Histogram-Based Visual Object Recognition for the 2007 Four-Legged RoboCup League*

Souzana Volioti and Michail G. Lagoudakis

Intelligent Systems Laboratory
Department of Electronic and Computer Engineering
Technical University of Crete
Chania 73100, Crete, Greece
{sbolioti,lagoudakis}@intelligence.tuc.gr

Abstract. RoboCup is an annual international robotic soccer competition. One of its most popular divisions is the Four-Legged League, whereby each team consists of four fully autonomous Sony AIBO robots and researchers focus solely on software development over a standard robotic platform. To compete successfully each robot needs to address a variety of problems: visual object recognition, legged locomotion, localization, and team coordination. This paper focuses on the problem of visual object recognition and suggests a uniform approach for recognizing the key objects in the RoboCup 2007 field: the two goals, the two beacons (colored landmarks), and the ball. The proposed method processes the color-segmented camera image and delivers the type, as well as an estimate of the distance and the angle with respect to the robot, of each recognized object in the current field of view. Our approach is based on a number of procedures which are used uniformly for all three types of objects: horizontal and vertical scanning of the image, identification of large colored areas through a finite state machine, clustering of colored areas through histograms, formation of a bounding box indicating possible presence of an object, and customized filtering for removing implausible indications. Our approach is compared against the approaches of two RoboCup teams (German Team 2004 and SPQR-Legged 2006) and is shown to be equally good or better in many cases. The proposed approach has been used successfully by Team Kouretes of the Technical University of Crete, Greece during various RoboCup events.

1 Introduction

In its short history, the RoboCup competition [1] has grown to a well-established annual event bringing together the best robotics researchers from all over the world. The initial conception by Hiroaki Kitano in 1993 led to the formation of the RoboCup Federation with a bold vision: "*By the year 2050, to develop a team of fully autonomous humanoid robots that can win against the human world*

* This project was partially supported by the European Marie-Curie International Reintegration Grant MCIRG-CT-2006-044980 awarded to Michail G. Lagoudakis.

Fig. 1. Four-legged league shot from RoboCup German Open 2007.

soccer champions". The uniqueness of RoboCup stems from the real-world challenge it poses, whereby the core problems of robotics (perception, cognition, action, coordination) must be addressed simultaneously under real-time constraints. The proposed solutions are tested on a common benchmark environment through soccer games in various leagues, thus setting the stage for demonstrating and promoting the best research approaches, and ultimately advancing the state-of-the-art in the area. Beyond soccer, RoboCup now includes also competitions in search-and-rescue missions (RoboRescue), homekeeping tasks (RoboCup@Home), robotic performances (RoboDance), and simplified soccer leagues for K-12 students (RoboCup Junior).

The Four-Legged League of the RoboCup competition (Figure 1) is among the most popular leagues, featuring four robot players (SONY AIBO robots) per team. The AIBO ERS-7 robot resembles a small pet dog and is equipped with a 576MHz 64-bit RISC processor, 64 Mb of RAM, a color CCD camera, three infrared distance sensors, touch sensors, wifi communication, and a number of servos on its legs, head, tail, and ears, providing 20 degrees of freedom in total. Games take place in a $4m \times 6m$ field marked with white lines, colored goals, and two colored beacons which serve as landmarks (Figure 2). Each game consists of two 10-minute halves and teams switch colors and side at halftime. There are several rules enforced by human referees during the game. For example, a player is punished with a 30-seconds removal from the field if he performs an illegal action, such as pushing an opponent for more than three seconds, grabbing the ball between his legs and head for more than three seconds, leaving the field, or entering his own goal area as a defender.

