Research at the Learning and Vision Mobile Robotics Group

J. Andrade, A. Grau, J. Climent, F. Moreno, R. Alquézar,
F. Serratosa, J. Vergés, T. Vidal, Y. Bolea, and A. Sanfeliu
Institut de Robòtica i Informàtica Industrial, UPC-CSIC
Llorens Artigas 4-6, 08028 Barcelona, Spain
lrobots@iri.upc.es

Abstract

This article presents the current trends on wheeled mobile robotics being pursued at ESAII-IRI-UPC/CSIC. It includes an overview of recent results produced in our group in a wide range of areas, including robot localization, color invariance, segmentation, tracking, visual servoing, audio processing and object and face recognition.

Keywords: robot localization, color invariance, segmentation, tracking, recognition, non-speech sound.

1 INTRODUCTION

The Learning and Vision Mobile Robotics Group, an interdisciplinary team of researchers, is the product of a joint effort between the *Institut de Robòtica i Informàtica Industrial* and the *Departament d'Enginyeria de Sistemas, Automàtica i Informàtica Industrial* at the *Universitat Politècnica de Catalunya*, and the *Departament d'Enginyeria Informàtica i Matemàtiques* at the *Universitat Rovira i Virgili.*

Headed by Prof. A. Sanfeliu, as of today, it embraces 5 professors, 2 posdoctoral associates, and 3 PhD students. The group, consolidated in 1996, has given rise to 3 PhD thesis and 7 final year projects. Within the last 7 years, the group has published 8 peer reviewed journal articles, 7 book chapters, 7 conference proceeding editorials, and presented articles in over 50 international conferences (40 indexed in SCI), and 20 national conferences. Furthermore, within the past two years, the mobile robotics platforms developed in our group have been portrayed numerous times on live and printed media [6, 8, 12, 22, 23, 24].

Financial support comes mainly from a continued set of projects from the Interministerial Council of Science and Technology (CICYT) TAP96-0629-C04-03, TAP98-0473 and DPI2001-2223.

2 CURRENT RESEARCH AREAS

During the last year, our efforts have been tailored at giving our mobile platforms the ability to navigate autonomusly in unknown structured settings. In this sense, we have contributed new insight in the classical simultaneous localization and map building problem, from a control systems theory point of view [1, 2, 7]. Furthermore, we have developed new feature validation techniques that improve the robustness of typical map building algorithms [3, 4]. Also, we are currently investigating on the use of non-speech sound as yet another feature source to aid to robot localization [5].

Moreover, very good results have been achieved in the tracking of subjects under varying illumination conditions and in cluttered scenes. On the one hand, we have mastered now the use of histogram based techniques to color segmentation and illumination normalization [9, 25, 26]. On the other hand, we

1a Jornada de Recerca en Automàtica, Visió i Robòtica
pp.117-128. Barcelona, Feb. 3-4, 2004 ISBN 84-7653-844-8

have tested different statistical estimation paradigms to track subject candidates using not only color information but shape as well [13, 14]. Most of our video demonstrations from last year show results in this topic.

With respect to 3d object recognition and subject identification, we have formalized and validated a compact representation for a set of attributed graphs. We call this new formulation a *function described graph*, and it borrows the capability of probablisitic modelling from random graphs. FDG's are the result of a long time effort in the search for a compact representation of multiple view model descriptors [15, 18, 16, 17, 19]. Also, we have considered the use of oriented matroid theory for shape representation and the indexing of object views [20, 21, 10].

In the following sections we summarize the key contributions from five selected publications [1, 13, 18, 10, 5]. Each of them tackles very different problems typically encountered in mobile robotics applications.

3 SIMULTANEOUS LOCALIZATION AND MAP BUILDING

To univocally identify landmarks from sensor data, we study several landmark representations, and the mathematical foundation necessary to extract the features that build them from images and laser range data. The features extracted from just one sensor may not suffice in the invariant characterization of landmarks and objects, pushing for the combination of information from multiple sources.

Once landmarks are accurately extracted and identified, the second part of the problem is to use these observations for the localization of the robot, as well as the refinement of the landmark location estimates. We consider robot motion and sensor observations as stochastic processes, and treat the problem from an estimation theoretic point of view, dealing with noise by using probabilistic methods.

