

Vision Software for a Humanoid Soccer Robot

David Prasser

19th October 2001

43 Aberfoyle St,
Kenmore,
Brisbane QLD 4069

19th October, 2001.

Professor Simon Kaplan,
Head of School,
School of Information Technology
and Electrical Engineering,
University of Queensland,
St Lucia QLD 4072.

Dear Professor Kaplan,

In accordance with the requirements of the degree of Bachelor of Engineering (Honours) in the division of Electrical Engineering, I present the following thesis entitled "Vision Software for a Humanoid Soccer Robot". This thesis project was conducted under the supervision of Dr Gordon Wyeth.

I declare that the work submitted in this thesis is my own, except as acknowledged in the text, and has not been previously submitted for a degree at the University of Queensland or any other institution.

Yours sincerely,

David Prasser.

Acknowledgements

Gordon Wyeth for starting the whole GuRoo project.

Mark Chang for his endless advice and of course the SH4 board.

The other vision people: Andrew Blower, Shane Hosking, and Adrian Ratnapala for their ideas and advice.

Mark Venz and Damien Kee for providing the RoboLab Public Library.

Damien Kee, Ilan Lacish and Mark Venz again for their help with LYX and L^AT_EX.

The rest of the GuRoo team: Andrew Smith, Emanuel Zelniker, Andrew Blower, Damien Kee, Mark Wagstaff, Anthony Hunter, Shane Hosking, Nathaniel Brewer, Timothy Cartwright, Jarad Stirzaker, and Bartek Babel.

And everyone else in the robotics lab for making it such a relaxed and friendly lab.

Abstract

This thesis describes the initial work in developing a software system to provide visual information for a humanoid soccer playing robot. The robot is intended to operate in a colour coded environment. The system is to report the location of significant objects for playing soccer - the ball, goals, edgelines and opponents. In addition to this it is important for the robot to determine its own location and orientation. The software runs on a Hitachi SH4 processor board and is intended to use an OV7620 CMOS camera to provide YUV or $YC_B C_R$ images.

Several approaches were investigated before the final system was designed. The system has a colour detection first stage that uses a UV lookup table. This is followed by a run length based grouping algorithm that constructs objects from the colour detected image. Finally simple heuristics are used to reject poorly defined objects.

The code was trialled on an SH4 processor with a resolution of 240 by 180 pixels and operates at frame rate of 10 frames per second performing colour segmentation and object classification. The image size is limited by the amount of memory on the processor board. Further increases in speed would be possible by transferring parts of the code from the SH4 processor to a field programmable gate array that is also on the board.

The use of colour lookup tables combined with a row based object growing gives a fast method for robotic vision that will be suitable for humanoid robot soccer.

Contents

1	Introduction	1
1.1	Robotic Vision	1
1.2	RoboCup	1
1.3	GuRoo, the University of Queensland Humanoid Project	2
1.4	Robocup Vision	3
1.5	Outline of Thesis	4
2	Literature Review	5
2.1	Robotic Vision	5
2.1.1	Greyscale Thresholding	5
2.1.2	Colour Thresholding or Colour Detection	6
2.2	Morphology Theory	6
2.2.1	Edge Detection	8
2.2.2	Hough or Radon Transform	8
2.3	Robotic Vision in Robocup	9
2.3.1	YUV Colour Space	10
2.3.2	HSI or HSV Colour Space	10
2.3.3	After Thresholding	11
2.3.3.1	Localisation of Objects	11
2.3.3.2	Localisation of the Robots	11
2.4	CM Vision	11
2.5	ViperRoos	12

3	Problem Definition	14
3.1	Object Recognition and Self Localisation	14
3.2	Frame Rate and Resolution	15
3.3	Hardware	15
3.3.1	Camera	16
3.3.2	Processor Board	16
4	Initial Approaches	18
4.1	Colour Detection	18
4.1.1	Lookup Tables	18
4.2	Orientation	19
4.2.1	Edge Detection	19
4.2.2	Edge Following	19
4.3	Optical Flow	20
5	Final SH4 Software	22
5.1	Colour Detection	22
5.1.1	Basic Colour Detection	22
5.1.2	Fast Colour Detection	24
5.1.3	32 Bit Operation	24
5.2	Morphological Techniques Used in the Code	25
5.3	Grouping	27
5.3.1	Run Length Encoding	27
5.3.2	Blobs	28
5.3.3	Grouping	28
5.4	Analysis	31
5.5	Localisation and Line Detection	32
5.6	Training	32
5.7	Memory Organisation	33

6	Testing and Results	35
6.1	Testing Methods	35
6.2	Training and Lookup Table	35
6.3	Colour Detection	36
6.3.1	Reliability and Robustness of Colour Detection	36
6.3.2	Speed of Colour Detection	36
6.4	Morphological Operations	37
6.5	Run Length Encoding	37
6.6	Grouping	38
6.7	Blob Analysis and Object Detection	38
6.8	Line Detection	39
6.9	Total Processing Time	39
7	Future Work	41
7.1	Optimisations	41
7.1.1	Use of the Field Programmable Gate Array	41
7.1.2	Detection of Green	43
7.1.3	Faster Morphology	43
7.1.4	Noise Rejection	43
7.2	Additions	44
7.2.1	Coordinate Transforms and iPaq Software	44
7.2.2	Adapting to the output of the OV7620 Camera	45
8	Conclusion	46
A	Colour Images	49

B Code	57
B.1 Imdemo.c	57
B.2 Segment.h	57
B.3 Segment.c	59

List of Figures

1.1	The Guroo	2
2.1	Morphological operations.	7
2.2	Sobel Edge Detectors	8
2.3	Radon Transform	9
3.1	Hardware Block Diagram of the Humanoid.	16
4.1	Chain Code Line Description	21
5.1	Subsampling	24
5.2	32 Bit Data Storage	25
5.3	Run length encoded elements.	30
5.4	Memory Usage in the SH4.	34
6.1	Examples of the Line Detection Process	40
7.1	Logic Circuit Implementation of Colour Detection.	42
A.1	YUV Image Representation	50
A.2	Trainer Lookup Table	51
A.3	An example colour lookup table.	52
A.4	Colour Detected Image	53
A.5	Eroded Images	54

A.6 An Example of Poor Colour Detection	55
A.7 Detected Objects.	56

List of Tables

3.1	Colour Coded Environment	15
4.1	Eight neighbour direction codes.	20
5.1	Colour Codes	23
5.2	Dilation Priority	33
6.1	Subsampling speed and error	37
6.2	Breakdown of Processor Time	39

List of Algorithms

1	Basic Colour Detection Algorithm	23
2	Fast Colour Detection Algorithm	25
3	Simple 1×3 erosion.	26
4	Fast 1×3 erosion.	26
5	Grouping	28
6	Matching	29

Chapter 1

Introduction

1.1 Robotic Vision

One of the significant problems in robotics is extracting information from the environment. All but the most primitive industrial robots have some form of external sensing, such as sonar, infrared, or contact switches. The most versatile method is through computer vision, an algorithmic analysis of the digital output of a camera.

The problem then becomes one of constructing a model of the real world within the computer from the digital image received from the camera. A large amount of work has been done in this field of image processing. Approaches and algorithms for image processing can be roughly divided into two sections: real time and off-line image processing. Typically real time image analysis is done more crudely than offline analysis because of constraints on the amount of computational time that can be allocated for each image. Most robotics applications require real time processing, although applications that are not time critical may have a low frame rate.

1.2 RoboCup

The Robot World Cup or RoboCup is an international competition in which various types of robots compete against each other in games of soccer. The goal of RoboCup is to promote research in robotics while at the same time providing a series of complex problems for robotics and artificial intelligence development. The new focus of RoboCup is

humanoid robots with several humanoids being demonstrated at the 2000 RoboCup in Melbourne and competitions with humanoids planned for 2002[20]. The University of Queensland is building a robot, the GuRoo, to compete in future humanoid competitions [21]. Among the many systems that need developing for the GuRoo is a vision system to allow the robot to locate the ball, opponents, and other important data. The GuRoo robot will provide a platform for testing embedded vision systems.

1.3 GuRoo, the University of Queensland Humanoid Project

To represent the University of Queensland in the humanoid league at future RoboCups a robot, the GuRoo¹, is being designed. The GuRoo is intended to be approximately 1.2m high and be capable of autonomous movement, with on-board power and computer systems. Additionally the Guroo will have a fully functional upper body, including arms and a head that can tilt and pan (Fig 1.1).




Figure 1.1: The Guroo

The GuRoo is the a humanoid robot currently under development at the University of Queensland. [21].

The GuRoo's central computer will be a Compaq iPaq handheld computer connected to the drive control processors via a Controller Area Network (CAN) [12]. The iPaq will be responsible for the generation of walking patterns and gameplay intelligence. Also on

¹GuRoo, the Grossly Underfunded Roo. The suffix Roo is traditionally used for the University of Queensland's robot soccer teams.

board the GuRoo will be a Hitachi SH4 processor that analyses and interprets the output of a CMOS camera mounted in the robots head[5]. The vision processor passes on its output into the expansion slot of the iPaq. This SH4 processor will be used for testing the vision software developed in this thesis.

1.4 Robocup Vision

The RoboCup scenario will be used to provide some examples of the tasks an embedded vision system would need to be able to complete. The vision system is responsible for determining the location of the robot on the playing field as well as locating the ball, the goals and any opponents. It must be able to do this fast enough and accurately enough to give the robots the ability to determine the velocity and to intercept the ball. Latency is also an issue as the longer it takes to process a frame the less valid it becomes.

To ensure that the real time vision is possible the RoboCup competition enforces certain requirements that make colour segmentation possible. For example the edge lines of the field are white and the ball is orange. This makes colour detection a useful way of segmenting the objects of interest in the field.

There are two approaches to the way vision is used robotic soccer at the moment: global vision and local vision. Global vision uses a camera external to the robots which contains the entire playing field within its field of view, this technique is used only by the small size league robots. The alternative is a local vision system where the camera is mounted on the robot. This technique immediately leads to several complications:

- The camera is not in a fixed position or orientation, this means that the robot must be able to perform transformations to convert from image coordinates to world coordinates.
- The camera will not have a complete view of the area the robot is operating in. Typically RoboCup robots that use local vision have an inter-robot communication strategy to compensate for this problem.
- The shapes of objects change as they rotate making feature based object recognition more difficult. The ball being rotationally invariant is immune from this problem.
- The size of objects change which also makes object recognition more difficult.