The main characteristic of the Four-Legged League is that no hardware changes are allowed; all teams use the exact same robotic hardware and differ only in terms of their software. This convention results to the league's characterization by a unique combination of features: autonomous vision-based player operation, legged locomotion and action, uniform robotic platform. Given that the underlying robotic hardware is common for all competing teams, research efforts have focused on developing more efficient algorithms and techniques for visual perception, active localization, omnidirectional motion, skill learning, and

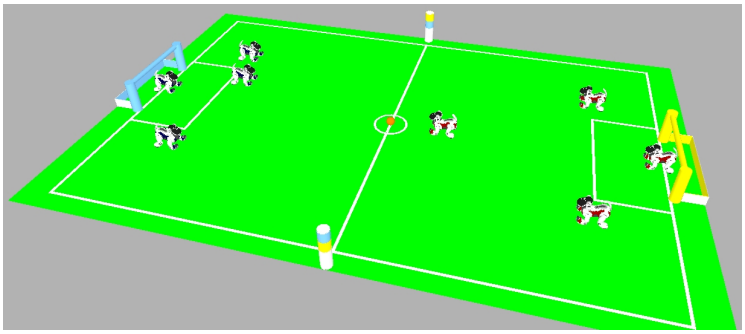

Fig. 2. The 2007 Four-Legged RoboCup League field [2].

coordination strategies. During the course of the years an independent observer could easily notice a clear progress in all research directions.

This paper focuses solely on the problem of object recognition within visual perception. Through its color CCD camera mounted on the head, each robot must be able to extract all information needed for successful play. The key objects in the RoboCup 2007 field are the two goals, the two beacons (colored landmarks), and the ball. Each object has its own appearance characteristics (color, shape, size) which discriminate it from the other objects. The goal is to identify such objects in the current field of view, if visible, and to estimate their distance and angle with respect to the robot's current pose. This information can subsequently be used for localization and decision making.

Briefly, the proposed object recognition method processes the color-segmented camera image; the image is scanned pixel by pixel, both horizontally and vertically along specific scan lines. Large areas of the same color are then determined through a finite state machine, which permits a specific tolerance of non-colored pixels within the colored areas. These areas are likely to be part of an object of interest. In order to figure this out, the colored areas are clustered through histograms, which expose the density of the colored areas and give an indication of their extent. Such indications help in finding out the dimensions of an object within the image. Then, after a series of transformations using the robot head pose, the object dimensions within the image, the camera parameters, and the real world dimensions of each object, an estimate of the distance and the orientation to each object with respect to the robot is computed. Some checks are also performed to filter out implausible indications. The proposed approach is compared against the solutions given by two RoboCup teams (German Team 2004 and SPQR-Legged 2006) by means of statistics in successful object recognition and processing time. The results indicate that our approach is equally good or even better in many cases.

The remainder of the paper is structured as follows: Section 2 describes the target field objects and states the problem. Section 3 describes the proposed histogram-based approach, while Section 4 discusses related work by other RoboCup teams. Finally, Section 5 summarizes the results of comparing our approach to other known approaches.

2 Problem Statement

This section states the problem we address in this paper and places our work in the broader context of the robot architecture in terms of input and output.

2.1 Raw Image Correction and Color Segmentation

Our object recognition method relies on visual input. Each AIBO robot is equipped with a color CCD camera which delivers YUV images at 30 frames per second with a resolution of 160×208 pixels. The viewing angle is 56.9 deg wide and 45.2 deg high. The camera is mounted on the robot head which can move independently of the robot body using its three degrees of freedom (head pan, head tilt, neck tilt). This freedom of motion offers an effective viewing angle which is approximately 240 deg wide and 175 deg high with respect to the robot body. Several hardware tests have shown that the raw image from the ERS-7 camera is radiometrically distorted (becoming blue away from the center and close to the corners), therefore we pass the image through an inverse transformation to cancel this distortion [3].

The corrected image is then segmented according to color. Given that there is only a handful of colors used for the objects in the field (green, orange, sky-blue, yellow, white), it is essential to map each pixel of the image to its correct color class and then use the color-segmented image for object recognition. Color segmentation has proved to be quite sensitive to lighting conditions, therefore RoboCup organizers arrange for ceiling lights that deliver approximately uniform lighting at about 1000 lux. Our color segmentation is based on a method that first transforms the image into the HSI format and then uses mainly the H-dimension to separate the colors, continuously adapting the separation thresholds every 20 frames for best separation [4]. This method is more effective when the raw image is saturated, therefore we typically set the camera parameters to high gain and medium shutter speed.