The main drawback we encounter is that current estimation techniques have been devised for static environments, and that they lack robustness in more realistic situations. To aid in those situations in which landmark observations might not be consistent in time, we propose a new set of temporal landmark quality functions, and show how by incorporating these functions in the data association tests, the overall estimation-theoretic approach to map building and localization is improved. The basic idea consists on using the history of data association mismatches for the computation of the likelihood of future data association, together with the spatial compatibility tests already available.

Special attention is paid in that the removal of spurious landmarks from the map does not violate the basic convergence properties of the localization and map building algorithms already described in the literature; namely, asymptotic convergence and full correlation.

We contribute also an in depth analysis of the fully correlated model to localization and map building from a control systems theory point of view. Considering the fact that the Kalman filter is nothing else but an optimal observer, we analyze the implications of having a state vector that is being revised by fully correlated noise measurements. We end up revealing theoretically and with experiments the strong limitations of using a fully correlated noise driven estimation theoretic approach to map building and localization in relation to the total number of landmarks used.

Partial observability hinders full reconstructibility of the state space, making the final map estimate dependant on the initial observations, and does not guarantee convergence to a positive definite covariance matrix. Partial controllability on the other hand, makes the filter beleive after a number of iterations, that it has accurate estimates of the landmark states, with their corresponding Kalman gains converging to zero. That is, after a few steps, innovations are useless. We show how to palliate the effects of full correlation and partial controllability.

Any map building and localization algorithm for mobile robotics that is to work in real time must be able to relate observations and model matches in an expeditious way. Some of the landmark compatibility tests are computationally expensive, and their application has to be carefully designed. We touch upon the time complexity issues of the various landmark compatibility tests used, and also on the desirable properties of our chosen map data structure. Furthermore, we propose a series of tasks that must be handled when dealing with landmark data association. From model compatibility tests, to search space reduction and hypothesis formation, to the actual association of observations and models.

Figure 1a shows some of the model compatibility heuristics devised for the validation of straight lines extracted from laser range data into walls. Frame b shows data as extracted from a laser range finder, and 2σ covariance ellipses around hypothesized landmark estimates. The third frame shows a virtual reality model of the map constructed during a run of the algorithm.

4 FUSION OF COLOR AND SHAPE FOR OBJECT TRACKING UNDER VARYING ILLUMINATION

Color represents a visual feature commonly used for object detection and tracking systems, specially in the field of human-computer interaction. For such cases in which the environment is relatively simple, with controlled lighting conditions and an uncluttered background, color can be considered a robust cue. The problem appears when we are dealing with scenes with varying illumination conditions and confusing background.

Thus, an important challenge for any color tracking system to work in real unconstrained environments, is the ability to accommodate variations in the amount of source light reflected from the tracked surface.

The choice of different color spaces like HSL, normalized color rgb, or the color space $(B-G, G-R, R+G+B)$, can give some robustness against varying illumination, highlights, interreflections or changes in surface orientation for an analysis of different color spaces. But none of these transformations is general enough to cope with arbitrary changes in

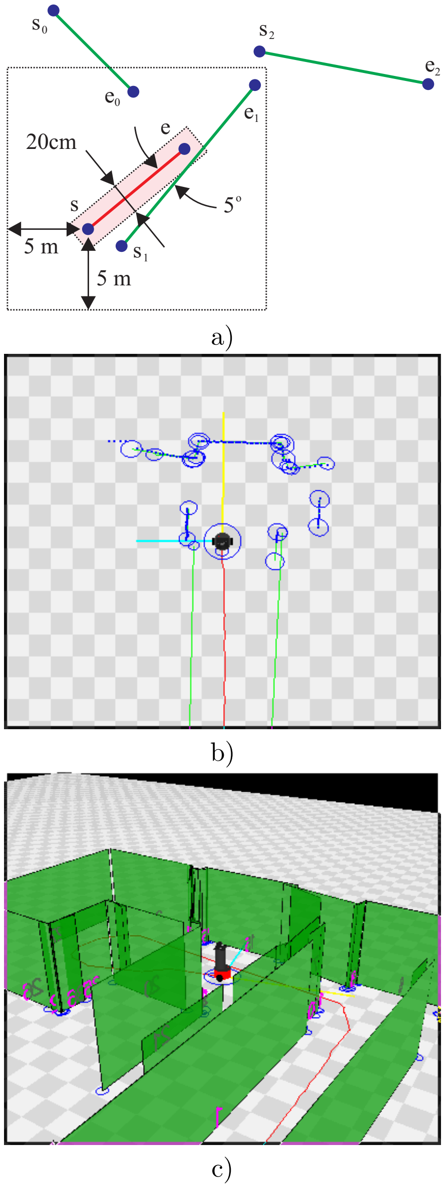

Figure 1: Concurrent mobile robot localization and map building. a) Hypothesis search range for walls extracted from a laser range scan. b) The blue dots indicate sensor raw data coming from a laser range finder. The green lines represent walls inferred from consecutive readings. The red lines indicate the estimated robot trajectory. c) Graphical representation of the map built.

illumination.

