

A Comprehensive Framework for Perception in Robotic Soccer

Luciano Oliveira¹, Augusto Loureiro² and Leizer Schnitman²

¹*Institute of System and Robotics
University of Coimbra,
Portugal*

²*Programme of Master in Mechatronic
Federal University of Bahia,
Brazil*

1. Introduction

The main difficulties of perception systems, mainly to be applied to robots which play soccer, may be listed as: recognition of objects in a very short time and provision of accurate information to the control system. To accomplish these tasks, many researches in the sensor fusion field have been carried through with the objective of determining complementary or redundant information of the world.

Throughout the years, several works have been proposed to tackle the problem of integrating sensor data to enhance the recognition performance of robotic systems (Bai et al, 2003; Fanny et al, 2004; Lanthier et al, 2004). Bai et al (2003) propose a fusion strategy based on Gaussian distribution over the space of robot position, with the main goal of estimating the robot position, using only range sensors. As the proposed method is based on an asynchronous sensor fusion, and it depends essentially on robot's dead reckoning, the strategy fails whereas the robot runs long distances. Ferrein et al (2005) combines sensor fusion techniques to estimate ball position in a robot soccer field, namely, a weight grid and a Kalman filter strategies. A comparative study was made and the proposed method overperformed traditional ones, although the work is very Robocup domain-specific. Lanthier et al (2004) use data from inexpensive sensors (sonars and infrared (IR) sensors) to enhance the accuracy of a stereo-based system, whose strategy of data fusion relies on an occupancy grid technique and a Kalman filter.

All those works propose different ways to combine sensor data in order to improve sensing performance. Considering all these elements, a generic framework which copes with the problem of perceiving objects in front of a mobile robot is then proposed. The framework has been applied in Robocup domains, but it can be easily transferred into other robotic fields, under just few adjustments.

The aim of this chapter is to present the proposed framework, which tries to establish an as much as possible tradeoff between accuracy and on-the-fly information. The system consists of a vision servoing module, presented on the top of the robot, in charge of perceiving the

Source: Robotic Soccer, Book edited by: Pedro Lima, ISBN 978-3-902613-21-9, pp. 598, December 2007, Itech Education and Publishing, Vienna, Austria

relevant objects (ball and robots), and a set of IR distance sensors whose data are fused with the vision information in order to achieve the spatial location of objects in front of the robot. In the vision system, a cascade of boost rejection (Viola & Jones, 2001) with a Support Vector Machine (SVM) (Vapnik, 1995), at the final stage, are used to guarantee a more accurate classification, in a short time. To integrate all sensor data, a Takagi-Sugeno (TS) fuzzy logic based system tries to balance the better situation in which each sensor data may be used or integrated, and gives the final spatial location of the objects in the soccer field. The proposed framework brings threefold contributions: i) a low computation cost perception system; ii) to the best we know, a novel calibration system, with the use of a regression SVM to obtain a mapping between the world and image coordinate systems, without the need of a rigid transformation scheme, and iii) a fusion system which provides a robust sensor integration. A thorough analysis of each module has been carried through and results have been shown to highlight the proposed framework characteristics.

The rest of the chapter is structured as follows: in Section 2, the system architecture of the soccer robot used and the overall structure of the perception framework are given; in Section 3, a brief overview of the image classification methods is presented; in Section 4, the fuzzy engine used into the framework is discussed; in Section 5, the novel calibration method and its performance analysis are detailed; Section 6 shows some experimental results. Finally, Section 7 draws some conclusions and future works.

2. System Architecture of the Soccer Robot

In Fig. 1, the hardware architecture of the soccer robot, which have been worked with, is depicted. It is essentially a three wheel robot base with a set of IR sensors and a two degrees of freedom (DOF) servoing vision head.



Fig. 1. Hardware architecture of the soccer robot

So that robots play soccer, integration architecture of the perception, control and navigation modules is necessary. Thereby, the architecture proposed in (Costa & Bittencourt, 1999) is used as dorsal spine of aggregation of all robot modules. This architecture is illustrated in Fig. 2.