These preprocessing steps are essential for obtaining a color-segmented image, however they are not unique; other alternative choices could be used to produce the desired color-segmented image. A detailed coverage of image correction and color segmentation methods is outside the scope of this paper.

2.2 RoboCup 2007 Field Objects

The problem studied in this paper is the problem of object recognition using color-segmented images from the robot camera. The main objects of interest are the two goals, the two beacons, and the ball. The robot needs to recognize at least these objects in order to play soccer successfully; the ball is the central object around which the game revolves, the goals are the areas where the ball must go to or stay away from, and the beacons are essential landmarks in determining the current robot position and orientation in the field. These five objects are depicted in Figure 3. There are also other objects of secondary importance in the field; these are the lines marked on the fields and the partner/opponent robot

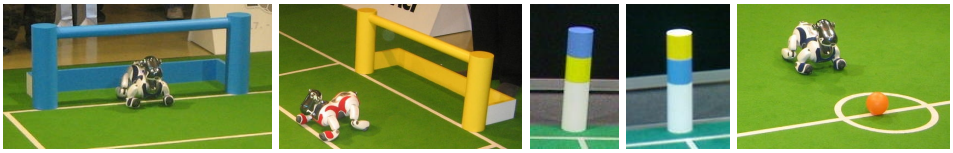

Fig. 3. Objects of interest in the 2007 field: goals, beacons, and ball.

players. Recognition of lines is not really helpful due to high perceptual aliasing (there are several areas in the field that yield the same line perception), and recognition of other players is a rather complicated task to be performed in real-time, given the multitude of possible robot poses and angles. The displayed goal structure was first introduced in 2007. In previous years, a completely different goal structure with compact colored areas was used, as well as four beacons instead of two marked with combinations of pink, yellow, and skyblue.

Each goal is either skyblue or yellow in color and is composed of four main pieces. Two cylindrical pieces of 30cm in length and 10cm in diameter stand for the vertical goalposts and support the upper horizontal goalpost which has a length of 80cm and a diameter of 5cm. There is also a rectangular part (10cm high) attached to the lower part of the vertical goalposts and delimits the goal region in the field. The beacons are cylindrical pieces 40cm in height and 10cm in diameter, which are painted white, skyblue, and yellow in two different combinations. Starting from the bottom, each beacon consists of a white part (20cm high) and then, depending on its type, two colored parts (10cm high). The beacons are placed at fixed positions on either side of the field middle line with the intention of facilitating the localization process, if recognized correctly. The ball is an orange plastic sphere 10cm in diameter. Its smooth surface typically leads to light reflections visible as yellow or white spots on the top part making its recognition even harder. Note that any of the above objects may be only partially visible through the camera. This is especially true for the ball when the robot is very close to it, a situation that occurs frequently. Even in such cases, it is imperative to obtain the correct type, distance, and angle of the object.

In order to address the recognition problem fully, the color-segmented camera images are combined with the pose of the robot at each moment, whereby by robot pose we mean the complete joint configuration of the robot. In particular, the pose of the robot's head, where the camera is placed, must be taken into account, so as to conclude the correct relationship between the camera image, the real world, and the orientation of the image with respect to the robot body.

3 Histogram-Based Object Recognition

This section covers the recognition process in detail. Recognition of the various objects relies on a few basic operations which are used in common. These basic operations are described first, followed by a description of the specialized recognition procedure for each object. Note that the parameter values given below are indicative only and may be changed at will. These numbers represent the choices that worked best for our specific recognition problem.

3.1 Basic Operations

The basic operations on the color-segmented image that are used uniformly for all objects are: construction of the horizon line, construction of scan lines, identification of large parts, construction of density histograms, and estimation of distance and angle. Note that these operations may be applied selectively on demand in each case; for example, vertical scanning may not be called, if horizontal scanning reveals no object indication.

Construction of the horizon line Since all objects in the field are aligned horizontally or vertically with respect to the horizon line, it is essential to construct first the horizon of the image. The horizon line is the intersection of the projection plane (the plane where the real world's objects are projected by the camera lens and is perpendicular to the optical axis of the robot camera) and a plane parallel to the ground at the height of the camera. To this end, the rotation of the camera relatively to the robot's body is necessary information; the rotation can be extracted accurately from the current robot pose and the forward kinematics of the robot.