Instead of searching for color constancy, other approaches try to adapt the color distribution over time. One such technique is to use Gaussian mixtures models to estimate densities of color, and under the assumption that lighting conditions change smoothly over time, the models are recursively adapted. Another option is to parameterize the color distribution as a random vector and to use second order Markov model to predict the evolution of the corresponding color histogram. These techniques perform much better than the mere change of color space representation, but have the drawback that they do not check for the goodness of the adaptation, which can still lead to failure.

The fusion of several visual modules using different criteria offers more reliability than methods that only use one feature. As an example, systems that track in real-time an individual might model the head of a person by an ellipse and use intensity gradients and color histograms to update the head position over time. In [14], color histograms are fused with stereo-vision information in order to dynamically adapt the size of the tracked head. These real time applications however are constrained only to the tracking of elliptical shapes.

A new methodology that addresses the problems present in the approaches described above, results in a robust tracking system able to cope with cluttered scenes and varying illumination conditions. The fusion is done using the CONDENSATION algorithm that formulates multiple hypothesis about the estimation of the object's color distribution and validates them taking into account the contour information of the object [11].

Four sets of sequence results are summarized in Figure 2 to illustrate the robustness of our system under different conditions. In the first experiment we show how our system is able to accommodate color by applying it over a synthetic sequence of circles moving around and changing randomly its color. In the upper left frame of the figure the path of the color distributions for the tracked circle is shown. The second experiment is to track a colored rectangle. It has to be pointed out that in the previous experiment we used the *RGB* color space, but in the present and subsequent experiments the color space used was the $(B - G, G - R, R + G + B)$ in order to provide robustness to specular higlights. The last two experiments, correspond to outdoor scenes, where although the change in illumination conditions is limited, they are useful to show that our method works with non-uniform shapes (third experiment of a beatle tracking), and in cluttered scenarios (fourth experiment of a snail tracking).

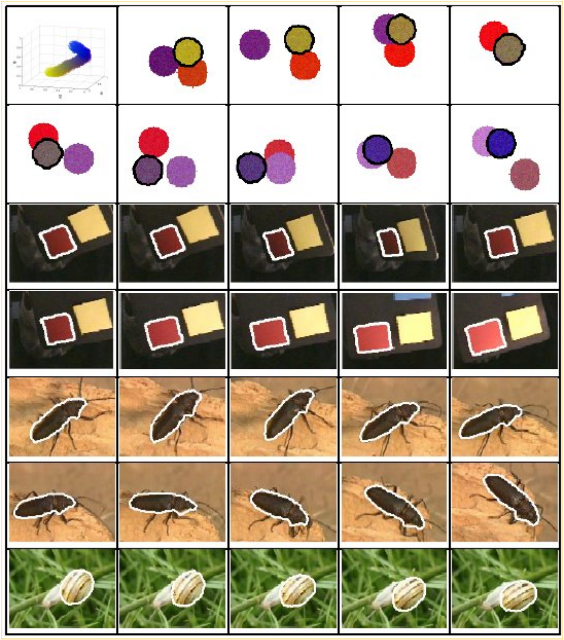

Figure 2: Results of the 4 experiments.

5 FUNCTION-DESCRIBED GRAPHS FOR MODELLING OBJECTS REPRESENTED BY SETS OF ATTRIBUTED GRAPHS

A *function-described graph* (FDG) is a model that contains probabilistic and structural descriptions of a set of attributed graphs (AGs) to maintain, to the most, the local features of the AGs that belong to the set and other AGs that are "near" them, as well as to allow the rejection of the AGs that do not belong to it or are "far" of them. Let us consider, as an example, the 3D-object modelling and recognition problem. The basic idea is that only a single FDG is synthesised from the graphs that represent several views of a 3D-object. Therefore, in the recognition process, only one comparison is needed between each model represented by an FDG and the unclassified object (view of a 3D-object) represented by a graph.