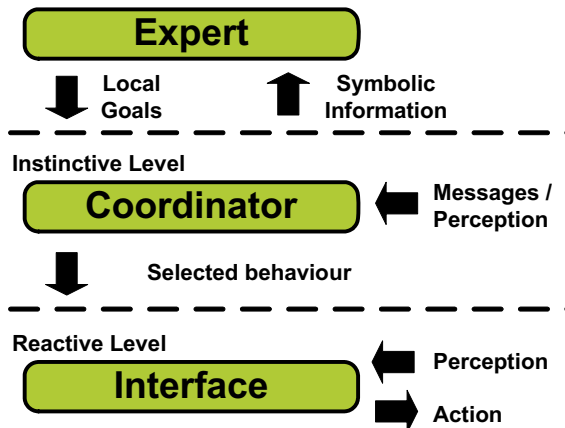


Fig. 2. Software architecture of our soccer robot

The *cognitive* module is a system based on symbolic knowledge which handles information come from *instinctive* level, as well as, asynchronous messages received from other robots (autonomous agents). This module gives global and local goals to the robot, as output. The *instinctive* level is in charge of identifying environment states and choosing the more appropriate behaviour for the robot’s current state and goals. The *reactive* level communicates directly to the perception system, receiving a frame containing information about spatial location and velocity of the targets (ball and robots) in front of the robot. Fig. 3 illustrates how this process works. Detailed information of the frame contents (Fig. 3) given by perception system to the control system is discussed in Section 4.

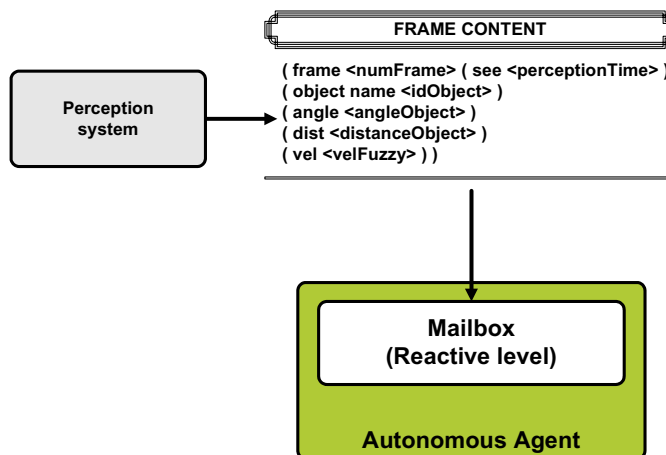


Fig. 3. Communication protocol between the perception system and the reactive level

2.1 The Perception Framework: Overall Architecture

The perception system consists of two threaded modules. After acquiring sensor data, each frame and a respective set of distance data are processed in the specific modules. The objective of this step is to prepare these simple data for the data fusion stage. Fig. 4 shows this process.




Fig. 4. Communication protocol between the perception system and the reactive level

Particularly, in the vision-based module, each acquired image goes to a feature extractor in order to achieve Haar-like features (Viola & Jones, 2001). These features are, then, classified using a cascade of weak classifier (Adaboost), with an SVM at the end, validating all the classification process. Then, after converting the raw information come from the IR distance and synchronizing it by a timing line strategy, all this information goes to a fuzzy processor, giving object information in form of frames (Minsk, 1975), illustrated in Fig. 3.

In the next sections, detailed information about the perception system, as well as, for completeness, a brief overview of the classification methods used is given.

3. Overview of the Classification Methods

In this section, an overview of the image classification methods applied in the vision system is highlighted, as well as, the way these methods have been combined in order to provide a more reliable object recognition system. This new approach represents an advance with respect of our first implementation in (Oliveira et al, 2005).

3.1 Adaboost

In (Viola & Jones, 2001), a complete system for face recognition is presented. Motivated by this approach, we decided to use Haar-like features in the same way to recognize the ball and robots in the soccer field by means of an SVM, at the end, in order to turn the vision task more robust. With the use of an SVM classifier is intended to decrease the high number of false alarm the Adaboost classifier, against Haar-like features, is prone to. For completeness, we present a brief description of the method and the way it has been used.