The robots requirements for a high frame rate can also lead to significant trade offs in terms of both reliability and accuracy with local vision. In the context of humanoid robot soccer local vision is required.

1.5 Outline of Thesis

Introduction: An introduction to the robotic vision and its application to RoboCup. This chapter also discusses the the GuRoo and its planned vision system.

Literature Review: A review of current practice in image processing and robotic vision techniques, with emphasis on RoboCup robot.

Problem Specification: This chapter gives a detailed description of the requirements and goals of the GuRoo's vision system. This section will also describe the available resources for the software such as camera and processor types.

Initial Approaches: A number of early solutions to the problems outlined in chapter four are described, with particular emphasis on why they were never developed further.

Final SH4 Software: This chapter describes the process used to transform the camera output into a coded representation of each colour. The methods used to extract the location and type of objects from the thresholded image will be described.

Results: In this chapter the success of the software will be evaluated in terms of frame rate, accuracy, and robustness.

Conclusion: Final results of the project are summarised and future improvements are described.

Chapter 2

Literature Review

2.1 Robotic Vision

The initial step in most robotic vision applications is segmentation, in which the original image is segmented into different regions containing objects of interest[11]. After segmentation more complex algorithms can be used to determine the relevance and type of each object as well as its centroid and any other parameters of interest.

The very first stage of segmentation is usually to simplify the image by a process known as thresholding. Thresholding reduces the amount of information in an image by converting each pixel into some sort of symbolic code representing what type of pixel it is[11, 14, 10]. For example in text analysis a greyscale image could be converted into a binary image which is true for pixels that are part of the text and false for pixels that are part of the background.

2.1.1 Greyscale Thresholding

Thresholding is a simple technique for segmenting images. The most primitive example of thresholding is converting a grey scale image into a binary image composed of background and foreground objects. For instance all pixels above a particular level of brightness are foreground while those below are background. The difficulty lies in finding the threshold value to ensure that there are no errors [11]. One technique for selecting the correct thresholds is to use a histogram of the greyscale values [11, 14]. This system

can be expanded to segment the image into more than just background and foreground by using ranges instead of a single value.

2.1.2 Colour Thresholding or Colour Detection

Another method of segmenting an image is through the colour information received from a colour camera [2, 3, 14]. In effect this is like thresholding but with three values, for instance red, green, and blue. Histograms can again be used to determine the ranges for thresholding [6]. However the computational complexity has now been increased by a factor of three.

2.2 Morphology Theory

Mathematical morphology is a useful tool for preprocessing images [10]. The morphological operators treat binary images as sets which are members of a two dimensional coordinate space.

The simplest two operations are erosion (\ominus) and dilation (\oplus). These operations have two operands, an input (A) and a structuring element (B). Erosion results in groups of pixels becoming smaller while dilation causes groups of pixels to become larger [10].

Erosion is defined mathematically as equation 2.1, where $(B)_x$ is the element B translated by x . Effectively an erosion results in a set of points for which B fits entirely inside A . An example of bitmap undergoing erosion is shown in figure 2.1. Erosion can be used to remove small groups of noise pixels from a colour detected image [10].

$$A \ominus B = \{x \mid (B)_x \subseteq A\} \quad (2.1)$$


Dilation causes a cluster of pixels to grow in size according to the structuring element B . The set theory definition is shown in equation 2.2 where \hat{B} is the reflection of B around its origin. The result of a dilation is the set of points that B would occupy as it is moved around the image but still overlaps by at least one pixel with A . Dilations can be used to link together elements that have been incorrectly separated by the colour detection or thresholding processes [10].

$$A \oplus B = \left\{x \mid \left(\hat{B}\right)_x \cap A \neq \phi\right\} \quad (2.2)$$


The final two common morphological operations are opening and closing. An opening (\circ) is simply an erosion followed by a dilation with the same structuring element (equation 2.3). The erosion removes noise pixels and the dilation restores the the thinning caused by the erosion. A closing (\bullet) is effectively the opposite of an opening being a dilation followed by an erosion (equation 2.4). The effect of a closing is to link together pixels that are separated by a small distance by dilation. The erosion then restores the objects to their original size, keeping the new linked regions intact[10].

$$A \circ B = (A \ominus B) \oplus B \quad (2.3)$$

$$A \bullet B = (A \oplus B) \ominus B \quad (2.4)$$



(a) Dilation



(b) Erosion

Figure 2.1: Morphological operations.

The input is on the left and the output is on the right. The structuring element is in the middle.

2.2.1 Edge Detection

One technique commonly used for determining the shape of objects is edge detection. A sudden change in the image gradient, ie when one pixel is significantly darker than some of its neighbours, corresponds to the edge between two objects or one object and the background [19]. Typically an approximation to the gradient operator such as Sobel edge detectors are used[11, 19]. The Sobel operators shown in figure 2.2 provide an indication of edge strength when convolved with a greyscale image[10, 14]. Finally the gradient is thresholded to eliminate small variations that are caused by noise[19]. The result is a bit plane containing ones that correspond to sharp changes in brightness in an image.

-1	0	1	-1	-2	-1
-2	0	2	0	0	0
-1	0	1	1	2	1

(a) x gradient

(b) y gradient

Figure 2.2: Sobel Edge Detectors

The Sobel edge detectors calculate the change in brightness when they are convolved with a greyscale image. The X gradient operator (G_x) calculates the change in brightness in the X axis and Y gradient operator (G_y) calculates the change in the Y direction. The sum of the two operators will provide the edge strength at the point the operators are applied [10, 14].

2.2.2 Hough or Radon Transform

The Hough or Radon transform is a powerful method for converting an image of edge pixels into a set of lines. A line can be represented by the equation $\rho = x \cos \theta + y \sin \theta$ [14]. A two dimensional array of accumulator cells is created with one dimension being the line parameter ρ and the other θ . For each pixel in the edge map a set of parameters can be developed that represent all possible lines passing through that pixel. The corresponding accumulator cells for these parameters are increased by one. After each pixel has been processed in this way the cells with large values correspond to dominant lines in the edge map image[14].

An alternative approach with exactly the same result is the Radon Transform. The radon transform uses a rotating vector in the centre of the image (figure 2.3). At each angle of rotation the edge detected pixels are projected onto the vector. If the intensity of the projection is plotted against the distance along the vector and the rotation of the vector then the same output as the Hough transform is obtained.




Figure 2.3: Radon Transform

As the vector rotates around the origin the summated projections of each edge pixel onto the vector are recorded. When the vector is perpendicular to the direction of a group of edge lines the output is a maximum for that group of edges. For example the group of pixels labelled b in the above diagram would be giving its maximal progression at this rotation. The maximum would be at distance d along the vector.

2.3 Robotic Vision in Robocup

Real time image processing for robotics is a particularly difficult area of work. The vision problems expected in Humanoid local vision are very similar to the problems encountered

in the local vision version of the small and medium size leagues. The fundamental system used for segmentation in these leagues is colour[2, 3, 4].

Typically in RoboCup competitions the colours of all significant objects in the competition are standardised and known before hand, for example the ball is orange and the goals are blue and yellow. In global vision systems the robots usually make use of coloured markings to distinguish between robots. In localised vision however the robots are restricted to being black in colour[16, 17]. The Medium Size league uses colour tags on the robots to provide a mechanism for robots to distinguish between each other[16].

2.3.1 YUV Colour Space

Many systems use the YUV colour space to perform colour segmentation[2]. The YUV is a different interpretation of colour to the more familiar RGB colour space, where a pixel is composed of varying strengths of red, green and blue light. In the YUV system Y represents the intensity of the light, while U and V are the blue and red chrominance respectively [14] A representation of YUV data is shown in figure A.1. The conversion from RGB to YUV is ideally performed in the camera.

The advantage of the YUV colour space is that changes in brightness theoretically only affect the Y value. This means that colour detection can be performed using only the U and V values. The U and V coordinates are almost independent of lighting changes.

2.3.2 HSI or HSV Colour Space

Some robots use a different colour space known as HSV or HSI (hue, saturation, brightness). This colour space offers better separation between different colours but at the cost of slower processing as the camera output must be transformed into the HSV colour space[3, 4]. In HSV colour format the hue parameter describes the colour of the pixel while the saturation value represents how white this hue is. Clearly the brightness describes how bright the pixel is. The advantage of HSI is that is fairly close model of a humans perception of light[14, 10]. A simple way to speed up the transformation is to only calculate the pixel's hue. The hue will describe the colour of the object and may be enough to distinguish between all colours except black and white (which can be easily done in RGB or YUV space) [4].

2.3.3 After Thresholding

After thresholding the locations and type of all significant objects must be extracted from the picture. These objects include the ball, opponents, walls, goals and in some cases special markers for localisation (used in the Sony Legged league). Approaches to this are varied but are usually simple and optimised for speed.

2.3.3.1 Localisation of Objects

A simplification that can be made to the process of transforming two dimensional data to three dimensional world coordinates is to realise that in robot soccer the ball and other objects can be almost guaranteed to be on the ground. This makes the transform solvable with monocular vision without having to use information relating to the size of the object [2]. The AgiloRoboCuppers team use this to help discard redundant information, by calculating the size of each object and comparing it to the expected size based on its distance calculated by trigonometry if there are significant discrepancies then the object is ignored[2]. The trigonometry is encoded as a lookup table to reduce the processor load.

2.3.3.2 Localisation of the Robots

In the Sony Legged league there are colour marker posts that allow the robots to determine their position by triangulation [20]. In the other leagues however there are no such conveniences, and odometry information is not sufficient to retain a correct idea of the robots place in the world. The usual approach is to use the edge lines of the field to update the robots estimate of its location. The edge lines can be acquired by looking at transitions from white to green codes in the colour detected image[2] or by using a linear edge operator[3]. Linear regression or a Hough transform can then be used to find the lines inclination, from this the robots position can be determined[4, 2, 3].

2.4 CM Vision

A good example of a fast colour based system for image analysis is the CM Vision software library developed at Carnegie Mellon University[8, 7]. It internally uses either a YUV or HSI colour space and transforms RGB inputs to YUV. Colour thresholding is

performed using a colour cube in YUV space. A colour cube is box-like region in the three dimensional colour space. Pixels within this region can be tagged as being one type of coloured object while pixels outside this region are not.