Random graphs are on the other hand, one of the earliest approaches used to represent a set of AGs. In this approach, AGs are extended to include probabilistic information. Wong et al. first defined the General Random Graphs (GRGs) for modelling classes of patterns described by AGs through a joint probability space of random variables ranging over pattern primitives (vertices) and relations (arcs). Due to the computational intractability of GRGs, caused by the difficulty in estimating and handling the high-order joint probability distribution, First-Order Random Graphs (FORGs) were proposed for real applications. Strong simplifications were made in FORGs to allow the use of random graphs in practical cases; more precisely, the following assumptions were made: a) the random vertices are mutually independent; b) given values for the random vertices, the random arcs are independent; c) the arcs are independent of the vertices except for the vertices that they connect.

FDGs can be seen as a type of simplification of the GRGs, different from FORGs, in which some structural constraints are recorded. A drawback of FORGs is that the strong assumptions about the statistical independence of nodes and arcs may lead to an excessive generalisation of the sample graphs when synthesising a FORG. To alleviate this weakness, a qualitative information of the joint probabilities of two nodes is incorporated into FDGs, thus improving the representational power of FORGs with a negligible increase of computational cost.

5.1 APPLICATION OF FDGS FOR MODELLING AND RECOGNITION OF OBJECTS

FDGs are applied here to 3D-object representation and recognition. The attribute of the vertices is the average hue of the region (cyclic range from 0 to 49) and the attribute of the edges is the difference between the colours of the two neighbouring regions.

We first present an experimental validation of FDGs using artificial 3D-objects in which the adjacency graphs have been extracted manually and afterwards we present a real application on an image database in which the graphs have been extracted automatically. The advantages of the experimental validation are that the results do not depend on the segmentation process and that we can use a supervised synthesis, since we know which vertices of the AGs represent the same planar face of the object. Thus, we can discern between the effects of computing a distance measure using different values of the costs on the 2nd-order relations. In the real application, we show the capacity of FDGs to keep the structural and semantic knowledge of an object despite the noise introduced by the segmentation process and an automatic synthesis.

5.1.1 SUPERVISED SYNTHESIS ON ARTIFICIAL OBJECTS

We designed five objects using a CAD program (see Figure 3). After that, we took five sets of views from these objects and from these views we extracted a total of 101 adjacency graphs. To obtain the AGs of the test set and of the reference set, we modified the attribute values of the vertices and arcs of the adjacency graphs by adding zero-mean Gausian noise with different variances. Moreover, some vertices and arcs where inserted and deleted randomly in some cases. The FDGs were synthesised using the AGs that belonged to the same 3D-object and using the synthesis given a common labelling from a set of AGs described in.

Table 1 shows the ratio of correctness for different levels of noise and applying several costs on the antagonisms and occurrences. We see that the bests

Table 1: FDGs ratio of correctness.

Num. Vertices	**Ins & Del**	**0**	**0**	**0**	**0**	**0**	**1**	**2**	**1**
Standard	**Deviation**	**0**	**2**	**4**	**8**	**12**	**0**	**0**	**8**
Cost Antag	**Cost Occu**								
Moderate	Moderate	100	98	97	95	92	89	85	83
High	Null	100	92	89	87	84	61	54	57
Null	High	100	91	89	88	85	62	59	59
High	High	100	95	90	86	80	60	53	56
Moderate	Null	100	92	91	91	87	80	75	75
Null	Moderate	100	95	92	91	86	81	77	76
Null	Null	100	90	89	88	86	70	67	68

Table 2: FDGs (moderate 2nd order costs), FORGs and 3-NN ratio of correctness.

Num Vertices Ins. or Del.	**0**	**0**	**0**	**0**	**0**	**1**	**2**	**1**
Standard Deviation	**0**	**2**	**4**	**8**	**12**	**0**	**0**	**8**
FDGs (Moderate costs)	100	98	97	95	92	89	85	83
Random Graphs (FORGs)	100	90	89	88	86	70	67	68
3-N.N. (Edit Op. Distance)	100	98	82	62	52	90	58	58

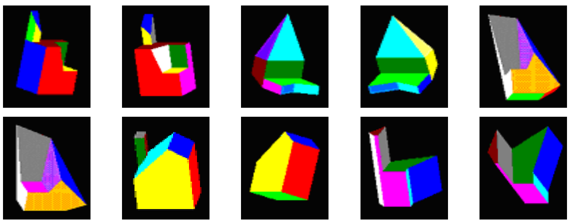

Figure 3: 10 views of 5 artificial objects designed with a CAD program.

results appear always when we use moderate 2nd-order costs. Furthermore, when noise increases, the recognition ratio decreases drastically when we use high costs but there is only a slight decrease when we use moderate costs. Moreover, in Table 2 we compare the FDGs (with 2nd-order costs) to other two methods. The FDG classifier always obtains better results than the 3-Nearest Neighbours and the Random Graph classifiers. The difference of the ratio between FDGs and the other two methods increases when the noise also increases. FORGs obtain better results than the 3-N.N. only when the noise is high.