Haar wavelate templates (or Haar-like features) have been firstly used in (Oren et al, 1997), along with an SVM classifier, in order to recognize people. The set of these templates was

further expanded in (Viola & Jones, 2001) to incorporate other kinds of contrast differences. Fig. 5 illustrates the common set of the templates.




Fig. 5. (a), (b) and (c) are templates for line features; (e) and (f) are for edge features and (g) is the center-surrounded feature

These feature templates are regard to contrast differences in the image pixels. Fig. 6 shows how these templates are applied in the proposed system.




Fig. 6. Templates used in the image pixels

Haar-like features are extracted from the image in an overlapping way. Hence, this approach leads to an overcompleted set of features feasible to be applied in object categorization tasks (Viola & Jones, 2001). For each feature, in the image, a weak classifier $h_j(x)$ is trained and has the following form:

$$h_j(x) = \begin{cases} 1, & \text{if } p_j f_j < \varphi_j \\ 0, & \text{otherwise} \end{cases} \quad (1)$$

where x is an $m \times n$ image subwindow which consists of a feature f_j , a threshold φ_j and a parity p_j which indicates the direction of inequality. Each of $h_j(x)$ reacts to a Haar-like feature. The final Adaboost classifier is composed by all weak classifiers and it is represented as in Fig. 7.

As in Fig. 7, information about all subwindows extracted from the feature extraction module is goes through a cascade of 15 weak classifiers (C1...C15).

As mentioned before, at the end of the rejection cascade, an SVM classifier has been employed in order to reinforce the decision made previously and decrease the number of Adaboost false alarms.




Fig. 7. A cascade of rejection classifiers with an SVM classifier at the end

3.2 Support Vector Machine

SVM is a deterministic learning machine which employs linear discriminant functions into the raw input vector, in case of linear SVMs, or into a high dimensional feature space, in case of non-linear SVMs (Vapnik, 1995). As a supervised method for data classification, its structures embodies a training and prediction stages. In the training stage, algorithms provided by optimization theory are applied in order to learn how to separate the input space; when necessary, a mapping to a higher dimensional feature space is accomplished to guarantee the linear separability of data in any circumstance. In prediction stage, the classifier has the following forms in (2) and (3).

$$\begin{aligned} w_i x_i + b, d_i = +1 \\ w_i x_i + b, d_i = -1 \end{aligned} \quad (2)$$

Considering a training sample $\Omega = \{(x_i, d_i)\}$, where $i = 1 \dots N$ samples, x_i is the i^{th} input element and d_i is the i^{th} desired output, represented by the set $\{+1, -1\}$. Then, the discriminant function for a linear SVM is given by (2), where $+1$ represents an object and -1 , a non-object. In case of non-linear SVM, (2) is written in the form.

$$\begin{aligned} \phi(w_i x_i) + b, d_i = +1 \\ \phi(w_i x_i) + b, d_i = -1 \end{aligned} \quad (3)$$

where ϕ represents a kernel function which implicitly maps, by means of an inner product, the input space to a higher dimensional space. The most usual kernel function can be listed as:

1. Polynomial: $\phi(x_i, x_j) = (\gamma x_i^T x_j^T + r)^T, \gamma > 0$
2. Gaussian: $\phi(x_i, x_j) = \exp(-\gamma \|x_i - x_j\|^2), \gamma > 0$
3. Sigmoid: $\phi(x_i, x_j) = \tanh(\gamma x_i^T x_j^T + r), \gamma > 0$

A function is only considered a kernel if it satisfies the Mercer's condition (Cristianini & Shawe-Taylor, 2003; Kecman, 2001), *i. e.*, must be semidefinite and positive.

4. Sensor Fusion Using Fuzzy

In this section, the fusion processor, shown in Fig. 4, is described. This processor is based on a Takagi-Sugeno (TS) fuzzy engine (Takagi & Sugeno, 1985), responsible to decide which data from the sensors are to be taken into account in order to guarantee the most accurate object information to the robot. The TS fuzzy system has been built from the spatial sensor location on the soccer robot, and is depicted in Fig. 8. The main idea of the system is to provide spatial location of objects in front of the robot in polar coordinates (θ , d) and every location is regard to the mass center of the soccer robot. This information comes as from the IR distance sensor as from the camera (according to a calibration scheme, described in Section 5).

