new average silhouette which contains information of the concept learned by robot j .

SIFT fusion: It is obtained by adding the complementary SIFT points of concept $SIFT_k^j$ to the set of SIFT points of concept $SIFT_i^i$. Also, each pair of coincident SIFT points of both concepts is averaged in terms of position and their corresponding SIFT descriptors.

The main idea to fuse coincident concepts is to improve their representation.

Complementary Concepts: A concept C_k^j contains complementary information if it differs with all known concepts by robot i , i.e., if both shape and local features are different to all known concepts by robot i , $C_1^i, \dots, C_{numObjs}^i$. That is, if $v_P^i < \alpha$ and $v_S^i < \alpha$.

A complementary concept C_k^j is fused with the collective concepts known by the robot i as follows:

PCA fusion: The new average silhouette is added and the new PC features are obtained by re-training the PCA using the updated set of average silhouettes.

SIFT fusion: The new SIFT features are simply added to the current set of SIFT concepts known by the robot i .

Confused Concepts: There are two types of confusion that can occur between concepts:

Different shape and similar local features (type 1): This type of confusion occurs when the new concept C_k^j is complementary by shape, Sil_k^j , to all the concepts known by the robot i , $Sil_1^i, \dots, Sil_{numObjects}^i$ but it is coincident by local SIFT features, $SIFT_k^j$, with at least one concept known by the robot i . That is, $v_{S_i}^i \geq \alpha$ and if $v_P^i < \alpha$.

Similar shape and different local features (type 2): This type of information occurs when concept C_k^j is coincident by shape, Sil_k^j , to at least one concept known by the robot i , but it is complementary using its local SIFT features, $SIFT_k^j$. That is, if $v_{P_i}^i \geq \alpha$ and $v_S^i < \alpha$.

In both types of confusion, type 1 or type 2, there can be two options:

- a) **Different objects:** Both concepts correspond to different objects.
- b) **Same object:** Both concepts correspond to the same object but they were learned by robots from different points of view.

In our current approach, both types of confusions are solved as complementary objects. The reason is that robot i cannot distinguish with its current information between both, different objects or same object, using only the individual and the shared concepts. To solve the ambiguity, as future work each robot should build autonomously a map and locate its position in the map. In addition, for each learned object, robots will locate them in the map. For confused objects a robot can move to the position of the object marked in the map to see the object from different perspectives in order to solve the conflict.

5 Experiments and Results

We performed several experiments to demonstrate the proposed algorithm. In section 5.1, we show the results of a general experiment that demonstrates the main features of the proposed approach. In section 5.2 we present the accuracy of the collective concepts versus the individual concepts.

In these experiments we used a robot team consisting of two homogeneous Koala robots (Figure 2) equipped with a video camera of 320×240 pixels, and an on board portable computer of 1 GB of RAM memory for the processing. For more than two robots our method can be applied straightforward. The only difference is that robots will need to consider the information from more than one robot, possibly reducing confused concepts.

5.1 Concept Acquisition and Testing

The mobile robots learn on-line representations of several objects while following a predefined trajectory without prior knowledge on the number or nature of the objects to learn. The idea of using pre-planned trajectories instead of making the robots wandering randomly, is that we can control the experimental conditions to show different aspects of the proposed methodology.

Each robot shares its individual concept as soon as it is learned to improve the representation of this concept or to include a new concept in the other robot.



Fig. 2. The robot team.

Figure 3 shows the training objects used in this experiment. As can be seen in the figure, some objects have the same shape but different texture, some have the same texture but different shape, some others are not symmetric in their shape. The objective of this experiment is to show the performance of the system to detect coincident, complementary and confused information under a wide variety of conditions.

In Figure 4 we present how the training objects were placed in the environment, and how the robots moved in the environment to see the different training objects. Robot 1 (R1) learned during individual training concepts for: *dolphin*, *can*, *water bottle* and *vase*. Robot 2 (R2) learned individual concepts for: *vase*, *soda bottle*, *bottle* and *cone*. Note that some objects are learned by both robots while others are only learned by one robot.



Fig. 3. Training objects. (a) *vase*, (b) *water bottle*, (c) *can*, (d) *dolphin*, (e) *soda bottle*, (f) *bottle* and (g) *cone*.