Construction of scan lines Each image is scanned horizontally and if needed vertically to identify segments of uniform color. The horizontal scan lines are constructed to be parallel to the horizon line and to cover the entire image, whereas the vertical scan lines are constructed to be perpendicular to the horizon line and may or may not cover the entire image depending on the target object. Each scan line is eventually transformed into a pixelated line over the image pixels that approximates the true geometric line.

Identification of large parts Scanning along any scan line, whether horizontal or vertical, is conducted using a finite state machine. This machine alternates between two states depending on whether the target color has been found (TargetColor state) or not (NoTargetColor state) along the line, under some transition tolerance to account for noise in the image. Persistence of the machine in the TargetColor state for more than a threshold number of pixels indicates the existence of a large part within the scan line that carries the target color. Through the state machine, the beginning and end of each such large part can be easily marked.

Construction of density histograms All large parts (if found) are grouped into a single structure that reveals their density along the scanning direction. Each histogram is a discretization of the entire viewing angle along the scanning direction and the value of each bin indicates the number of overlapping large parts at that particular angle. The peak of the histogram indicates the highest density of large parts along the scanning direction.

Estimation of distance and angle If an object has been successfully identified on the image and its size in pixels has been determined, then using the known real dimensions of the object, the camera parameters and characteristics, and common geometry one can easily estimate the distance of the object from the camera, as well as its horizontal angle with respect to the camera focal

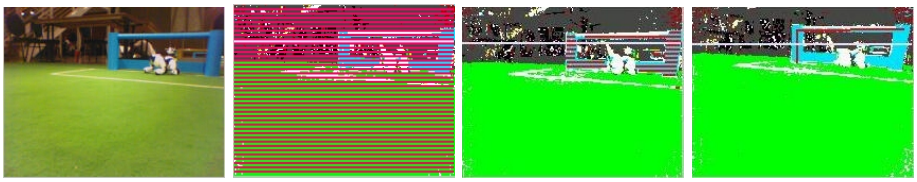

Fig. 4. Basic operations: raw image, horizon line, scan lines, large parts, indications.

point. Furthermore, utilizing knowledge of the current robot pose, the angle with respect to the robot body can also be estimated.

Figure 4 demonstrates some of these basic operations in the context of a horizontal scanning for a goal. Note that all operations are applied to the color-segmented image.

3.2 Goal Recognition

The recognition procedure for goals aims at recognizing the vertical and the horizontal goal posts separately. To this end, the image is first scanned horizontally in its entirety using 40 uniformly-distant horizontal scan lines and the large parts found are collected into the horizontal histogram. A peak in this histogram (if any) above a prespecified threshold (10% of the total number of scan lines) indicates the presence of a vertical post. If two peaks are found, then both vertical posts are visible and the entire goal is contained within the image (horizontally, at least); in this case, the horizontal goal indication extends between the two posts and receives high confidence on both sides. If only one peak is found then only one vertical post is visible and therefore the goal is only partially visible in the image; in this case, the horizontal goal indication extends from that post to the appropriate end of the image as indicated by the density of large parts and receives high confidence only on the side of the vertical post. If no peaks are found, then either no goal is currently visible in the image, or the robot is so close to the goal that only sees the central goal part in the image; in these two cases, either no horizontal goal indication is formed at all, or the one formed comes with a low confidence.

The procedure continues with a vertical scanning using 50 uniformly-distant vertical scan lines that extend only above and below the horizon to a fixed viewing angle ($\pm 30\,\mathrm{deg}$). This is dictated by the fact that the goals lie around the horizon line and cannot be far away from it. Note that if the current pose of the robot is such that the horizon line is not visible, vertical scanning may not take place at all, if no vertical scan lines lie within the image. Vertical scanning proceeds exactly as the horizontal one, accumulating large parts into the vertical histogram and identifying two, one, or no vertical goal posts as before. A vertical goal indication is similarly derived from the vertical scanning with the appropriate confidence level.