Hence, the decision of taking the angle or distance information from the IR distance sensor, from the camera, or combining both of them, must be made by the TS fuzzy system.

According to Fig. 8, the further the object is from the IR distance sensor (S1 to S5), the more accurate is the determination of the object angle. This is verified in the following way: the spread of each IR distance sensor is made by way of a 3 cm wide cylinder (distance between receptor and emitter); for instance, an object between A_1 and A_1' will have a less accurate angle determined by the distance sensors than an object between A_2 and A_2' , if only IR distance sensors are considered. In other words, the angle provided by the IR distance sensor will be more accurate in region A_2 and A_2' than in region A_1 and A_1' , since the first one is narrower.




Fig. 8. Determination of the fuzzy structure based on sensor location in the robot
Nevertheless, considering the precision scale, fuzzy sets are defined, according to Fig. 9, with the aim of determining the best angle information to the soccer robot




Fig. 9. Fuzzy sets determined from sensor physical disposal in Fig. 8. d is measured in cm

From these fuzzy sets, eight rules were proposed:

- R1: IF distance = SHADOW THEN distance = $d(\text{camera})$
- R2: IF distance = SHADOW THEN angle = $a(\text{camera})$
- R3: IF distance = VERYCLOSE THEN distance = $d(S_i)$
- R4: IF distance = VERYCLOSE THEN angle = $a(\text{camera})$
- R5: IF distance = CLOSE THEN distance = $d(S_i)$
- R6: IF distance = CLOSE THEN angle = $a(\text{camera}) * 0,5 + a(S_i) * 0,5$
- R7: IF distance = FAR THEN distance = $d(S_i)$
- R8: IF distance = FAR THEN angle = $a(S_i)$

The functions $a(\cdot)$ e $d(\cdot)$ represent, respectively, the angle and distance obtained by the camera and IR distance sensors. The real values of distances and angles, after evaluation of the rules, are determined by:

$$S = \frac{\sum \psi_i z_i}{\sum \psi_i}, \quad (4)$$

where ψ_i is the T-norm of each antecedent and z_i is the result of the function $f(x, y)$, responsible for describing the relationship between the fuzzy sets of the antecedent.

At the end of the fusion process, each object is identified and located by means of a internal representation frame structure (Minsky, 1975):

```
( frame <numFrame> ( see <timePerception> )
  ( object_name <idObject> )
  ( angle <angleObject> )
  ( dist <distanceObject> )
  ( vel <fuzzyVel> ) )
```

For each image frame $\langle numFrame \rangle$, all objects are located and identified by an object name $\langle idObject \rangle$ and three supplied characteristics: angle relative to the center of the base of the robot $\langle angleObject \rangle$, object distance regards to the front of the robot $\langle distanceObject \rangle$ and fuzzy velocity of each object $\langle fuzzyVel \rangle$.

The fuzzy velocity $\langle fuzzyVel \rangle$ is determined by (5). Fig. 10 shows the fuzzy sets used.

$$\langle fuzzyVel \rangle = [\mu_l(difP), \mu_m(difP), \mu_h(difP)], \quad (5)$$

where $\mu_i(difP)$ are membership functions, and i represents each fuzzy set of linguistic variable velocity (*low, medium and high*). $difP$ is the difference between centroid location of an object in relation of frames n and $n - 1$.




Fig. 10. Fuzzy sets for fuzzy velocity (fuzzified by pixel information). This information is just fuzzified and passed by the control module as an estimative

This velocity information is not intended to be accurate since the control module will translate it into a treatable knowledge within other fuzzy controllers. Hence, the universe of discourse represents the diagonal of the image frame given by the camera sensor. This information is just fuzzified and passed to the control module as an estimative.

5. Calibration Method

Based on the features extracted from the image (height, width, area, centroid), object location is determined in polar coordinates (θ, d) , where θ represents the angle and d is the distance relatives to the robot mass center.