Rather than use a series of comparison operations to determine whether or not a pixel is within the colour cube CM Vision uses a lookup table. Instead of a three dimensional $256 \times 256 \times 256$ table CM Vision uses three vectors one for each axis. Each colour coordinate is looked up in its respective vector and the result is ANDed together. This technique reduces a 16 MB lookup table to a 768 byte set of three tables. The cost of this system though is that it requires three memory lookups instead of one and can only represent rectangular regions in the colour space[8, 7].

After colour detection the system converts the data to a run length encoded data set. In run length encoding horizontal runs of pixels of the same colour are recorded simply as the number and colour of the pixels. From this the software performs four connected region growing. Four connected or four neighbour refers to the four pixels directly above, below and to the sides of the pixel in question. Initially each run length element is tagged as a the start of a blob and the software runs through each row comparing it with adjacent rows. Any run that is adjacent under four connectivity to the current one has its pointer updated. Two runs through the image are required to group all the runs into distinct regions. The final stage involves grouping similar regions based on colour, proximity and density, that links any regions that have been accidentally separated[8, 7].

2.5 ViperRoos

The University of Queensland's local vision team, The Viperroos faces similar problems to the GuRoo's vision system. The system used last year was based on YUV transformed data from an RGB camera[9].

The Viperroos use a two dimensional lookup table in UV space to distinguish between all the colours with the exception of black and white. The Y value is thresholded into three sections: black, colour, and white. Pixels that fall into the colour region are then thresholded via UV lookup.

The Viperroos' software then forms regions from the colour detected image and calculates the regions' centres. To avoid the multiplications and divisions involved in calculating a mathematical centroid the Viperroos define the centre to be the geometric centre of the blob.

The Viperroos currently use an SH3 processor running at 104 MIPS which produces around 12 frames per second. This year the camera will be upgraded to one with YUV output and the processor will be replaced by an SH4 with 360 MIPS. These changes should move the Viperroos into the 25-35 frame rate range[9].

This information is of particular interest as the GuRoo will be using the Viperroos new processor board and cameras.

Chapter 3

Problem Definition

The vision software should be useful to both the ViperRoos and the GuRoo as both face similar problems at the early stages of the software, ie fast colour segmentation. Additionally they will be using the same hardware which makes using the same software much more attractive.

The GuRoo's only external sensor is the camera mounted in its head. From this sensor all of the information needed must be derived. At present requirements for accuracy are unknown as the robot has not been constructed or the environment in which it is to operate specified. The medium and small size leagues are used as starting points for developing specifications for the vision system.

3.1 Object Recognition and Self Localisation

The location of the ball, obstacles, edge lines and goals needs to be determined relative to the robot. From this information the robot can work out its own location in the environment as well. No specifications for accuracy exist because the nature of the environment in which the robot operates is undefined.

Using the rules for other local vision leagues for RoboCup leads to the conclusion that the environment will be extensively colour tagged with significant objects having distinct and defined colours (Table 3.1)[17, 16]. Obviously the software must be capable of colour recognition to meet these requirements.

A capability originally suggested for the humanoid vision system was to be able to provide information suitable for the control systems of the robot to use as part of the stability

system. This would require calculating velocities and positions of the head with respect to the environment.

Object	Colour
Field/Operating Surface	Green
Ball	Red
Edge of Field Lines	White
Opponents	Black
Goals	Blue or Yellow

Table 3.1: Colour Coded Environment

In the colour coded environment the robot operates in each object of interest has an assigned colour.

3.2 Frame Rate and Resolution

The University of Queensland's other local vision robots, the ViperRoos face similar problems and uses the same hardware as the GuRoo. The ViperRoos' code is currently running on an SH3 processor with a RGB camera and has a frame rate of about 16 frames per second (fps) or 62.5 ms with a resolution of 32 by 128 pixels[9]. About 20 ms of time per frame is used to perform a costly RGB to YUV colour space transform. In terms of pixels per second the SH3 operates at around 64 kilopixels per second. The costly YUV conversion can be eliminated by using a YUV output camera reducing the processing time for a frame by 20ms. This would increase the number of pixels analysed by the system to 96.4 kilopixels per second. The ViperRoos' new SH4 processor will run at about three times as fast as the old processor giving a speed of approximately 300 kilopixels per second.

Any software system should at least run through the pixels at the same rate as the old ViperRoos software if not faster. For general three dimensional vision instead of using the ViperRoos letter box type image (32×128) an image with the usual $1\frac{1}{3} : 1$ aspect ratio of a computer screen should be used instead (640×480).

3.3 Hardware

The hardware consists of three physically separate elements the camera board, processor board and the iPaq (Fig 3.1). Communication between the boards is run through a field

programmable gate array (FPGA). Data flow is one way although there is the possibility of using the iPaq to alter the program settings on the SH4[12].




Figure 3.1: Hardware Block Diagram of the Humanoid.

This block diagram shows the flow of information through the humanoid's computer system. Data flow to and from the SH4 is through the FPGA allowing reconfigurable connections between input/output pins [5, 12].

3.3.1 Camera

The camera that will be providing images for processing is an Omnivision OV7620 CMOS camera[5]. There are several reasons for using this particular camera. Firstly it is a CMOS camera which means that it has a digital output and low current consumption both of which are necessary for mobile robots. Secondly the OV7620 has on board subsampling. The previous ViperRoos system performed subsampling inside the SH3 processor, in other words a full resolution picture is transmitted to the processor which discards the majority of the data. Subsampling by the camera reduces the amount of unnecessary data flowing across the bus to the processor. Finally and most importantly it can produce a YUV output[13, 5]. This removes the costly RGB to YUV conversion stage from the software.

3.3.2 Processor Board

The processor board is based around a Hitachi SH4 microprocessor. The SH4 is capable of running at up to 360 million instructions per second (MIPS) as opposed to the SH3 board it replaces which ran at 104 MIPS[9]. Aside from this approximately threefold increase in speed the SH4 processor board also has 512 KB of static ram (SRAM) as well as static-dynamic RAM (SDRAM). Also this board is equipped with a Spartan field programmable gate array (FPGA) which sits between the SH4 and the camera. The FPGA gives the option of carrying out some processing on the data before it reaches the SH4. It

is also planned to use the FPGA to buffer data from the camera as it comes into the SH4 [5].

Chapter 4

Initial Approaches

There were several abortive attempts at solving the problems outlined in chapter 3 before the techniques finally used were developed. All of these techniques were based on YUV images which were trialled in MATLAB.

4.1 Colour Detection

The original segmentation system used separate histograms of the red and blue chrominance. From these histograms the upper $x\%$ of the red chrominance and the lowest $x\%$ of the blue chrominance were used as the thresholds for the red ball. The parameter x must be adjusted to produce reasonable results. The difficulty with this technique is that the image must be histogrammed and the thresholds recalculated for each image. Secondly while this technique is good at finding red it is not so effective at locating pixels belonging to the ball as it does not combine information from the U and V coordinates. A faster solution is to use lookup tables.

4.1.1 Lookup Tables

A system similar to the CM Vision lookup table was trialled. The CM Vision system uses three vectors to describe rectangular three dimensional regions in YUV space. The problem with this technique is that adding adaptive Y thresholds is difficult and that the rectangular regions in UV space appeared to be less useful than arbitrary regions.

4.2 Orientation

To locate the robot in the in space it was originally planned to make detailed use of the edge lines of the field. By using edge detection techniques an image showing the edge lines can be produced. From this several methods could be used to calculate the line descriptors of the edges. One approach would be to use the Radon or Hough transform (2.2.2), but it was feared this may be too computationally expensive. So a different approach using edge following was trialled.

4.2.1 Edge Detection

The use of traditional edge detection operators such as Sobel or even Canny[14] detectors was avoided for several reasons. Firstly the the operators would require a large number of calculations to performed per pixel (at least four for a 2×2 operator). Secondly the results of these edge detections would be a set of edges for all significant changes in brightness while only the edges of the field lines need to be segmented. Instead of using a mathematical edge operator some other techniques were tried:

- Eroding the segmented white objects in the image and then XORing this with the original segmentation. This produces the set of pixels that are the boundary of all white objects. The problem with implementing this approach is that the entire boundary of the white objects are detected, while only the lower part is needed.
- Creating an edge map of all the white colour coded pixels that are directly above a green coded pixel. This works quite well and would be the best method for determining the edge lines if the field is segmented. In the final approach though the green area of the image is not segmented (Chapter 6).

4.2.2 Edge Following

The edge following system starts from a random point around on the edge map and begins tracing the edge. As the algorithm moves from pixel to pixel the direction is encoded and stored in a null terminated list (a chain code). As the algorithm moves between pixels in an eight neighbour manner there are eight possible directions to be encoded 4.1. After traversing a set of edge pixels in a straight line the start and end coordinates of the line will have been found.

Direction	Code
→	1
↗	2
↑	3
↖	4
←	5
↙	6
↓	7
↘	8

Table 4.1: Eight neighbour direction codes.

The edge following system represents the direction from pixel to pixel using these numerical codes.

As the algorithm moves along the system it uses simple heuristics to determine whether or not the line it is following is a straight line. While tracing a line the algorithm permits only two adjacent directions to be followed. This detects if the line doubles back or significantly changes direction but does not detect a line that only changes direction by 45° or less (figure 4.1). A solution to this problem is to develop an approximation to the gradient and use probability techniques to detect gradient changes[15]. This area of research was abandoned when it was found that the colour detection was not able to provide a straight edge line for a pixel by pixel approach like this.

4.3 Optical Flow

Optical flow is a technique for measuring the movement of pixels from one frame to the next[11]. Optical flow could provide some indication as to the velocity and stability of the head. However it is computationally expensive and was dropped at an early stage when gyroscopes were announced for internal sensors of the GuRoo. Because of the planned internal sensors and concerns with speed and latency within the vision system the use of optical information to assist with the control or balance of the GuRoo was abandoned.




Figure 4.1: Chain Code Line Description

The line following correctly detects the the straight line segment of the upper group of pixels but does not detect to the change in gradient of the lower line. The start and end points of the groups represented by the arrows show what the computer believes is a completely straight line.

Chapter 5

Final SH4 Software

In this chapter the final software running on the SH4 processor is described. It consists of several stages that execute one after another: colour detection; morphological operations; run length encoding; grouping; and object analysis. The software was prototyped in MATLAB and then recoded in C for a PC platform. Once the code was debugged and optimised it was ported to the SH4 using the GNU C compiler for the SH series[1].