During the final goal recognition phase, it is decided whether the horizontal and vertical indications can be combined to offer a single goal indication (a bounding box). Three filters are used to reject implausible indications: (a) the

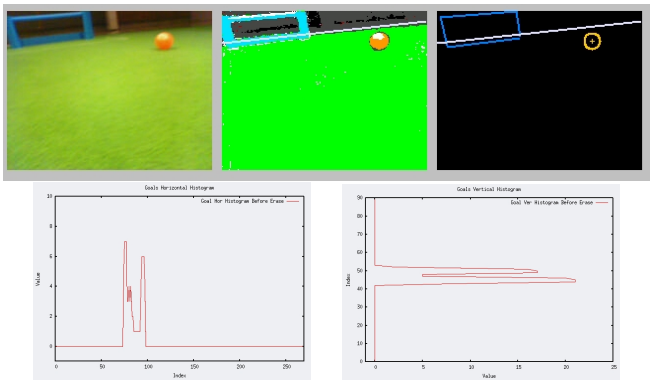

Fig. 5. Goal recognition: raw image, color-segmented image, recognized objects

confidence filter rejects indications that come with low horizontal and vertical confidence, that is, no goal post has been identified; (b) the ratio filter rejects indications that come with high confidence, but their horizontal vs. vertical size ratio fails to comply (within an interval) to the true size ratio of the real goal; and, (c) the distance filter rejects indications that yield an invalid distance estimation (negative or larger than the field size). If a set of indications passes successfully through these filters, the combined indication is used to determine the distance and the angle of the recognized object with respect to the robot body and the resulting information is passed to the higher level of cognition. The entire procedure is repeated in a interleaved manner for both the skyblue and the yellow colors. Figure 5 shows an instance of successful goal recognition along with the corresponding horizontal and vertical histograms. The two peaks in each histogram indicate the corresponding goal posts.

3.3 Beacon Recognition

The recognition procedure for beacons aims at recognizing areas where the two colors are aligned on top of each other. To this end, the image is first scanned horizontally in its entirety using 60 uniformly-distant horizontal scan lines and the large parts found are collected into horizontal histograms, one for each color (yellow and skyblue). Then, we seek to find a peak in the product of the two histograms; the idea is that only areas where the two colors coincide will yield large product values. Before multiplication, each entry in the two histograms is also checked for having a value above a prespecified threshold (10% of the total number of scan lines) and whether the values in the two histograms are not too far from each other (each beacon colored area must yield approximately the same number of large parts). If such a peak is found, it clearly indicates two aligned colored areas of approximately the same size, likely a beacon indication. The horizontal extent of the mode corresponding to the peak in the product of histograms is calculated to mark the horizontal limits of this indication.

If a horizontal indication is found, the procedure proceeds with a vertical scanning using only 7 uniformly-distant vertical scan lines that extend horizon-

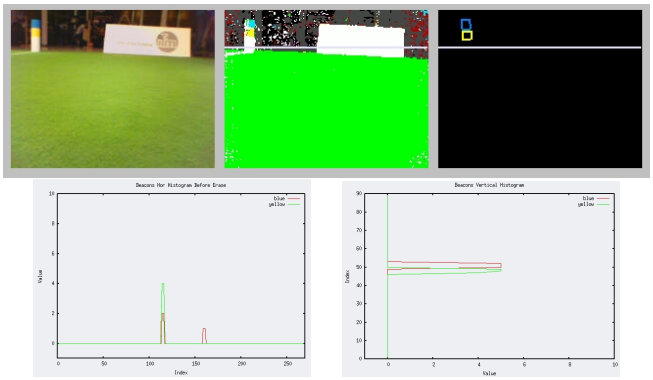

Fig. 6. Beacon recognition: raw image, color-segmented image, recognized objects

tally only within the limits of the horizontal indication and vertically only above and below the horizon to a fixed viewing angle (± 30 deg). Again, this is dictated by the fact that the beacons, like the goals, lie around the horizon line and cannot be far away from it. Note that if the current pose of the robot is such that the horizon line is not visible, vertical scanning may not take place at all, if no vertical scan lines lie within the image. Vertical scanning proceeds exactly as the horizontal one, accumulating large parts into two vertical histograms and identifying one peak in each, one for each color. The relative positioning of the picks indicates which color is above and which one is below.