While learning a new concept, each robot has to decide whether to fuse the current concept with a previously known concept or include it as a new one. Table 1 shows the probability vectors of the PCA features based on shape (v_P^1) and of the SIFT features (v_S^1) obtained by Robot 1. A similar table is obtained by Robot 2 (for details see [16]). In Table 1 the coincident information is represented in bold.




Fig. 4. Concept acquisition.

Table 1. Probability vectors PCA (v_P^1) and SIFT (v_S^1) obtained by R1.

Objects	New (collective concepts R1)													
	PCA							SIFT						
	<i>Dol-phin</i>	<i>Vase</i>	<i>Can</i>	<i>Soda bottle</i>	<i>Water bottle</i>	<i>Bot-tle</i>	<i>Co- ne</i>	<i>Dol-phin</i>	<i>Vase</i>	<i>Can</i>	<i>Soda bottle</i>	<i>Water bottle</i>	<i>Bot-tle</i>	<i>Co- ne</i>
<i>Dol-phin</i> _{R1}	-	-	-	-	-	-	-	-	-	-	-	-	-	-
<i>Vase</i> _{R2}	0.19	-	-	-	-	-	-	0.09	-	-	-	-	-	-
<i>Can</i> _{R1}	0.31	0.26	-	-	-	-	-	0.12	0.12	-	-	-	-	-
<i>Soda bottle</i> _{R2}	0.36	0.28	0.58	-	-	-	-	0.28	0.11	0.40	-	-	-	-
<i>Water bottle</i> _{R1}	0.43	0.28	0.53	0.73	-	-	-	0.15	0.59	0.20	0.20	-	-	-
<i>Bottle</i> _{R2}	0.31	0.17	0.56	0.61	0.58	-	-	0.08	0.15	0.65	0.04	0.12	-	-
<i>Vase</i> _{R1}	0.25	0.69	0.42	0.43	0.41	0.32	-	0.16	1.00	0.23	0.10	0.08	0.09	-
<i>Cone</i> _{R2}	0.31	0.01	0.28	0.28	0.33	0.43	-	0.05	0.28	0.43	0.10	0.14	0.09	-

We used the defined criteria in Section 4.2 to recognize coincident, complementary or confused concepts, with $\alpha = 0.65$ as threshold value, and the probability vectors of Table 1.

For instance, in column labeled with PCA of Table 1 it is shown how the probabilities of objects of R1 are affected using only PCA over the shapes of objects, as both robots encounter and learn concepts while traversing the environment. In the first row, R1 learns about the concept *dolphin* and acquires it. In the second row, R2 then learns about *vase* and shares this concept to R1 (that is expressed with a subindex R2). The probability, according to the PCA features to be a *dolphin* is 0.19 (second row). R1 learns the object *can*, which has a probability of 0.31 to be a *dolphin* and a probability of 0.26 to be a *vase*, which was learned by R2 and shared to R1 (third row). In the fifth row, R1 learns about a *water bottle* but it confuses with the *soda bottle* learned and

shared before by R2. As can be seen from Figure 3, both objects have the same shape and consequently the PCA features are not able to discriminate between these two objects. This is not the case for the SIFT features, which prevent R1 to consider it as the same object. In the seventh row, R1 learns about *vase* which was already learned and shared by R2, and in this case both concepts are merged.

Robot 1 detects the three types of possible information and fuses them as mentioned in section 4.2. The first four learned objects (*dolphin*, *vase*, *can* and *soda bottle*) were detected as complementary because all values of their corresponding vectors $v_{P(x,l)}^1$ and $v_{S(x,l)}^1$ are less than α threshold. The next two objects (*water bottle* and *bottle*) were detected as confused (type 2 and type 1, respectively), because either their SIFT features, or their shape features are similar, according to the α threshold. Object *vase* learned by Robot 1 was detected as coincident with the learned object *vase* learned by Robot 2 and shared to Robot 1 ($v_{P(7,2)}^1 = 0.69$ and $v_{S(7,2)}^1 = 1.00$). Finally object *cone*, unnoticed for Robot 1, is detected as complementary.