5.1 Colour Detection

In the initial colour detection stage the YUV colour image is converted into a new image that contains eight bit codes representing what type of colour class each pixel belongs to. The colour codes used for humanoid vision are listed in table 5.1. The colour codes are set as powers of two so that each class of objects has its own bit plane. In other words the output of the colour detection can be thought of as eight single bit planes with the same width and height as the original image. As a result the system is able to detect eight different colours.

5.1.1 Basic Colour Detection

The final technique is based on a two stage system that first uses the pixel brightness to classify the pixel into one of three classes: black; white; or other. This is done with two simple thresholds, pixels below one threshold are black while pixels above the other threshold are white. The other pixels between the two thresholds are candidates to be

some other colour (yellow, green, blue) and require another level of attention to determine which. Colours apart from black and white are determined by a lookup table with U and V coordinates.

The Y thresholds are recalculated for each frame as offsets from the sample mean of the image. The lookup table values are constructed from a trainer program (section 5.6). The result of this process is that each YUV pixel is mapped to a single value describing the colour class of the pixel. The simplified implementation of this algorithm is shown as algorithm 1.

Algorithm 1 Basic Colour Detection Algorithm

```

mean ← mean of y values in image
low threshold ← mean - low offset
high threshold ← mean + high offset
for each pixel
  read Y
  if Y < low threshold then
    output ← BLACK
  else
    if Y > high threshold then
      output ← WHITE
    else
      read U
      read V
      output ← lookup table[U][V]

```

Object	Colour	Code
?	Any other	0
Field	Green	1
Wall/Edgeline	White	2
Ball	Red	4
Opponent/Obstacle	Black	8
Blue Goal	Blue	16
Yellow Goal	Yellow	32

Table 5.1: Colour Codes

After colour thresholding the output image contains these coded values describing each pixels colour. There are two remaining colour codes (64 and 128) that could be used to represent more colours.

5.1.2 Fast Colour Detection

The colour detection process can be accelerated by reducing the number of pixels that the operation is performed on. Instead of processing every pixel only every fourth pixel is used. If there is a change in output between last two pixels tested then the algorithm backtracks and processes the pixels in between. Otherwise it is assumed that these are all of the same value (algorithm 2). This increases the speed of the algorithm but also increases the number of pixels that are misclassified (Table 6.1).

The misclassification is caused by the fact that this algorithm can skip over small groups of pixels or small gaps between pixels (figure 5.1). In fact the algorithm cannot guarantee the segmentation of any object smaller than the subsampling size (4 pixels). However because of the backtracking the actual location and general shape of a segmented object is not damaged.




Figure 5.1: Subsampling

The first line shows the output of pixel by pixel colour detection. The second line shows the effect of $4\times$ subsampling with backtracking. The 0th and fourth pixels are examined first. The output is different so the 1st, 2nd and 3rd pixels are also tested. The final (8th) pixel is read, because it is the same as the previous subsampled pixel (4th) so the intervening pixels are kept as 1 without ever being thresholded.

5.1.3 32 Bit Operation

As the SH4 is a 32 bit processor reading 24 bit YUV colour information is not particularly efficient because of problems with word alignment and the fact that multiple reads are required. As can be seen in algorithm 1 either one or three memory reads are made per pixel. A much more efficient implementation can be made by reading the pixel as 32 bit value (figure 5.2). Each pixel then needs only one memory read. Furthermore the Y value thresholds can just be scaled and compared to the whole 32 bit pixel as the Y value is in the most significant byte (ignoring the unused byte). Also instead of using U and V as array indices (table[U][V]) which leads to the computer calculating the address as

Algorithm 2 Fast Colour Detection Algorithm

```

end_address ← image_size - CD_SUBSAMPLE
index = 0
while index < end_addr
  next_index ← index + CD_SUBSAMPLE
  output ← colour_detect(next_index)
  if output = last_output
    each output between index and next_index ← output
    index ← next_pixel
  else
    colour_detect each pixel between index and next_index
    index ← next_index

```

$table + U \times 256 + V$ the low 16 bits of the pixel data can be used as a one dimensional index. This is the only stage where the 32 bit architecture required change from the original code.




Figure 5.2: 32 Bit Data Storage

The YUV colour information is padded to 32 bits from 24. The magnitude of the 32 bit value approximately 2^{16} times the Y value of pixel. The lower 16 bits are used as the index into the lookup table.

5.2 Morphological Techniques Used in the Code

Morphological operations can be used to decrease the amount of noise in the image after the colour detection phase. To remove the noise quickly only an erosion is used as opposed to an opening. A 1×3 structuring element is used to clean up the colour detected data. The 1×3 element was chosen because it was the minimum size necessary to give good noise rejection.

A simple implementation of 1×3 erosion is shown in algorithm 3. This algorithm makes use of the fact that each colour code is in a separate bit position so an erosion can be

thought of as happening on eight bitmaps in parallel. This implies the AND in algorithm 3 is a bitwise AND instead of a logical AND.

Algorithm 3 Simple 1×3 erosion.

```

for i = 1 to size of image -1
  output[i] = input[i-1] AND input[i] AND input[i+1]

```

The morphological system can be sped up by reducing the number of memory accesses. It is clear from algorithm 3 that each memory location is accessed three times, so by keeping previously read data it is possible to reduce the memory reads by three. The C code for this erosion is shown as algorithm 4.

Algorithm 4 Fast 1×3 erosion.

```

i ← 0
for y ← 0 to height
  out[i] ← 0
  i++
  a ← input[i-1]
  b ← input[i]
  c ← input[i+1]
  for x ← 1 to width - 4 step x+=3
    out[i] ← a AND b AND c
    i++
    a ← input[i+1]
    out[i] ← b AND c AND a
    i++
    b = input[i+1]
    output ← c AND a AND b
    i++
    c ← input[i+1]
  next x
  input[i] ← 0
  i++
  input[i] ← 0
  i++
next y

```

5.3 Grouping

The grouping algorithm converts the colour detected pixels into a list of structures describing each four connected group of pixels. There are two passes in this algorithm, the first pass turns the colour detected image into a list of run length encoded (RLE) elements. These elements are then used to construct a set of blob objects that describe large blobs found in the image. The advantage of using RLE is that the grouping algorithm will be simpler if it operates on blocks of pixels instead of individual pixels. The other advantage is that the RLE compressed data will be convenient for transmitting to the iPaq, so a colour detected image can be displayed on the iPaq's screen[12].

5.3.1 Run Length Encoding

The first stage of the grouping algorithm is to convert from an array of pixels to a null terminated list of run length encoded elements. The run length encoding elements contain the following data members:

begin: The x coordinate corresponding to the beginning of a run.

end: The x coordinate of the end of the run.

y: The y coordinate of the run.

colour: The colour of the pixels in the run.

tag: The number of the run counting from the top left corner.

blobpointer: A pointer to a blob structure (section 5.3.2).

The run length encoding process creates a list of runs in memory that encapsulates each horizontal block of pixels. Unlike normal run length encoding this systems constrains the runs to appear on one line. The tag number is set to be the number of the run counting from the top left hand corner of the image. The tag value zero is reserved for terminating the list of runs. The blobpointer member is set to null when a run is created.

5.3.2 Blobs

A blob is an four connected region of pixels of the same colour. At present a blob contains the following data members:

area: The number of pixels in the blob.

colour: The colour code of the pixels in the blob.

xmax, xmin, ymax, ymin: The bounding dimensions of the blob.

Blobs are the output of the grouping process and the input to the final analysis stage of the vision process. Originally the sum of the moments of each pixel in the blob was stored within the blob structure so that the exact centroid of the blob could be recovered. To reduce the computational load this feature was dropped.

5.3.3 Grouping

After run length encoding blobs in the image are formed through a grouping process. Each run is examined one at a time and compared to the runs on the following row (algorithm 5). Rows that are four connected to the current row and the same colour are part of the same blob object. The basic approach to constructing these blobs is shown in algorithm 6. If both run elements have their blob pointers at null then a new blob object is manufactured for them and their pointers set to it. If only one of the two run elements has a non-null blob pointer then the empty pointer is set to point to the valid blob and the blob parameters are updated. The final case when both of the run elements have set blob pointers then the newest blob is deleted and both of them are set to the older blob.

Algorithm 5 Grouping


```

r ← start of RLE list
while r.tag ≠ 0
  n ← next RLE element for which n.y ≠ r.y
  if n.y = r.y
    for r ← each RLE element from r to n
      match(r against the RLE elements starting from n)
  r ← r+1

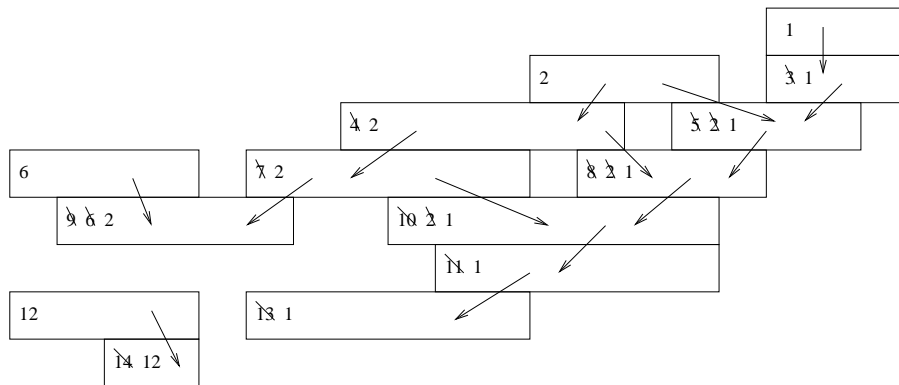
```

Algorithm 6 Matching

```
r is the element being compared to the next row of elements
n is the first element of the next row such that  $n.y = r.y + 1$ 
rowbelow  $\leftarrow$  n
while rowbelow.y = r.y + 1 {
  if r and rowbelow overlap in four connected space
    if rowbelow.tag > r.tag {
      if r.blobpointer = NULL
        create a new blob and attach it to r
      merge rowbelow.blobpointer with r.blobpointer
      rowbelow.blobpointer  $\leftarrow$  r.blobpointer
      rowbelow.tag  $\leftarrow$  r.tag
    } else {
      if rowbelow.blobpointer = NULL
        create a new blob and attach it to rowbelow
      merge rowbelow.blobpointer with r.blobpointer
      rowbelow.blobpointer  $\leftarrow$  r.blobpointer
      rowbelow.tag  $\leftarrow$  r.tag
    }
  rowbelow  $\leftarrow$  rowbelow + 1
}
```



(a) Run length encoded data before grouping.