During the final beacon recognition phase, it is decided whether the horizontal and vertical indications found can be combined to offer a single beacon indication. Three filters are used to reject implausible indications: (a) the proximity filter rejects indications where the vertical peaks of the two colors are far apart and not adjacent to each other; (b) the ratio filter rejects indications whose horizontal vs. vertical size ratio fails to comply (within an interval) to the true size ratio of the real beacon; and, (c) the distance filter rejects indications that yield an invalid distance estimation (negative or larger than the field size). If a set of indications passes successfully through these filters, the combined indication is used to determine the distance and the angle of the recognized beacon with respect to the robot body and the resulting information is passed to the higher level of cognition. Figure 6 shows an instance of successful beacon recognition along with the corresponding horizontal and vertical histograms. The two peaks in each histogram indicate the corresponding areas for the skyblue and the yellow colors. Notice that the two peaks coincide in the horizontal histogram, but are adjacent in the vertical one.

3.4 Ball Recognition

The recognition procedure for the ball is a bit trickier given that, when the robot is close to the ball (this is often the case), the ball cannot be viewed in its entirety in the image. The image is first scanned horizontally in its entirety using

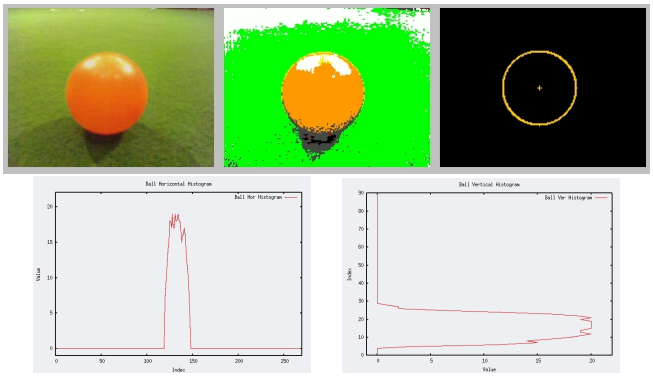

Fig. 7. Ball recognition: raw image, color-segmented image, recognized objects

40 uniformly-distant horizontal scan lines and the large parts of orange color found are collected into the horizontal histogram. Then, the image is scanned vertically using 50 uniformly-distant vertical scan lines that extend only below the horizon line (the ball cannot lie above the horizon line—this is also a way of avoiding false positives caused by spectators around the field wearing orange clothes) and the large parts of orange color found are collected into the vertical histogram. Then, we seek to find the peaks in both the horizontal and the vertical histograms. The resulting point (if found) indicates a central point of an orange blob in the image which might correspond to a ball.

In the next phase, a ray scanning begins from the central point along eight directions in the image (the four orthonormal and the four diagonal) aiming at identifying at least three points on the periphery of the ball. If a ray scan reaches the end of the image before the end of the orange color, it is rejected as it corresponds to a direction where the ball is hidden. However, if a transition to another color (with appropriate tolerance) is observed, then the last orange point is kept as belonging to the periphery of the ball. From three such points the center and the radius of the circle marking the ball can be calculated using the middle perpendiculars. In particular, the center of the circle is the intersection point of the middle perpendiculars of the circle chords formed by the three points on the periphery. The radius is calculated as the distance of the center to any of the three points. Note that the center of the ball may lie outside the image.

The resulting circle indication goes through two filters: (a) the horizon filter rejects indications which place the center of the ball above the horizon; and, (b) the distance filter rejects indications that yield an invalid distance estimation for the ball (negative or larger than the field size). If the circular indication passes successfully through these filters, it is used to determine the distance of the recognized ball and its angle with respect to the robot body; this information is then passed to the higher level of cognition. Figure 7 shows an instance of successful ball recognition along with the corresponding horizontal and vertical histograms. This is the simple case where the ball lies entirely within the image.