To determine these polar coordinates, two regression SVMs have been employed: with respect to θ , a mapping function between the pixel coordinate of the centroid of the object and the concerned physical angle; with respect to d , a mapping between the height of the object in the image and the respective distance between the robot and the concerned object.

To obtain the mapping function centroid-angle, *radii* and parallel lines are drawn in a white paper, according to Fig. 11. Objects with a determined centroid pixel have been placed in each intersection in order to gather information to train the regression SVM.

Firstly, a set of pairs (centroid pixel, angle) is achieved from the centroid of each object in front of the robot and the relative angle to the mass center of the robot.




Fig. 11. Calibration map used to extract angles and distances from the image

The vision head is composed by two servo-motors, and, thus, it has two DOF: pan and tilt movements. In order to taper the angular distortion effect in the image, which occurs when the vision head takes different positions related to z axis, an angle value is added to the previous pair (centroid pixel, angle), according to Fig. 12.

In resume, in each angle (20, 35, 50 and 65 degrees), a set of pairs (centroid pixels) is determined and a tuple (centroid pixel, angle, angle of head) is given to the regression SVM.




Fig. 12. Angles of the servoing vision head used to enhance the regression SVM performance. The position (0, 0, 0) corresponds to the initial position in which the vision head starts when it is switched on

Information about distance is then extracted to complete the perception information when the objects are found in shadow areas. To obtain more accurate information of the distances, in case of a shadowed object, the robot may rotate its body to correct the estimated distance given by the image. This information is achieved by a simple relation between object height and distance.

5.1 Performance Analysis

Evaluation of the system has consisted to measure the difference between the true and the obtained values. To do this, the calibration map, in Fig. 11, was used once, but, now, with specific angles (-16.5, -8, 0, 8 and 16.5) and distance values (25, 30, 35, 40, 45, 50 and 55) and for each angle shown in Fig. 12. Results are, then, summarized in Table 1.

Table 1 shows the angles obtained by the regression SVM. The angles was gathered through four different angular positions of the vision head, and with objects in different places in the classification map shown in Fig. 11.

After data have been gathered to populate Table 1, analysis of the system error has been carried through and Table 2 summarizes the results. In Table 2, the absolute deviation of the measurements is placed regards to the mean values found in Table 1. The mean values of the absolute deviation were computed and they corresponds to the mean deviation obtained in each angular position of the vision head, as in Table 3.

It is worth noting a systematic positive error in the final correction values found, which indicates a higher error to the right side of the vision system. Hence and according to the final **mean error** of the system (see Table 3), the value of 1.88 degrees has been added to the angles obtained by the regression SVM. The final error could be explained by one or more of the items in the list below:

- **Radial distortion of the camera lens** - higher distortion of the camera lens to the right side;
- **Calibration error** - an error in building the calibration map;
- **Regression SVM** - lack of sufficient generalization or information to the regression SVM.

The values in the Table 3 can be better visualized with the respective values of the standard deviation in each point as shown in Fig. 13.