(b) Propagation of tag numbers through run elements.

Figure 5.3: Run length encoded elements.

Each run element has a tag number which is originally set to the index of the run (a). Figure (b) shows the RLE elements after the grouping algorithm has passed through. Each element is matched against the elements in the row directly beneath it. When the elements are four connected the tag number is used to determine which element inherits from the other, the arrow indicating the direction of inheritance. Every RLE element connected by an arrow therefore contributes to the area and bounding rectangle of the blob.

The tag member of the run structure is used to determine which blob is deleted when regions are merged. Smaller tag numbers indicate that the blob associated with it is older. The merging process is simply a matter of increasing the bounding rectangle and area to accommodate the information from the new run or blob.

The grouping algorithm does not produce information showing which runs are members of each blob. The only information that is produced is the bounding rectangle and the number of pixels in the blob. While further passes of the grouping system would eventually be able to provide the exact shape of the blob there does not appear to be any need for this information.

5.4 Analysis

The final stage of feature extraction is the rejection of incorrect matches caused by errors in the first colour segmentation stage. The analysis stage examines the list of blobs and produces a set of objects that the software considers to be the best candidates for each type of object.

The first stage is to determine the approximate height of the wall. The area above the wall can then be ignored. This will remove the possibility of false matches in the area above the wall.

The wall object is the largest white blob in the image. Any object higher than a fixed offset above the wall object is ignored. The blue and yellow goals are likewise defined as the largest blue or white objects except that they must be in the valid region described by the wall. All of these objects must also be over a minimum size or they are not detected at all. This prevents a few noise pixels from being detected as an object.

Obstacles are defined as all black objects in the valid area of the screen that are above a minimum size and a list of obstacles is constructed containing all of these objects.

The ball is a special case in that it is rotationally invariant so it should maintain a fixed unity aspect ratio from any view point. Therefore the difference between the width and the height can be used as a simple test for the ball. This rejects red objects that do not have the correct dimensions for the ball.

5.5 Localisation and Line Detection

The two most powerful cues for localisation are the blue and yellow goals. From the angle of the goal and knowledge of the approximate height of the robot an idea of the robot's position on the field can be calculated. Another source of information for localisation are the edge lines. A fast way of obtaining orientation from the edge lines is to look for horizontal lines. An efficient method to do this is to use a Radon transform that only looks for lines that are horizontal.

Resuming the discussion in 2.2.2 it is clear that to only look for horizontal lines the projections need only be calculated when the vector is vertical. This is equivalent to simply summing the number of edge pixels in each horizontal row. The summation of edge pixels can be stored in a one dimensional array giving a profile of edge strength related to y coordinate. A significant horizontal line in the image will then translate into a large spike in the profile. A maximum detection system can then be used to find this point.

The maximum detection system is quite simple. The maximum point of the profile is first found. If this maximum is greater than a minimum threshold value and also greater by a minimum value than the second largest value in the profile than it is considered to correspond to a horizontal line. Aside from the detection of these cues for orientation no further work has been done in the field of localisation.

5.6 Training

The trainer lets the user provide representative samples of each colour class and construct a UV lookup table from these samples. The user can select a rectangular region of the image and specify it as being of a certain colour.

The selected region is passed to a two dimensional histogramming algorithm that counts the occurrence of each combination of U and V coordinates. This histogram is consequently of size 256×256 . This histogrammed data is then thresholded and assigned a code value by the user to produce a lookup table for one colour class. The user can interactively make further histograms of other colour classes and superimpose them onto the lookup table. This results in table like figure A.2(a).

Finally to link up the regions in the table and account for any sort of colour shift the whole lookup table is dilated with 3×3 a cross shaped mask (figure A.2(b)). As each

bitplane is dilated separately it is possible for the lookup table to contain two colour codes for the same U and V coordinate. As this would result in a pixel being detected as two colours the dilation is post processed to select only one colour code for each cell. Firstly if the cell in the table already contains a colour value then that cell is not permitted to be changed by the dilation. This prevents one colour from overwriting another colour during the dilation. Remaining overlaps are arbitrated by order of priority (table 5.2).

The finished lookup table can be trialled on a test image which duplicates the SH4 vision code to the point where it can show the location of segmented objects.

Priority Number	Colour Code
1	Red
2	Black
3	Blue
4	Yellow
5	White
6	Green

Table 5.2: Dilation Priority

The lookup table dilation uses this priority list to ensure that each position in the lookup table corresponds to only one colour code.

5.7 Memory Organisation

The image processing software requires a large amount of data to be stored within the 512 KB of on board memory. The YUV image; colour detected image; and the eroded image will use a large part of the memory. In addition to this there is also a large amount of symbolic information to be stored such as the run length encoded elements and the blob structures. Memory is allocated to all of these elements using fixed size buffers. Buffers are used instead of dynamically allocated data structures such as linked lists for several reasons, the chief being that a buffer is simpler to move through. These buffers use up about 400 KB of memory leaving the remaining memory for the program space, local variables and the stack.

Buffer	Size	Memory
YUV Image	240 × 180	168.75 KB
Lookup Table	256 × 256	64 KB
Colour Detected Image	240 × 180	42.2 KB
Eroded Image	240 × 180	42.2 KB
RLE Elements	3072	72 kB
Blobs	512	12 KB
Objects	64	768 Bytes
Edge Profile	180	720 Bytes
Total		401.8 KB

Figure 5.4: Memory Usage in the SH4.

The remaining 110 KB of the 512 KB of memory is used for the program, local variables and the stack.

Chapter 6

Testing and Results

6.1 Testing Methods

As there was no digital camera input for testing the code with the SH4 pictures were taken with a commercial handheld Casio camera. These 640×480 RGB colour images were converted to YUV format and subsampled to 240×180 using MATLAB and saved as raw binary files suitable for download to the SH4 board.

The SH4 code's performance was tested by using a serial port program to download and upload data from buffers on the SH4. The frame rate can be measured by toggling a LED on the SH4 board every time a frame is processed. The frequency of the LED, which is twice the frame rate, can be measured with an oscilloscope.

6.2 Training and Lookup Table

The training software was used (section 5.6) to construct a lookup table for detecting red, blue and yellow. The white and black regions of the image are detected by thresholding the Y coordinate. These thresholds are set at ± 50 above and below the brightness mean which is recalculated for each new frame of data.

The final lookup table shows the location in UV space of each of the three colours of interest (See figure A.3). The horizontal axis corresponds to increasing red chrominance from left to the right and the blue chrominance increases down the vertical axis. As expected then the region that detects blue is located at the bottom of the image and the

red region is located on the top right of the image. It is interesting to note that there is a relationship between the red and blue chrominance of the ball. Roughly speaking as the red chrominance of the ball increases the blue chrominance decreases. This relationship is preserved because the lookup table uses arbitrary shapes as opposed to rectangular regions.

6.3 Colour Detection

6.3.1 Reliability and Robustness of Colour Detection

Two examples of the colour detection system in action can be seen in figure A.4. It is clear that in general the performance of the colour detection stage is fairly good however, there is a problem detecting yellow in the image (figure A.4 (b)). This is caused by problems with the brightness thresholding of the image. Basically the Y values of yellow approach those of white leading to confusion as to the difference between white and yellow. Other less bright colours such as red, blue or green are generally detected well.

Aside from misinterpreting yellow as white, the detection of black and white is very successful. This indicates that readjusting the Y thresholds for each frame is an intelligent strategy, especially as the lighting level within the environment can change dramatically.

Another problem encountered during colour detection is that particularly dark shades of green can be detected as red (figure A.6). This leads to extra red in the colour detected image that causes problems later in the image processing system. It appears these problems are caused by a poorly constructed red region in the colour lookup table. Redefining the colour table to be more conservative in red, ie making the red region smaller in the table, may solve this problem.

6.3.2 Speed of Colour Detection

The software was tested on the SH4 processor board and the frame rate recorded for a few test images. The performance of the fast colour detection system at various rates of pixel skipping is given in table 6.1. The fast colour detection by pixel skipping gives reasonable accuracy with a general speed saving of about 12 ms when thresholding every fourth pixel. This is an acceptable error for reducing the speed by approximately 25%.

The speedup is not as large as expected, probably because the branching interferes with the pipeline of the SH4. A subsampling rate of about 20 gives the fastest speed and increasing the subsampling rate beyond 20 does not reduce the processing time. The limiting factor on the subsampling though is that the system is not guaranteed to detect any object narrower than the sampling interval.

Subsampling	Time (ms)	Pixels in Error
none	40	0
6	23.2	331
5	25.2	251
4	28	240
3	32.8	230
2	35.6	117

Table 6.1: Subsampling speed and error

This is the time taken to perform colour detection on a particular 240×180 pixel image. The figure pixels in error is the number of pixels that are different to the colour detected image created with no subsampling. In general terms the error is particularly small (331 errors is 0.8% error for a 240×180 image).

6.4 Morphological Operations

The removal of noise by morphological operations using a 1×3 erosion was generally successful particularly in removing small amounts of red noise. The benefits of the noise reduction can be seen in figure A.5.

The slow erosion process listed in algorithm 3 takes approximately 25 ms to complete. The erosion faster algorithm (4) only reduces the execution time by three milliseconds. A better method for implementing an erosion is outlined in section 7.1.3.

6.5 Run Length Encoding

Run length encoding is another costly operation taking 49 ms to encode the test image into about 180 run length elements. Clearly further operations that use the RLE data will operate quickly as the amount of data to process has been dramatically reduced. Also it is apparent that the buffer allocated for the RLE data is much larger than it needs to be,

as space has been allocated for 3072 elements. The reason for this is that if the colour detection fails badly or is miscalibrated the number of RLE elements can become very large and potentially overwrite other sections of the program.

6.6 Grouping

The grouping stage which converts the RLE elements into blobs executes quickly and efficiently as it uses symbolic data as its input rather than an image array. Under the current grouping system an isolated RLE element will not become a blob because the system only creates a blob when it combines two or more RLE elements. This effectively means that there is no need for a morphological erosion in the vertical direction.