4 Related Work

This section reviews the work of four representative RoboCup Four-Legged teams in object recognition. Given that the type of goals used before 2007 were different, we only review beacon and ball recognition, which remained largely unchanged.

The German Team is a collaboration of the Humboldt University of Berlin, the University of Bremen, the Technical University of Darmstadt, and the University of Dortmund and has been participating in RoboCup competitions since 2001. In German team's object recognition [3], the horizon line is calculated first and the image is scanned using non-uniform horizontal and vertical scan lines. These scan lines are denser near the horizon and sparser away from it. For beacon detection (2004 type), the scan lines look for pink pixels, which are clustered to form a horizontal indication of pink color. From that indication, 4 scan lines are constructed perpendicular to it and scan for the other color (yellow or skyblue) of the beacon. After that, different sanity and comparison checks are performed to verify the existence of a beacon and, if they come true, the center of the pink area is calculated. From this center point, the image is scanned in four directions to detect the beacons' borders. Regarding ball detection, the scan lines search for the largest orange part of the image and, when found, its center is calculated. The center is the starting point of eight scan lines that scan horizontally, vertically, and diagonally in order to find edges of the ball. Some further color checks are performed to avoid false ball perceptions and, the ending points of the scan are saved in a list. The center and radius of the ball are found using the Levenberg-Marquardt method taking into account points from this list.

The rUNSWift team is based at the University of New South Wales in Sydney, Australia, and first participated in the RoboCup competition in 1999. Object recognition for rUNSWift [5] begins with feature detection. During this phase, the image is scanned using grids of scan lines constructed around the horizon line and each pixel's value is compared with its previous one in the scan line. If a rapid value transition occurs an object feature is detected. For example, changing from green to orange indicates a ball feature. The features detected are grouped together to create the complete object. Beacon recognition (2004 type) relies on detecting pink features. An eight-point asterisk scan is performed around each beacon feature to figure out the beacon type. Then, the height and centroid of the pink features are calculated and from them the height of the whole beacon. Ball recognition is based on detecting ball edge features to form the outline of the ball, according to a generalization of Siegel's repeated median line fitting algorithm to circles. These features are combined in all possible triplets in order to form a circle using the middle perpendiculars.

Cerberus is a four-legged league team based at Bogazici University in Istanbul, Turkey and has been participating in the RoboCup competition since 2001. Regarding object recognition, first the regions of interest that may contain an object are determined. The Cerberus team applies a region growing algorithm on the raw image [6]. Regions of the same color are merged, if they are adjacent to each other, and are blobbed to form areas likely to contain a specific object (candidate regions). For accurate processing, each candidate region is rotated

Table 1. Recognition performance of the proposed and other approaches.

Object	Approach	True Positives	False Negatives	True Negatives	False Positives
Goals	Proposed	**61%**	**39%**	**93%**	**7%**
Beacons	Proposed	**34%**	**66%**	**99%**	**1%**
	GT2004	28%	72%	99%	1%
	SPQRL2006	31%	69%	99%	1%
Ball	Proposed	**79%**	**21%**	**100%**	**0%**
	GT2004	67%	33%	99%	1%
	SPQRL2006	51%	49%	100%	0%

according to the robot's head pose. The team has introduced the bounding octagon of a region, which intends to reduce the rotation cost, as octagons are more easily rotated than rectangular boxes. To detect an object, a number of checks is performed on each candidate region. First, the neighboring regions should meet the specifications of the field colors. For instance, the bottom neighboring region of the ball should be green. Furthermore, the regions' size is taken into account and compared with the actual size and dimensions' ratio of the real objects.

The SPQR-Legged team is based at the University of Rome "La Sapienza" in Rome, Italy and has been participating to RoboCup competitions since 1998. Object recognition begins by scanning the whole color-segmented image trying to find regions of the same color. This process consists of two phases: first, the image is scanned for different particles of the same color and then the particles of the same color are blobbed in order to create a region where most pixels belong to a single color class. So, for recognition of particular objects, one needs to scan only the relevant colored regions. Each region is a rectangular box within the image frame. Once the blobbing of candidate regions is finished, each region is scanned for an object according to the color that the region contains. Concerning recognition of the beacons and the ball, SPQR-Legged has been using the approaches of the German Team applied separately to each candidate region and not to the whole image as the German Team does.