To test the performance of the individual concepts and the collective concepts acquired by each robot, the concepts were used in an object recognition task. Each robot followed a predefined trajectory to recognize objects in the environment. The objects were detected by the robot team in the following order: *cone*, *water bottle*, *vase*, *bottle*, *soda bottle* and *dolphin*. Once an object is detected, the robot (i) evaluates its class using the PCA (v_P^i) and SIFT (v_S^i) probability vectors and combines both probabilities using a Bayesian approach, summarized in formula 6.

$$P_{B_l}^i = \frac{v_{P_l}^i \times v_{S_l}^i \times P_u}{\left(v_{P_l}^i \times v_{S_l}^i \times P_u\right) + \left((1 - v_{P_l}^i) \times (1 - v_{S_l}^i) \times (1 - P_u)\right)} \quad (6)$$

where P_u is a uniform probability distribution ($P_u = \frac{1}{numObjs^i}$), $v_P^i = p(PCA \text{ projection} \mid Class = i)$, $v_S^i = p(SIFT \text{ matching} \mid Class = i)$, P_B^i is the Bayesian probability vector ($p(Class = i \mid PCA \text{ projection}, SIFT \text{ matching})$), and l is the index of the Bayesian probability vector, where the maximum size of the probability vector is $numObjs^i$.

We show in Table 2 the precision of the object recognition task using the individual (columns labeled with R1 and R2) and collective (columns labeled with R1-R2 and R2-R1) concepts. Precision is calculated as the number of well classified images from the total number of test images that contain an object during one experiment. The precision is presented in two ways, one considering the total number of objects, and the other one taking into a count only the number of objects used during the individual training (reported in parentheses). As can be seen the collective concepts produce a significantly better precision.

5.2 Accuracy of the Individual and Collective Concepts

In this section we compare the results of the individual concepts with that of collective concepts. In each experiment, a different set of objects was used, and

Table 2. Precision in the object recognition task using the individual and collective concepts acquired by each robot.

	R1	R2	R1-R2	R2-R1
PCA	47.82 % (89.70 %)	52.12 % (98.00 %)	95.65 %	98.93 %
SIFT	50.00 % (93.87 %)	52.13 % (98.00 %)	95.65 %	98.93 %
Bayes	48.91 % (91.83 %)	52.12 % (98.00 %)	95.65 %	100.00 %

both robots learned the same set of objects. Therefore, all the shared concepts were coincident, that is, robots learned both individually and collectively the same number of concepts. At the end of each experiment the robots learned four concepts that were tested in a predefined sequence.

In Table 3, we present the average percentages of accuracy using the individual and collective concepts in an object recognition task. The accuracy indicates the average percentage of well classified images during the whole set of experiments, in this case six, when using the PCA and SIFT features and the Bayesian approach.

As it can be observed in Table 3, the accuracy that indicate the quantity of well classified images using the collective concepts for the object recognition task, is in general better than the accuracy using the individual concepts. For PCA, SIFT and Bayes there is an improvement in the accuracy of 2.56 %, 13.79 % and **20.62 %**, respectively. This demonstrates that the collective concepts have better coverage than the individual concepts because they contain information acquired from different points of view, which allows a better recognition of test objects.

In Table 4, we present the average percentages of false positives for both, the individual and the collective concepts acquired by the robots. From the results presented in Tables 3 and 4 we conclude that the collective concepts have better quality than the individual concepts.

In general for the individual and the collective concepts, we observed an improvement in the accuracy when using the Bayesian approach. The average profit in the percentages of classification using the Bayesian approach using the collective concepts with regard to the individual concepts is of **14.63 %**.

Table 3. Accuracy of individual and collective concepts.

	PCA	SIFT	Bayes
Individual	84.94 %	67.88 %	80.18 %
Collective	87.18 %	81.12 %	94.81 %

Table 4. Average percentages of false positives of individual and collective concepts.

	PCA	SIFT	Bayes
Individual	14.42 %	0.64 %	0.64 %
Collective	13.14 %	0.00 %	0.00 %

6 Conclusions and Future Work

In this paper we have introduced a new on-line learning framework for a team of robots. Some of the main features of the proposed scheme are:

- The robots do not know in advance how many objects they will encounter. This poses several problems as the robots need to decide if a new concept or a shared concept, is of a previously learned object or not.
- The representation of objects are learned on-line while the robots are traversing a particular environment. This is relevant for programming autonomous robots.
- Three possible cases in which to merge concepts and how to merge them were identified.

The detection of coincident concepts avoids producing multiple concepts for the same object. The detection of complementary concepts allows to detect and learn unknown objects not seen by a particular robot. The detection of confused concepts allows to fuse information: (i) when the objects have different shape and similar SIFT features, and (ii) when the objects have similar shape and different SIFT features. These cases are particularly difficult to deal with because the objects may be genuinely different or may be the same but seen from different points of view by the robots.