One minor bug detected in the code was when a blob was merged with itself the blob was deleted, this was corrected by preventing the blobs from being merged with themselves.

6.7 Blob Analysis and Object Detection

Blob analysis is the most difficult and error prone section of the system. The system can easily make errors, usually these are failures result in no object being detected rather than an imaginary object being found. Two successful examples of the output are shown in figure A.7. The shadow under the ball is detected as an obstacle.

The detection of the goals is particularly successful although that is not a great accomplishment as there should only be one goal in any image and they are large objects. Even though yellow is not particularly well segmented, the large size of the object is sufficient to allow simple detection of the yellow goal. Obstacles are also easily segmented for the same reasons. Although an occasional lighting problem causes a few parts of the field to be detected as black causing imaginary black objects to appear.

The difficulty in the image segmentation process is the red ball. Usually the source of error is poor detection of red, for example figure A.6. The rejection of red noise is not particularly successful at the present stage causing the blob analysis stage to fail to correctly locate the ball. Another source of error is that sometimes the software decides that lines on the field such as the centre line are in fact the edge line. This incorrect decision causes the algorithm to ignore the majority of the image.

6.8 Line Detection

The process for detecting horizontal lines is quite successful at rejecting inappropriate lines (figure 6.1). As the system will not report a line unless the maximum of the edge profile is significantly greater than the second largest value most errors are avoided. For these trials the minimum acceptable peak on the edge profile is 35 and the minimum difference between the maximum and the second largest peak is 17. The software is however unable to differentiate between the lines in the middle of the field and the edge lines of the field. The process of developing the edge profile is slow taking 24 ms to complete, but the peak detection stage was practically instantaneous.

The problem with the line finding at the present stage is that it does not use employ any information from other stages of the process. In particular during the blob analysis stage the bounding rectangle of the edge line is found. If the search was limited to finding horizontal lines in this region the line finding would be faster and less likely to make mistakes.

6.9 Total Processing Time

In total the system operates at 10 fps without the edge detection code in place or at 8 fps with the edge detection code. The breakdown of time spent in the various stages of the image processing system is shown in table 6.2. It is clear that the parts of the system that operate on symbolic data execute much faster than those parts which use image buffers. The code then operates at 348 kilopixels per second which is more than expected (section 3.2).

Task	Time (ms)
Colour Detection	28
Erosion	22
RLE	49
Grouping	Negligible
Analysis	Negligible
Edge Detection and Summation	25
Edge Analysis	Negligible
Total Time	124 ms

Table 6.2: Breakdown of Processor Time

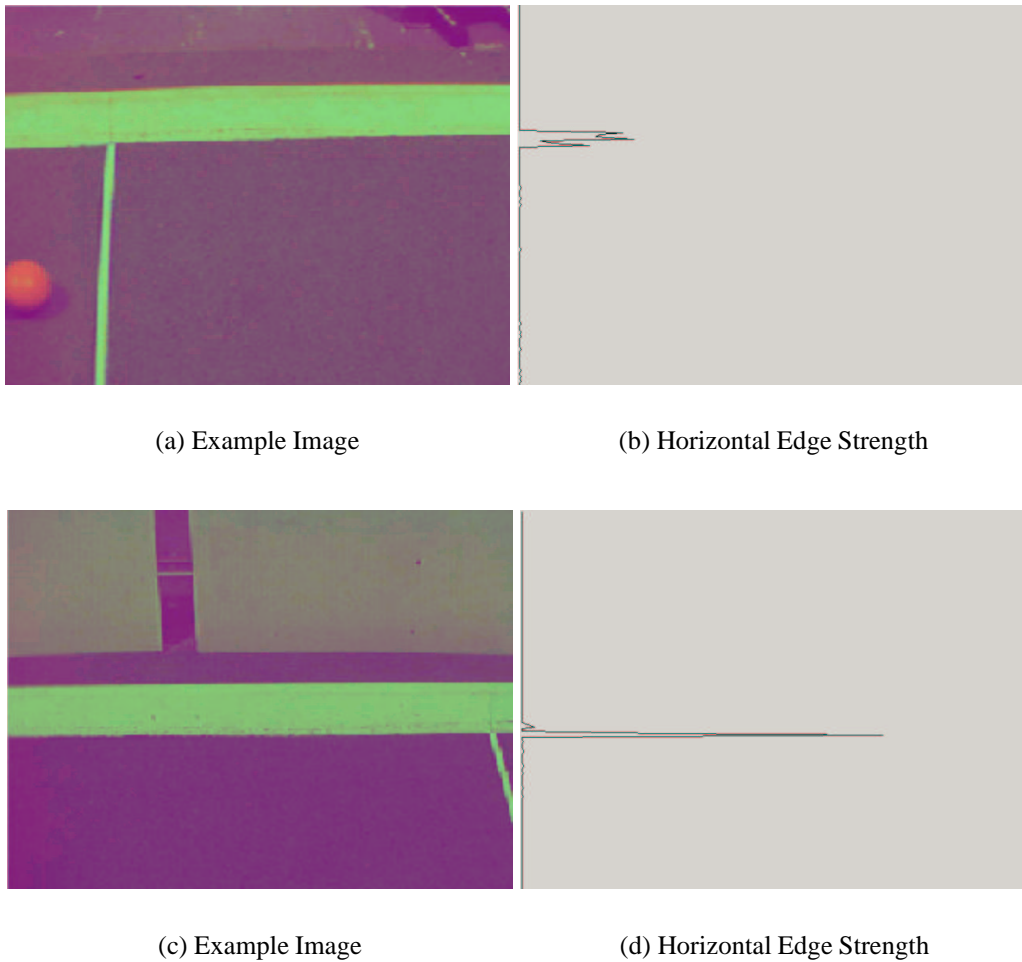


Figure 6.1: Examples of the Line Detection Process

In the first image **(a)** there is no dominant peak in the edge strength profile **(b)**. The second image **(c)** does have a definite peak in the edge profile **(d)**. The edge detection software recognises this as a horizontal line in the original image.

Chapter 7

Future Work

The vision system presented here would need many changes to be developed into a system that would be usable in a RoboCup scenario. Additionally there are many changes that could be made to provide incremental improvements to the process either increasing the frame rate or reducing the error.

7.1 Optimisations

Several optimisations can be made to the code that would reduce the time taken to process a frame or the amount of error in the output of the system.

7.1.1 Use of the Field Programmable Gate Array

The most significant change to the system would be moving the colour lookup stage from the SH4 to the board's field programmable gate array (FPGA). Currently the colour lookup process takes about 30 ms to complete (table 6.1). This comprises taking a sample mean of the frame's brightness and performing colour detection on each pixel. The colour detection could be implemented in a ROM as shown in figure 7.1. The brightness (Y) data could be used in the FPGA as well to move the colour detection stage to the FPGA. Colour detection on the FPGA could operate at the same rate that YUV data passes into the FPGA, so each pixel would be analysed and the result sent to the SH4. Not only would this technique reduce the processing by 30ms for a 240×180 image but it would also reduce bus traffic and the amount of memory needed to store the image in the SH4.

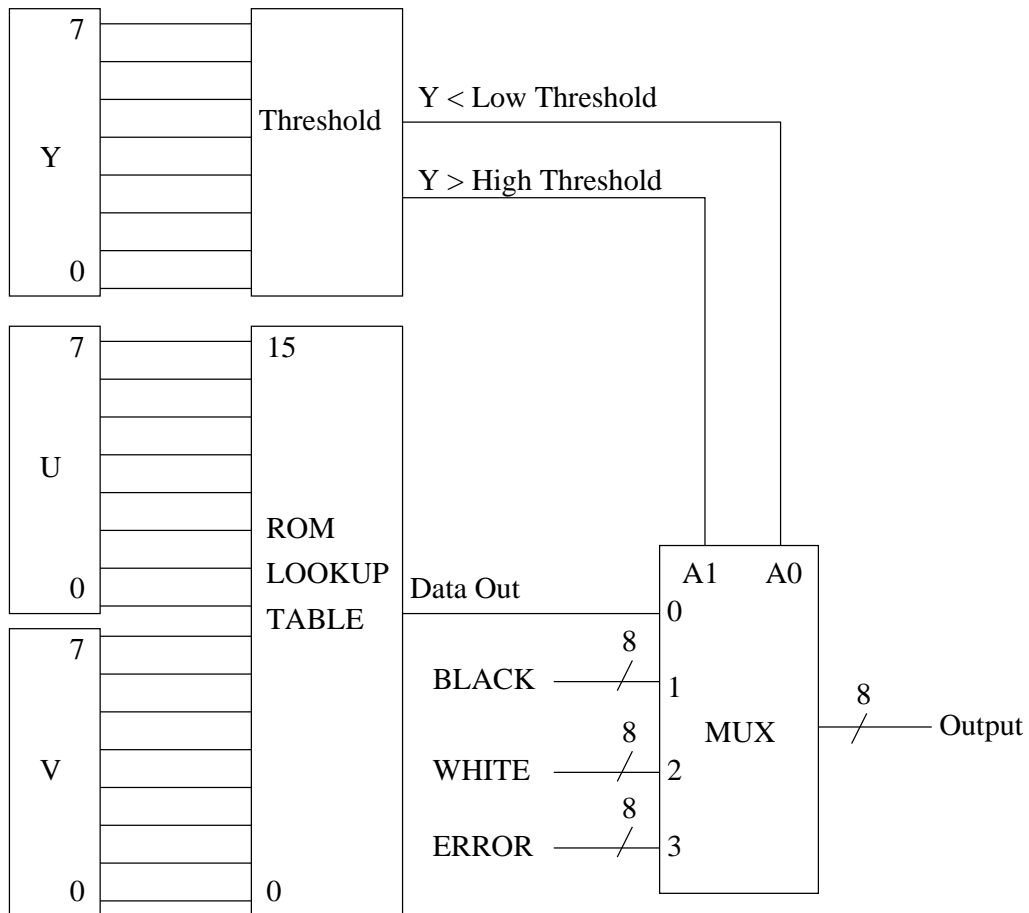


Figure 7.1: Logic Circuit Implementation of Colour Detection.

The U and V values are used as an address into a ROM that contains the values of the lookup table. The thresholded Y value is used to control a multiplexer (MUX) that selects between the lookup table's output and the constant values BLACK and WHITE . The final input to the multiplexer is reserved as ERROR because it should be impossible to have both $Y < \text{Low Threshold}$ and $Y > \text{High Threshold}$.