Vision head = 20 degrees (angles obtained by regression SVM)								
True angle	In 25 cm	In 30 cm	In 35 cm	In 40 cm	In 45 cm	In 50 cm	In 55 cm	Mean
-16.5°	-16.0°	-15.0°	-15.0°	-15.0°	-15.0°	-15.0°	-15.0°	-15.1°
-8°	-9.0°	-10.0°	-9.0°	-9.0°	-9.0°	-9.0°	-9.0°	-9.1°
0°	-1.0°	-1.0°	-1.0°	-1.0°	-1.0°	-1.0°	-1.0°	-1.0°
8°	5.0°	4.0°	4.0°	5.0°	5.0°	5.0°	5.0°	4.3°
16.5°	14.0°	14.0°	14.0°	14.0°	14.0°	14.0°	14.0°	14.0°
Vision head = 35 degrees (angles obtained by regression SVM)								
-16.5°	-15.0°	-15.0°	-16.0°	-16.0°	-16.0°	-16.0°	-16.0°	-15.9°
-8°	-10.0°	-10.0°	-10.0°	-10.0°	-10.0°	-10.0°	-10.0°	-10.0°
0°	-1.0°	-1.0°	-1.0°	-1.0°	-1.0°	-1.0°	-1.0°	-1.0°
8°	5.0°	4.0°	4.0°	6.0°	6.0°	6.0°	6.0°	5.6°
16.5°	14.0°	14.0°	14.0°	14.0°	14.0°	14.0°	14.0°	14.0°
Vision head = 50 degrees (angles obtained by regression SVM)								
-16.5°	-18.0°	-18.0°	-17.0°					-17.7°
-8°	-11.0°	-10.0°	-11.0°					-10.7°
0°	-2.0°	-2.0°	-2.0°					-2°
8°	4.0°	5.0°	7.0°					5.3°
16.5°	13.0°	12.0°	12.0°					12.3°
Vision head = 65 degrees (angles obtained by regression SVM)								
-16.5°	-17.0°	-17.0°						-17°
-8°	-8.0°	-9.0°						-8.5°
0°	-1.0°	-2.0°						-1.4°
8°	4.0°	4.0°						4°
16.5°	13.0°	12.0°						12.5°

Table 1. Results for each vision head angle assumed in the training phase (20, 35, 50 and 65)

Vision head = 20 degrees		
True angle	Mean	Absolute deviation
-16.5°	-15.1°	1.4°
-8°	-9.1°	1.1°
0°	-1.0°	1.0°
8°	4.3°	3.7°
16.5°	14°	2.5°
Vision head = 35 degrees		
-16.5°	-15.9°	0.6°
-8°	-10.0°	2.0°
0°	-1.0°	1.0°
8°	5.6°	2.4°
16.5°	14.0°	2.5°
Vision head = 50 degrees		
True angle	Mean	Absolute deviation
-16.5°	-17.7°	1.2°
-8°	-10.7°	2.7°
0°	-2.0°	2.0°

8°	5.3°	2.7°
16.5°	12.3°	4.2°
Vision head = 65 degrees		
-16.5°	-17.0°	0.5°
-8°	-8.5°	0.5°
0°	-1.4°	1.5°
8°	4.0°	4.0°
16.5°	12.5°	4.0°

Table 2. Absolute deviation between the mean of the angle values obtained by the regression SVM and the real angles (calibration map)

Vision head angle	Erro correction
20°	1.4°
35°	1.46°
50°	2.56°
65°	2.1°
Mean	1,88°

Table 3. Mean error correction of each vision head angular position




Fig. 13. Graph of the overall error correction

6. Experimental Results

A version of the proposed framework using an SVM multi-classifier to recognize the objects in the soccer field (same team robot, the other team robot and ball) can be found in (Oliveira

et al, 2005). To evaluate this system, Table 4 summarizes results over different illumination values, according to Robocup rules for F180 robot competition (Robocup, 2007).

Illumination (lux)	Classification Rate (%)
570	87.75
660	84.08
780	84.87
800	86.36
920	87.72
970	90.00

Table 4. Classification rate over different illumination values, measured by means of a luximeter

A luximeter was used to evaluate the illumination over the objects in the field. In order to decrease the illumination changing effect, the YCrCb colour space was used. After the object pixels have been classified, a cluster algorithm was applied in order to give the centroid of the object, considering: for the robots, the colours of the balls on the top of the robot, which discriminates the same team and the other team robots; and the ball.

Considering different kernels, Table 5 shows the overall performance classification rate over 800 luxs of illumination.

Kernel type	Classification Rate (%)
Linear	77.6
Gaussian	84.6
Sigmoid	78.5
3 deg. polynomial	64.2

Table 5. Use of different kernel types and its respective classification rate in 800 luxs

It is worth noting that Gaussian kernel has given the best classification rate and has motivated its use for classification and the results in Table 4.

Considering all the aforementioned, and motivated by the speed and good results in object recognition (Viola & Jones, 2001) by the boost classifiers, we have decided to apply a new classification approach, described in Section 3.1.