5 Results

This section attempts an evaluation of the proposed approach and a comparison to related approaches of two other RoboCup teams: GermanTeam 2004 and SPQR-Legged 2006. It was not possible to compare goal recognition, since the present goals have been introduced in 2007 and are significantly different than those used in previous RoboCup competitions. The statistics of object recognition using all approaches is shown in Table 1. Evaluation is based on three log files[1] recorded during the RoboCup German Open 2007 competition totalling several hundreds of frames.

GT2004 approach exhibits better performance than SPQRL2006 in recognizing the ball, whereas SPQRL2006 performs better beacon recognition. Our

[1] Log files contain continuous camera frames along with the robot pose at each frame.

proposal compares favorably to both of them; it improves on both beacon and ball recognition. The average processing time of our approach stands at 51 ms per frame, where GT2004 stands at 42 ms and SPQRL2006 at 52 ms. GT2004 achieves the lowest processing time, since the grids of scan lines focus only on regions that are most likely to contain objects of interest, whereas SPQRL2006 needs more time per frame to determine the candidate regions and process separately each one of them. Our approach exhibits a processing time similar to that of SPQRL2006, however with a significant increase in performance. The increase compared to GT2004 is due to the larger and uniform grids of scan lines we use.

False negatives are mostly due to the rapid movement of the robot's head, which can degrade the original appearance and shape of the object within the image. In addition, bad color segmentation, shadows, and reflections within the image can lead to misunderstanding of the object shape, which may lead to a false negative. Regarding the relatively low percentage of true positives for the beacons, this is likely due to their relatively small colored area. On the other hand, ball perception reaches excellent performance perhaps due to the fact that the ball is the only orange object in the field, in contrast to the skyblue and yellow colors which appear in multiple places.

6 Conclusion

The robots operate autonomously in the four-legged league, therefore all processes (object recognition, motion control, behavior control, localization, etc.) must run in real time on the onboard processor. This limitation prevents the use of sophisticated machine vision techniques for object recognition. In this paper, we proposed a simple approach for efficient recognition of the objects in the 2007 field that yields good performance while observing the real-time requirements. The proposed approach has been used successfully by Team Kouretes of the Technical University of Crete, Greece during various RoboCup events.

References

1. Kitano, H., Asada, M., Kuniyoshi, Y., Noda, I., Osawa, E., , Matsubara, H.: Robocup: A challenge problem for AI. AI Magazine **18**(1) (1997) 73–85
2. Four-Legged Technical Committee: Robocup four-legged league rule book (2007)
3. Rofer, T., Laue, T., Burkhard, H.D., Hoffmann, J., Jungel, M., Gohring, D., Lotzsch, M., Duffert, U., Spranger, M., Altmeyer, B., Goetzke, V., von Stryk, O., Brunn, R., Dassler, M., Kunz, M., Risler, M., Stelzer, M., Thomas, D., Uhrig, S., Schwiegelshohn, U., Dahm, I., Hebbel, M., Nistico, W., Schumann, C., Wachter, M.: GermanTeam 2004. Technical report, Humboldt-Universitat zu Berlin, Universitat Bremen, Technische Universitat Darmstadt, and University of Dortmund (2004)
4. Iocchi, L.: Robust color segmentation through adaptive color distribution transformation. In Lakemeyer, G., Sklar, E., Sorrenti, D.G., Takahashi, T., eds.: RoboCup. Volume 4434 of Lecture Notes in Computer Science., Springer (2006) 287–295
5. North, A.: Object recognition from subsampled image processing. Honor's thesis, University of New South Wales, Sydney, Australia (2005)
6. Akin, H.L., Mericli, C., Gokce, B., Geleri, F., Tasdemir, N., Celik, B., Celik, M.: Cerberus'06 team report. Technical report, Bogazici University, Turkey (2006)