5.1.2 UNSUPERVISED SYNTHESIS ON REAL LIFE OBJECTS

Images were extracted from the database COIL-100 from Columbia University. It is composed of 100 isolated objects and for each object there are 72 views (one view each 5 degrees). Adjacency graphs are obtained by color segmentation. Figure 4 shows 20 objects at angle 100 and their segmented images with the adjacency graphs. The test set was composed by 36 views per object (taken at the angles 0, 10, 20 and so on), whereas the reference set was composed by the 36 remaining views (taken at the angles 5, 15, 25 and so on). FDGs were synthesised automatically using the AGs in the reference set that represent the same object. The method of incremental synthesis, in which the FDGs are updated while new AGs are sequentially presented, was applied. We made 6 different experiments in which the number of FDGs that represents each 3D-object varied. If the 3D-object

Figure 4: The 20 selected objects and the segmented images with the AGs.

was represented by only one FDG, the 36 AGs from the reference set that represent the 3D-object were used to synthesise the FDG. If it was represented by 2 FDGs, the 18 first and consecutive AGs from the reference set were used to synthesise one of the FDGs and the other 18 AGs were used to synthesise the other FDG. A similar method was used for the other experiments with 3, 4, 6 and 9 FDGs per 3D-object.

Similarly to the previous experimental results, the correctness is higher when 2nd-order relations are used with a moderate cost. The best result appears when each object is represented by 4 FDGs, that is, each FDG represents 90 degrees of the 3D-object. When objects are represented by 9 FDGs, each FDG represents 40 degrees of the 3D-object and 4 AGs per FDG, there is poor probabilistic knowledge and therefore the costs on the vertices and arcs are coarse. Moreover, when objects are represented by only 1 or 2 FDGs, there are too much spurious regions (produced in the segmentation process) to keep the structural and semantic knowledge of the object.

6 ORIENTED MATROID THEORY FOR SHAPE REPRESENTATION

In the referenced work [10], an indexing technique that combines region connection calculus and oriented matroid theory is presented. More precisely, the type of connectivity between connected regions of views is described by means of the formalism of the region connection calculus, whereas the topological properties of the disconnected regions of the views are encoded into a data structure called set of cocircuits. The set of cocircuits, that is one of the several combinatorial data structure referred to as oriented matroids, encodes incidence relations and relative position of the elements of the image and gives local and global topological information about their spatial distribution. Reasoning with region connection calculus is based on composition tables, while the other representation permits algebraic techniques to be used. These two descriptions merged together are used as an index of the database. Since index tables are by definition discrete, the discrete nature of these two representations nicely fits with this technique.

This indexing method is employed for hypothesis generation in 3D object recognition from single views and can be regarded as a qualitative counterpart of the geometric hashing technique.

Extracting the oriented matroid of a view is not straightforward since the regions that form an image cannot be reduced to points, taking for instance their centroids, without losing essential topological information for object recognition. Therefore, in the method the convex hull of each region is used to represent the region itself. Then, pairs of the resulting convex polygons are considered and the oriented matroid is extracted based on the spatial location of the other convex regions of the image with respect to two lines arising in the process of merging the convex hulls of pairs non-overlapped regions. Consider, for instance, the ordered pair of convex regions (S, T) of the object 1 of Fig. 5. It is easy to see that the convex hull of these two convex planar non-overlapped polygons is a polygon whose set of vertices is included in the union of the set of vertices of S and T. On the contrary, the set of edges of the convex hull of S and T is not included in the union of their set of edges. Indeed, two new "bridging edges," e_1 and e_2, appear as illustrated in Fig. 5.a. Actually, efficient algorithms for merging convex hulls are based on finding these two edges .

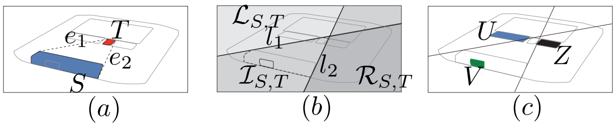

Figure 5: Steps of encoding of the combinatorial properties of a view of an object into a chirotope.