The performance of the system was again evaluated in the same conditions. Table 6 shows the last results in different illumination values (in *luxes*). As can be seen in Table 6, the classification performance rate increased with a respective increasing in speed (from 5 to 10 fps), since the Gaussian SVM is only applied in the final stage of the cascade rejection of the Adaboost (15 stages used. For more information, see Section 3.1).

Illumination (lux)	Classification Rate (%)
570	88.01
660	90.34
780	91.23
800	91.67
920	93.45
970	94.06

Table 4. Classification rate over different illumination values, measured by means of a luximeter

7. Conclusions

A framework for perception in robotics soccer has been presented. The proposed framework has shown been effective in recognize coloured objects in the soccer field but it might be slightly changed to be used in different approaches, through a new set of training samples and fuzzy rules.

A TS fuzzy engine has been used to integrate information from different sensors, particularly, an IR distance sensor and a camera. To integrate these sensor data, a novel calibration method, based on a regression SVM, was developed and it has shown a robust mapping between the calibration map and the obtained values in camera space, with a low average error of 1.88 degrees.

Also, the vision system was evaluated, and the new scheme, with an addition of a cascade of boost rejection and an SVM, has given better performance than in (Oliveira, 2005). The Adaboost classifier decreased the computation cost for the object recognition task and the SVM, used at the last stage of the cascade, reinforce the decisions taken by the Adaboost.

Future work has been conducted to a temporal image fusion by means of a tracking system, which will allow a robot to analyze of the behaviour of the robots and the ball. Moreover, this system will help to enhance the classification performance by decreasing yet more the number of false alarms.

8. References

- Bai, T. ; Gu, J. ; Cheng, G. ; Madjalawieh, O. & Liu, P. (2003) A New Landmark Framework for Mobile Robot Localization and Asynchronous Sensor Fusion. *International Symposium on Computational Intelligence in Robotics and Automation*, IEEE, pp. 1451–1456
- Costa, A & Bittencourt, G. (1999) From a concurrent architecture to a concurrent autonomous agents architecture. *Third International Workshop in RoboCup (IJCAI'99)*, pp. 85–90, Springer, Lecture Notes in Artificial Intelligence.
- Cristianini, N.; Shawe-Taylor, J. (2003) An Introduction to Support Vector Machine an Other Kernel-based Learning Algorithms, *Cambridge University Press*.
- Ferrein, A. ; Lakemeyer, G. & Hermanns, L. (2005) Comparing Sensor Fusion Techniques for Ball Position Estimation. *RoboCup*.

- Kecman, V. Learning and Soft Computing (2001), *The MIT Press*.
- Oliveira, L.; Costa, A. & Schnitman, L. (2005) An Architecture of Sensor Fusion for Spatial Location of Objects in Mobile Robotics. *Progress in Artificial Intelligence*, v. 1, pp. 462-473, Springer, Lecture Notes in Computer Science.
- Lanthier, M. ; Nussbaum D. ; Sheng A. (2004) Improving Vision-Based Maps By Using Sonar and Infrared Data, *10th Int. Conf. on Robotics and Applications, IASTED*, pp. 118-123.
- Minsky, M. (1975) A framework to represent knowledge. In: *The Psychology of Computer Vision*, McGraw-Hill, pp. 211-277.
- Oren, M.; Papageorgiou, C.; Sinha, P.; Osuna, E. & Poggio, T. (1997) Pedestrian Detection Using Wavelet Templates. *Conference on Computer Vision and Pattern Recognition*, IEEE Computer Society, pp. 193-199.
- Robocup (2007) [<http://small-size.informatik.uni-bremen.de/rules:main>].
- Takagi, T. & Sugeno, M. (1985) Fuzzy Identification of Systems and Its Applications to Modelling and Control. *Transactions on Systems, Man and Cybernetics*, v. 28, p. 15-33, IEEE.
- Vapnik, V. (1995) The Nature of Statistical Learning Theory. *Springer Verlag*.
- Viola, P. & Jones, M. (2001) Rapid Object Detection Using a Boosted Cascade of Simple Features. *Conference on Computer Vision and Pattern Recognition*, IEEE Computer Society, pp. 511-518.