7.1.2 Detection of Green

The reason that green is not detected by the system when working on robot soccer is that it was thought that the large amount of green in each frame would lead to a significantly larger number of run length encoding elements to process in the grouping stage and would lead to no further useful information. However tests show that the effect of adding green detection to the whole colour and blob growing system only increased the frame time by a few milliseconds. The advantage of detecting green is that white pixels with green pixels directly underneath would be good candidates for being part of the edge line.

7.1.3 Faster Morphology

The current morphological system is based on a logical 1×3 erosion of the colour detected image. This erosion uses logical operators but a different solution could be developed. A 1×3 erosion is equivalent to shortening each run length element by one at both ends. This leads to the idea of implementing a complete set of $1 \times n$ morphological operations on the run length encoded data instead of the colour detected image.

Rather than use addition and subtraction to perform these operations it would better to apply them as the run length elements are constructed. For example if an RLE element is only of length one then removing it from the list would be equivalent to an opening of the image with a 1×3 structuring element. A closing could then performed by running the through the list of RLE elements and joining ones that are almost adjacent. This has the potential to reduce the processing time by up to 22 ms.

7.1.4 Noise Rejection

Further work can definitely be done rejecting incorrect matches for red. Some other techniques were considered to remove but never implemented either for reasons of time or apparent complexity.

- With a definite measure of distance from the robot to the ball the approximate size of the ball can be calculated. The difference between the actual and expected sizes gives another measure of error[2]. This technique was not implemented because none of the physical constants such as the height of the robot were available.

- Relating the area of the ball to the bounding rectangle. The trainer program currently displays red blobs with a bounding rectangle and circle with a radius based on the blob area. This gives an indication of the relationship between bounding box and area. Tests show that the relationship is fairly solid and will be suitable for rejecting noise as the relationship noticeably breaks when applied to noisy pixels.
- The final technique considered was to measure the edge length of the blobs and relate this to the area of the pixel. This would require measuring the edge though which makes it less attractive than the previous technique.

7.2 Additions

7.2.1 Coordinate Transforms and iPaq Software

To be useful for robotic motion planning the image coordinates of each object must be transformed into three dimensional coordinates relative to some fixed location. To do this the position and orientation of the camera relative to the robot must be known. For some robots like the ViperRoos this is a fixed position but for more complex robots such as the GuRoo the camera is capable of independent motion. Consequently the vision software must have access to the output of the internal sensors of the robot that report the angles of the limbs and head.

On the GuRoo this information will be passed to the iPaq so it would make sense for the coordinate transform stage of the vision system to be performed here. The two dimensional image coordinates of objects would be transmitted to the iPaq which would then build a transform matrix from the internal sensors information. The matrix would be used to calculate the three dimensional positions of all of the objects.

The PC based trainer is planned to be ported to the iPaq. The trainer would be able to receive the YUV images from the SH4 and use these as test images for training. In addition to this the run length encoded data from the SH4 could be transmitted quickly to the iPaq which could then decode this data and display it on the screen for debugging purposes.

7.2.2 Adapting to the output of the OV7620 Camera

At this point in time the OV7620 camera is not interfaced with the SH4 processor board and the system uses test images taken from a hand held digital camera. These images are converted from RGB to YUV and subsampled using MATLAB. The OV7620 is expected to have at least slightly different characteristics to the trial images. The lookup table would have to be recalibrated and the offsets used for calculating the brightness thresholds would have to be changed.

The use of a real camera would introduce problems of distortion and focusing to the system. Also the availability of a real time data feed would make errors in the system much more apparent than they currently are.

Chapter 8

Conclusion

A basic architecture for colour segmentation has been developed. The system is able to perform fast colour detection using a combination of lookup tables and brightness thresholds, that has been trialled on an embedded processor. The colour detected image has been successfully converted into blobs using a two stage process of run length encoding and grouping.

The YUV colour space is extremely useful for reducing the effect of brightness changes in the image. The other advantage is that YUV allows the use of unchanging lookup tables in U and V which provide an extremely fast way of performing colour detection.

The speed of execution is slightly better than expected (section 3.2). The majority of time is spent performing colour detection, morphological operations and run length encoding. There are many opportunities to increase the speed of the system, including using a field programmable gate array to perform colour detection (section 7.1.1).

The system is capable of detecting most objects although more work does need to be done preventing incorrect matches. Some techniques to correct these problems are outlined in section 7.1.4.

A GUI system has been developed that allows the user to construct a colour lookup table using test images. The performance of the colour detection and later stages of the software can be examined within the GUI environment. This software has also proved useful during the development and debugging of the vision software.

The software is at the stage where it could provide useful information to robot navigating a colour coded environment. However the system needs to be made much more robust before it can operate reliably.


Bibliography

- [1] M. Abe, *Building Cross Development Environment Targetting SH4 System*, www.linuxsh.sourceforge.net/docs/abe/2001320-gcc2.97/README_E.php3, August, 2001.
- [2] T. Bandlow et al, *Fast Image Segmentation, Object Recognition and Localisation in a RoboCup Scenario*, Lecture Notes in Artificial Intelligence Volume 1856, Springer, Berlin, 2000.
- [3] R. Bartelds, et al, *Clockwork Orange: The Dutch RoboSoccer Team*, RoboCup 2001, Springer-Verlag, Berlin, 2001.
- [4] A. Berry, *Soccer Robots with Local Vision*, Honours Dissertation, Univ. of Western Australia, Dept. of Electrical and Electronic Engineering, 1999.
- [5] Blower, A., *Development of a Vision System for a Humanoid Robot*, Undergraduate Thesis, University of Queensland, 2001.
- [6] Browning, B., *Robust Vision For Robot Soccer*, Undergraduate Thesis, University of Queensland, 1999.
- [7] Bruce, J., *Realtime Machine Vision Perception and Prediction*, Undergraduate Thesis, Carnegie Mellon University, 2000.
- [8] Bruce, Balch, Veloso, *Fast and Inexpensive Colour Segmentation for Interactive Robots*, Proceedings of the 2000 IEEE/RSJ International Conference on Intelligent Robots and Systems, 2000.
- [9] M. Chang, et al, *ViperRoos 2000*, RoboCup 2000: Robot World Cup IV Lecture Notes in Artificial Intelligence 2019, Springer-Verlag, Berlin, 2001.

- [10] R. Gonzalez and R. Woods, *Digital Image Processing*, Addison-Wesley, Reading Massachusetts, 1992.
- [11] B. Horn, *Robot Vision*, The MIT Press, Cambridge Massachusetts, 1986.
- [12] Hosking S, *High Speed Peripheral Interface*, Undergraduate Thesis, University of Queensland, 2001.
- [13] Omnivision, *OV7620 Product Specifications - Revision 1.3*, 2000.
- [14] W. Pratt, *Digital Image Processing, Second Edition*, John Wiley & Sons, New York, 1991.
- [15] D. Paulus and J. Hornegger, *Applied Pattern Recognition: a Practical Introduction to Image and Speech Processing in C++*, Verlag Vieweg, Braunschweig Germany, 1998.
- [16] *RoboCup Medium Size League Rules*, smart.uni-ulm.de/ROBOCUP/f2000/rules01/rules2001.html, accessed 12 August 2001.
- [17] *Robocup Small Size League Rules*, parrotfish.coral.cs.cmu.edu/robocup-small, accessed 12 August 2001.
- [18] *Robocup Website*, www.robocup.org, accessed 12 August 2001.
- [19] R. Schilling, *Fundamentals of Robotics, Analysis and Control*, Prentice Hall, New Jersey, 1990.
- [20] P. Stone, *RoboCup-2000 The Fourth Robotic Soccer World Championships*, AI Magazine, Vol 22, Number 1, Spring 2001.
- [21] Wagstaff M, *Mechanical Design and Internal Sensors for a Humanoid Robot*, Undergraduate Thesis, University of Queensland, 2001.


Appendix A

Colour Images



(a) RGB Image

(b) Y Component




(c) U and V Components


(d) Colour as a function of U and V

Figure A.1: YUV Image Representation

The YUV representation of the image **(a)** can be easily described in two components. The Y component is simply the brightness of each pixel **(b)**. The UV components describe the colour of the pixel **(c)**. The variation of colour with U and V is shown in figure **(d)**. V is the horizontal axis and U is the vertical axis with 0,0 being the top left corner.



(a) Lookup Table Before Dilation



(b) Lookup Table After Dilation

Figure A.2: Trainer Lookup Table

The trainer program uses samples of colours provided by the user to construct the initial lookup table **(a)**. The table is then dilated to fill up the gaps in the image **(b)**.






Figure A.3: An example colour lookup table.
The lookup table used for testing, it is configured to detect red, blue and yellow. The horizontal access is V and the vertical access is U with a scale 0 to 255 on both axes.




(a) View of Blue Goal




(b) View of Yellow Goal

Figure A.4: Colour Detected Image

The blue, white and black areas are successfully detected. The bright yellow is often interpreted as white.




(a) View of Blue Goal



(b) View of Yellow Goal

Figure A.5: Eroded Images

The erosion by a 1×3 mask removes any pixel groups whose dimensions are smaller than 1×3 .



(a) Colour Detected

(b) YUV image and Object Locations





Figure A.6: An Example of Poor Colour Detection

The colour detection fails here because the area of the field on the right of the image is being segmented as red. The brightness normalised version of this image shows that the UV data in this region is borderline between red and green. A better defined lookup table may correct this problem.



(a) View of Blue Goal

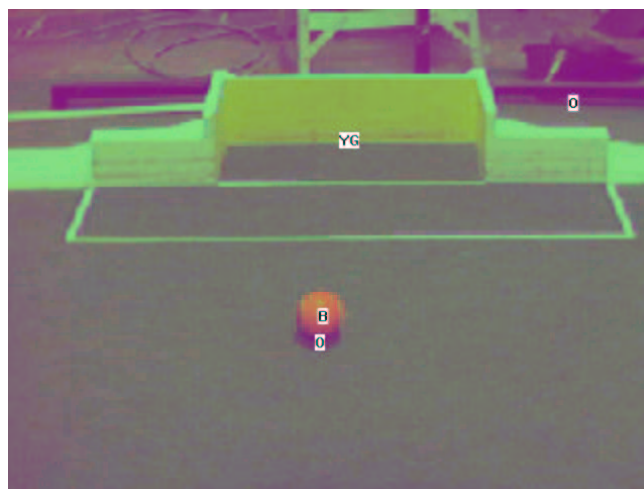


Figure A.7: Detected Objects.