6.1 INDEXING VIEWS OF 3D OBJECTS

The process of indexing a database of views of a set of objects starts with some preliminary choices, namely the number of colors in which the hue is quantized and the number of regions having the same color that will be taken into account. These choices, of course, depend on the properties of the images of the database. Then, the views are segmented according to these choices and the convex hull of each region is computed. As a consequence, the resulting image is a composition of convex polygonal regions that can be disjoint or partially or completely overlapped. In Fig. 6 two views of two objects are represented in which a color quantization with 6 colors white (W), red (R), yellow (Y), green (G), blue (B) and black (N) has been applied and up to two regions with the same color are taken into account. Let (W, R, Y, G, B, N) be the ordered tuple of colors considered.

The type of connection between the existing regions is described using the formalism of the region connection calculus. For each pair of disconnected regions the set of cocircuits of the view is computed. This is done for each view of the database. Then, this information extracted from the views of the database is combined and used for indexing a unique table whose entries are spatial combination of features of regions and the records contains the views that contains that combination.

In Table 3 an example of index table of the views 1 of the two objects 1 and 2 represented in Fig.6, denoted $v_{1,1}$ and $v_{1,2}$, respectively, is reported.

In the first column in correspondence with a certain ordered couple of regions (S, T) the relation between pairs of regions S and T in terms of region connection calculus is reported. When S and T are connected the set of cocircuits is not computed, and the symbol "$\emptyset$" means that the combination is not present in any view; in both cases the corresponding row is empty. When S and T are disconnected the set of cocircuits is computed and the symbol "$*$" in correspondence with a certain feature means that it is not present in the view.

Given a database of views of a set of 3D objects and a view v_i of one of them, not necessarily contained in the database, its set of cocircuits is computed. Each cocircuit is used to access the table that constitutes the index of the database.

For each view j of the object k, $v_{j,k}$, found at that address of the table, the elements (i, k) of an image-object views association matrix are increased of 1.

Figure 6: Two views of two objects whose topological properties are indexed in Table 3.

The final result of the indexing is therefore an association matrix in which the value of the element (i, k) indicates the strength of the hypothesis of associating the image v_i with the object k of the database.

In order to find the view of the database that best matches the image v_i, a similar strategy can be employed. In both cases the view v_i will be associated with the view or with the object that has the maximum number of correspondences with v_i in terms of cocircuits. It is easy to see that this method for hypothesis generation, that can be regarded as a qualitative version of the geometric hashing technique, is also robust to partial occlusions ofthe objects. Indeed, if a region of an image is occluded the set of cocircuits can still be computed and, therefore, the number of correspondences with the views or with the objects of the database can still be calculated. In this case, obviously, the selectivity of the indexing decreases.

Different advances of this work have been published in [20], [21] and [10].

7 NON-SPEECH SOUND FEATURE EXTRACTION

Sound offers advantages for information systems in delivery of alerts, information duration, encoding of rapidly incoming information, representing position in 3-D space around a person and his localization. Hearing is one of the human being most important senses. After vision, it is the sense most used to gather information about our environment. Despite this, little research has been done in the use of sound by a computer to study its environment. The research that has been done focuses mainly on speech recognition, while research into other types of sound recognition has being neglected. In robotics, non-speech audio has been ignored in front of artificial vision, laser beams and mechanical wave sensors beyond the audible spectrum. But the study and modeling of non-speech audio can help greatly in robot navigation and localization in the space domain. The existing research in non-speech sound is incipient and focuses on signal processing techniques for feature extraction with the use of neural networks as a classification technique. We propose a new technique based on pattern recognition techniques in order to locate a robot in the space domain by non-speech audio signals. The feature space will be built with the coefficients of model identification of audio signals. Due to their non-stationary property, wavelet decomposition is needed as a preprocessing step. We also propose a technique (transform function) to convert the samples in the feature space into the space domain, based in the sound derivative partial equation.