In general, the object recognition using the collective concepts had a better performance than using the individual concepts in terms of accuracy. This occurs because the collective concepts consider information from multiple points of view producing more general concepts.

As future work we propose to integrate schemes to object segmentation for dynamic environments. For instance, using an object segmentation based on distance as in [17]. Use a different set of features and identify possible conflicts between more than two kinds of features. We also plan to incorporate planning of trajectories to autonomously allocate the environment among robots. We also plan to add strategies to solve some confusions in shared concepts by taking different views from these objects. Finally, we plan to incorporate our algorithm for robot localization and search of objects, and to test our work for robot teams with three or more robots.

Acknowledgments. The first author was supported by the Mexican National Council for Science and Technology, CONACYT, under the grant numbers 212422 and 84162.

References

1. Asada M., Uchibe E., Noda S., Tawaratsumida S., Hosoda K.: Coordination of Multiple Behaviors Acquired by Vision-Based Reinforcement Learning. Proceedings of the International Conference on Intelligent Robots and Systems, pp. 917–924 (1994).
2. Mataric M. J.: Reinforcement Learning in the Multi-Robot Domain. *Autonomous Robots*, vol. 4 (1), pp. 73–83 (1997).
3. Howard A., Parker L. E., Sukhatme G. S.: Experiments with a Large Heterogeneous Mobile Robot Team: Exploration, Mapping, Deployment and Detection. *International Journal of Robotics Research*, vol. 25, pp. 431–447 (2006).
4. Parker L. E.: Distributed Algorithms for Multi-Robot Observation of Multiple Moving Targets. *Autonomous Robots*, vol. 3 (12), pp. 231–255 (2002).
5. Wessnitzer J., Melhuish C.: Collective Decision-Making and Behaviour Transitions in Distributed Ad Hoc Wireless Networks of Mobile Robots: Target-Hunting. *Advances in Artificial Life, 7th European Conference*, pp. 893–902 (2003).
6. Fernández F., Borrajo D., Parker L. E.: A Reinforcement Learning Algorithm in Cooperative Multi-Robot Domains. *Journal of Intelligent and Robotic Systems*, pp. 161–174 (2005).
7. Lowe D. G.: Distinctive Image Features from Scale Invariant Keypoints. *International Journal of Computer Vision*, vol. 60 (2), pp. 91–110 (2004).
8. Ekvall S., Jensfelt P., Kragic D.: Integrating Active Mobile Robot Object Recognition and SLAM in Natural Environments. *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 5792–5797 (2006).
9. Steels L., Kaplan F.: AIBO’s first words. *The Social Learning of Language and Meaning. Evolution of Communication*, vol. 4 (1), pp. 3–32 (2001).
10. Mitchell T. M.: *Machine Learning*. McGraw-Hill Science/Engineering/Math (1997).
11. Mitri S., Pervolz K., Surmann H., Nuchter A.: Fast Color Independent Ball Detection for Mobile Robots. *Proceedings of the IEEE International Conference Mechatronics and Robotics*, pp. 900–905 (2004).
12. Saebera Lopes L., Chauhan A.: How many words can my robot learn?. *Interaction Studies*, vol 8(1), pp. 53-81 (2007).
13. Montesano L., Montano L.: Identification of Moving Objects by a Team of Robots Based on Kinematic Information. *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, vol. 1, pp. 284–290 (2003).
14. O’Beirne D., Schukat M.: Exploration and Object Recognition Using Cooperative Robots. *Proceedings of the International Conference on Imaging Science, Systems and Technology*, pp. 592-598 (2004).
15. Cheung G. K. M., Baker S., Kanade T.: Visual Hull Alignment and Refinement Across Time: A 3-D Reconstruction Algorithm Combining Shape-From-Silhouette with Stereo. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, vol. 2, pp. 375–382 (2003).
16. Palacios-García A. C., Muñoz-Meléndez A., Morales E. F.: Collective Learning of Concepts using a Robot Team. *Proceedings of the 7th International Conference on Informatics in Control, Automation and Robotics*, vol. 2 (Robotics and Automation), pp. 79–88 (2010).
17. Méndez-Polanco J. A., Muñoz-Meléndez A., Morales E. F.: People Detection by a Mobile Robot Using Stereo Vision in Dynamic Indoor Environments. *Proceedings of the 8th Mexican International Conference on Artificial Intelligence*, pp. 349–359 (2009).