We propose a new approach to localization in the space domain using non-speech audio signals that will be applied on a robot in an industrial environment, the approach follows the next steps: 1) measurement

	Connection	W	R	Y	G_1	G_2	B_1	B_2	N	Objects
WR	DC	0	0	$*$	0	0	0	-	+	$v_{1,1}$
WY	DC	0	$*$	0	0	$*$	0	0	-	$v_{1,2}$
WG_1	NTPP									$v_{1,1}$
WG_1	DC	0	$*$	0	0	$*$	0	0	0	$v_{1,2}$
WG_2	DC	0	0	$*$	0	0	+	0	0	$v_{1,1}$
WB_1	DC	0	0	$*$	0	0	0	0	0	$v_{1,1}$
WB_1	NTPP									$v_{1,2}$
WB_2	DC	0	0	$*$	+	+	+	0	+	$v_{1,1}$
WB_2	NTPPi									$v_{1,2}$
WN	DC	0	0	$*$	-	-	-	-	0	$v_{1,1}$
WN	DC	0	$*$	+	+	$*$	0	0	0	$v_{1,2}$
RY	$\emptyset$									
RG_1	NTPP									$v_{1,1}$
···	···	···	···	···	···	···	···	···	···	···
B_2N	DC	+	0	$*$	-	-	-	0	0	$v_{1,1}$
B_2N	DC	-	$*$	+	+	$*$	+	0	0	$v_{1,2}$

Table 3: Index of the topological properties of the two views $v_{1,1}$ and $v_{1,2}$ of the two objects represented in Fig. 6.

and data preprocessing. 2) MAX model signal identification by wavelet transformation; 3) feature selection, feature extraction and its correspondence with the space domain. Non-speech audio signals generated by any audio source (industrial machinery, appliances, etc.) are continuous by nature. Preliminary, non-speech signal preprocessing includes sampling the analog audio signal with a specific frequency and to convert it into a discrete set of samples. The sampling interval should be chosen in such a way that essential information be preserved.

Some experiments have been carried out with our robot MARCO in a real environment with a CNC milling machine as a non-speech audio source. The rate of spatial recognition for unknown positions is near 95% in average respect to the actual position.

With the methodology used in this work we have achieved some interesting results that encourage us to continue investigating in this research field. The introduction of more that one audio source is a new challenge. The experimental results show a narrow correspondence with the sound physical model and this demonstrates a high reliability of the proposed methodology. This work has been presented in [5].

References

[1] J. Andrade-Cetto. *Environment Learning for Indoor Mobile Robots.* PhD thesis, UPC, Barcelona, Apr. 2003.

[2] J. Andrade-Cetto and A. Sanfeliu. Concurrent map building and localization on indoor dynamic environments. *Int. J. Pattern Recogn.*, 16(3):361–374, May 2002.

[3] J. Andrade-Cetto and A. Sanfeliu. Concurrent map building and localization with landmark validation. In *Proc. 16th IAPR Int. Conf. Pattern Recog.*, volume 2, pages 693–696, Quebec, Aug. 2002. IEEE Comp. Soc.

[4] J. Andrade-Cetto and A. Sanfeliu. Concurrent map building and localization with temporal landmark validation. In *Proc. IEEE Int. Conf. Robot. Automat.*, pages 1576–1581, Taipei, Sep. 2003.

[5] Y. Bolea, A. Grau, and A. Sanfeliu. Non-speech sound feature extraction based on model identification for robot navigation. In A. Sanfeliu and J. Ruiz-Schulcloper, editors, *Progress in Pattern Recognition, Speech and Image Analysis*, volume 2905 of *Lect. Notes Comput. Sci.*, Havana, Nov. 2003.

[6] Canal Blau TV. *The IRI Robots in Canal Blau*, May 2002.

[7] A. Checa, J. Andrade, and A. Sanfeliu. Construcció de mapes per robots móbils equipats amb sensors làser de profunditat. Technical Report IRI-DT-03-01, IRI, UPC, Mar. 2003.

[8] El Periódico de Catalunya. *El Robot Doméstico Llama a la Puerta*, Apr. 2002.

[9] A. Grau, J. Climent, F. Serratosa, and A. Sanfeliu. Texprint: A new algorithm to discriminate textures structurally. In T. Caelli, A. Amin, R. P. W. Duin, M. Kamel, and D. de Ridder, editors, *Proc. IAPR Int. Workshop Syntactical Structural Pattern Recog.*, volume 2396 of *Lect. Notes Comput. Sci.*, pages 368–377, Windsor, Aug. 2002. Springer-Verlag.

[10] A. Grau, F. Serratosa, and A. Sanfeliu. Object indexing based on oriented matroid theory and region connection calculus for scene recognition. Accepted for publication in *Int. J. Discrete Applied Mathematics.*, 2004.

[11] M. Isard and A. Blake. Condensation - conditional density propagation for visual tracking. *Int. J. Comput. Vision*, 29(1):5–28, Aug. 1998.

[12] La Politècnica, Num 4. UPC. *Criatures Cibernètiques*, Nov. 2001.

[13] F. Moreno, J. Andrade-Cetto, and A. Sanfeliu. Fusion of color and shape for object tracking under varying illumination. In F. J. Perales, A. J. C. Campilho, N. P. de la Blanca, and A. Sanfeliu, editors, *Proc. 1st Iberian Conf. Pattern Recog. Image Anal.*, volume 2652 of *Lect. Notes Comput. Sci.*, pages 580–588, Puerto de Andratx, Jun. 2003. Springer-Verlag.

[14] F. Moreno, A. Tarrida, J. Andrade-Cetto, and A. Sanfeliu. 3d real-time head tracking fusing color histograms and stereovision. In *Proc. 16th IAPR Int. Conf. Pattern Recog.*, volume 1, pages 368–371, Quebec, Aug. 2002. IEEE Comp. Soc.

[15] A. Sanfeliu, R. Alquézar, J. Andrade, J. Climent, F. Serratosa, and J. Vergés. Graph-based representations and techniques for image processing and image analysis. *Pattern Recogn.*, 35:639–650, Mar. 2002.

[16] F. Serratosa, R. Alquézar, and A. Sanfeliu. Estimating the joint probability distribution of random vertices and arcs b means of second-order random graphs. In T. Caelli, A. Amin, R. P. W. Duin, M. Kamel, and D. de Ridder, editors, *Proc. IAPR Int. Workshop Syntactical Structural Pattern Recog.*, volume 2396 of *Lect. Notes Comput. Sci.*, pages 252–260, Windsor, Aug. 2002. Springer-Verlag.

[17] F. Serratosa, R. Alquézar, and A. Sanfeliu. Synthesis of function-described graphs and clustering of attributed graphs. *Int. J. Pattern Recogn.*, 16(6):621–655, 2002.

[18] F. Serratosa, R. Alquézar, and A. Sanfeliu. Function-described graphs for modeling objects represented by sets of attributed graphs. *Pattern Recogn.*, 36:781–798, 2003.

[19] F. Serratosa, A. Sanfeliu, and R. Alquézar. Modelling and recognizing 3d-objects described by multiple views using function-described graphs. In *Proc. 16th IAPR Int. Conf. Pattern Recog.*, volume 2, pages 140–143, Quebec, Aug. 2002. IEEE Comp. Soc.

[20] E. Staffetti, A. Grau, F. Serratosa, and A. Sanfeliu. Oriented matroids for shape representation and indexing. In F. J. Perales, A. J. C. Campilho, N. P. de la Blanca, and A. Sanfeliu, editors, *Proc. 1st Iberian Conf. Pattern Recog. Image Anal.*, volume 2652 of *Lect. Notes Computs. Sci.*, pages 1012–1019, Puerto de Andratx, Jun. 2003. Springer-Verlag.

[21] E. Staffetti, A. Grau, F. Serratosa, and A. Sanfeliu. Shape representation and indexing based on region connection calculus and oriented matroid theory. In G. Saniti di Baja, editor, *Proc. 11th Int. Conf. Discrete Geom. Comput. Imagery*, volume 2886 of *Lect. Notes Comput. Sci.*, pages 267–276, Naples, Nov. 2003. Springer-Verlag.

[22] Televisió de Catalunya. *Robot Casolà*, May 2003.

[23] Televisión Española. *Informe Semanal*, Jan. 2002.

[24] UPC Libraries. *More Than Just Machines*, Jan. 2003.

[25] J. Vergés-Llahí and A. Sanfeliu. Colour constancy algorithm based on the minimization of the distance between colour histograms. In F. J. Perales, A. J. C. Campilho, N. P. de la Blanca, and A. Sanfeliu, editors, *Proc. 1st Iberian Conf. Pattern Recog. Image Anal.*, volume 2652 of *Lect. Notes Comput. Sci.*, pages 1066–1073, Puerto de Andratx, Jun. 2003. Springer-Verlag.

[26] J. Vergés-Llahí, A. Tarrida, and A. Sanfeliu. New approaches to colour histogram adaptation in face tracking tasks. In *Proc. 16th IAPR Int. Conf. Pattern Recog.*, pages 381–384, Quebec, Aug. 2002. IEEE Comp. Soc.