The objects in the image detected by the vision software. The B represents the ball, O represents obstacles(black objects), and the YG and BG represent the yellow and blue goals. The shadow of the ball is interpreted as an obstacle.

Appendix B

Code

The image processing support functions are contained in `segment.c` and `segment.h`. The main frame processing function from `imtest.c` shows the order in which the functions are called.

B.1 Imdemo.c

Program B.1: `imtest.c`

```
1 void imdemo() {
2     quick_lookup(SRAMDATA.yuv_image, SRAMDATA.table, SRAMDATA.color_detected);
3     //lookup_bytable(SRAMDATA.yuv_image, SRAMDATA.table, SRAMDATA.color_detected);
4     erode13(SRAMDATA.color_detected, SRAMDATA.eroded, 240, 180);
5     rlecode(SRAMDATA.eroded, SRAMDATA.rle_encoded);
6     end_blobs = group(SRAMDATA.rle_encoded, SRAMDATA.blobs);
7     end_obstacles = analyse(SRAMDATA.blobs, end_blobs, SRAMDATA.obstacles, &ball, &bluegoal, &yellowgoal);
8     edgeprofile(SRAMDATA.eroded, SRAMDATA.profile);
9     findedge1ine(SRAMDATA.profile);
10 }
```

B.2 Segment.h

Program B.2: `segment.h`

```
1 #ifndef SEGMENT_H
2 #define SEGMENT_H
3
4 #include "image.h"
5 // #include <stdio.h>
6 // #include <stdlib.h>
7
8 #ifdef WIN32
9     #include <stdio.h>
10 #else
11     #include <sh/gio.h>
12 #endif
13
14 // Memory Allocation
15 #define ENCODEDSPACE 4096
16 #define NUMBERBLOBS 512
17 #define NUMBEROBSTACLES 64
18
19 // Image size
```

```

20 #define TABLEWIDTH 256
21 #define TABLESIZE (TABLEWIDTH * TABLEWIDTH)
22
23 #define IMWIDTH 240
24 #define IMHEIGHT 180
25 #define IMSIZE (IMWIDTH*IMHEIGHT)
26 #define IMBYTESIZE (IMSIZE*3)
27
28
29 // Color Codes
30 #define FIELDCODE 1
31 #define WALLCODE 2
32 #define BALLCODE 4
33 #define BLACKCODE 8
34 #define BLUECODE 16
35 #define YELLOWCODE 32
36
37 // Minimum Sizes
38 #define MINSIZE 6
39
40 #define MINWALLSIZE 140
41 #define MINOBSTACLESIZE 6
42 #define OBJECTLISTSIZE 10
43 #define TOPWALLOFFSET 0
44 #define MINGOALSIZE 10
45 #define MINEDGELINE 35
46 #define MINEDGEDIFF (MINEDGELINE/2)
47
48 /* Aspect ratio of a square for the chosen
49 image dimensions
50 (IMWIDTH/640)/(IMHEIGHT/480) */
51 #define SQUARERATIO 1
52
53 #define MIN(x, y) (((x) < (y)) ? (x) : (y))
54 #define MAX(x, y) (((x) > (y)) ? (x) : (y))
55
56 // #define USEPRINTF
57
58 void find_centroid(image* segmented_image);
59
60 struct blobtag {
61     // struct blobtag * next;
62     // struct blobtag * prev;
63     int area;
64     int color;
65     int xmax, xmin, ymax, ymin;
66 };
67
68 typedef struct blobtag blob;
69
70 struct rletag {
71     int begin;
72     int end;
73     int color;
74     int tag;
75     int y;
76     blob* blobpointer;
77 };
78 typedef struct rletag rle;
79 struct objecttag {
80     int distance, bearing, color;
81 };
82
83 typedef struct objecttag object;
84
85 void lookup_bytable(unsigned int* the_image, UCHAR* table,
86                    UCHAR* out_image); //, int* profile);
87
88 blob* group(rle* input, blob* blobs);
89 void rlecode(UCHAR* segmented_image, rle* output);
90 void reconstruct(rle* input, image* output);
91 void color_reconstruct(rle* input, UCHAR* output);
92
93 // void dilate22(UCHAR* the_image, UCHAR* out_image);
94 void code2rgb(UCHAR* input, UCHAR* output, int size);
95 // object* analyse(blob* blobdata, blob* endblob, object* results, object* pball);
96 object* analyse(blob* blobdata, blob* endblob, object* results,
97                object* pball, object* pygoal, object* pbgoal);
98
99 int findedgeline(int* profile);
100 void erode33(UCHAR* input, UCHAR* output, int width, int height);
101 void erode13(UCHAR* input, UCHAR* output, int width, int height);
102
103 void histogram_region(UCHAR* the_image, int sx, int sy, int ex, int ey, int* hist);
104 void thresh(int* hist, int code, int level, UCHAR* table);
105
106 void edgeprofile(UCHAR* target, int* profile);
107
108 #ifdef WIN32
109     void quick_lookup(UCHAR* the_image, UCHAR* table, UCHAR* out_image);
110     void lookup_bytable(UCHAR* the_image, UCHAR* table,
111                        UCHAR* out_image);
112 #else
113     void quick_lookup(unsigned int* the_image, UCHAR* table, UCHAR* out_image);
114     void lookup_bytable(unsigned int* the_image, UCHAR* table,
115                        UCHAR* out_image);
116 #endif
117 #endif
118 /* SEGMENT_H */

```

B.3 Segment.c

Program B.3: segment.c

```

1  #include "segment.h"
2  #define WORD long int
3  #include "angletable.h"
4  // #include "imageio.h"
5
6  #ifndef WIN32
7      #include <../vbsh4/test.h>
8      #include <sh/sci.h>
9  #endif
10
11 // #include <assert.h>
12
13 // y thresholds
14 unsigned int ytlow, ythigh;
15
16
17 // Sample Mean Constants
18 #define YMEAN_SAMPLEPOWER 10
19 #define YMEAN_NUMSAMPLE 1024
20 #define YMEAN_INCREMENT (IMSIZE/YMEAN_NUMSAMPLE)
21
22 #define CD_SUBSAMPLE 4
23
24 #ifdef WIN32
25
26     int ymean(UCHAR* the_image) {
27         UCHAR* impoiter = the_image;
28         int ytotal=0;
29         int i,y;
30         for (i = 0; i < YMEAN_NUMSAMPLE; i++) {
31             y = *impoiter;
32             ytotal+=y;
33             impoiter += YMEAN_INCREMENT*3;
34         }
35         return (ytotal >> YMEAN_SAMPLEPOWER);
36     };
37
38 #define Y_POSOFFSET (50)
39 #define Y_NEGOFFSET (50)
40
41     inline int lookup_pixel(UCHAR* impoiter, const UCHAR* table) {
42         UCHAR y = *impoiter;
43         UCHAR u = *(impoiter+1);
44         UCHAR v = *(impoiter+2);
45         // unsigned int out;
46         if (y > ythigh) return WALLCODE;
47         if (y < ytlow) return BLACKCODE;
48         return table[u*256+v];
49     }
50
51     void quick_lookup(UCHAR* the_image, UCHAR* table,
52                     UCHAR* out_image) {
53         unsigned int output;
54         unsigned int last_output = 0;
55         const UCHAR* end_addr = the_image + (IMBYTESIZE);//used to be -subsample
56         UCHAR* outpointer = out_image;
57         UCHAR* impoiter = the_image;
58         UCHAR* next_pixel;
59         // UCHAR* end_out = out_image + IMSIZE;
60         // int y,u,v;
61         int mean = (ymean(the_image));
62         int i;
63         ythigh = (mean) + Y_POSOFFSET;
64         ytlow = (mean) - Y_NEGOFFSET;
65         *outpointer = lookup_pixel(impoiter, table);
66         while (impoiter < end_addr) {
67             next_pixel = impoiter+CD_SUBSAMPLE*3;
68             output = lookup_pixel(next_pixel, table);
69             if (output == last_output) {
70                 for (i = 0; i < CD_SUBSAMPLE; i++) {
71                     *(++outpointer) = output;
72                 }
73                 // printf("s");
74                 impoiter= next_pixel;
75             } else {
76                 for (i = 0; i < CD_SUBSAMPLE-1; i++) {
77                     impoiter+=3;
78                     outpointer++;
79                     *outpointer = lookup_pixel(impoiter, table);
80                 };
81                 // then do the last one which has already been looked up
82                 outpointer++;
83                 impoiter= next_pixel;
84                 *(outpointer) = output;
85                 last_output = output;
86             }
87             // assert(outpointer < end_out);
88         };
89     };
90
91 #else
92 #define Y_POSOFFSET (50<<16)
93 #define Y_NEGOFFSET (50<<16)
94
95     int ymean(unsigned int* the_image) {
96         unsigned int* impoiter = the_image;    int ytotal=0;

```

```

98         int i,y;
99         for (i = 0; i < YMEAN_NUMSAMPLE; i++) {
100             y = ((*impointer)>>16)&0xFF;
101             ytotal+=y;
102             impointer += YMEAN_INCREMENT;
103         }
104         return (ytotal >> YMEAN_SAMPLEPOWER);
105     };
106
107     inline int lookup_pixel(unsigned int* impointer, const UCHAR* table) {
108         unsigned int data = *impointer;
109         unsigned int y = (data);
110         //unsigned int out;
111         if (y > ythigh) return WALLCODE;
112         if (y < ytlow) return BLACKCODE;
113         return table[data & 0x0000FFFF];
114     }
115
116
117
118     void quick_lookup(unsigned int* the_image, UCHAR* table,
119                     UCHAR* out_image) {
120         unsigned int output;
121         unsigned int last_output = 0;
122         const unsigned int* end_addr = the_image + (IMSIZE); //used to be -subsample
123         UCHAR* outpointer = out_image;
124         unsigned int* impointer = the_image;
125         int* next_pixel;
126         // UCHAR* end_out = out_image + IMSIZE;
127         // int y,u,v;
128         int mean = (ymean(the_image))<<16;
129         int i;
130         ythigh = (mean) + Y_POSOFFSET;
131         ytlow = (mean) - Y_NEGOFFSET;
132         *outpointer = lookup_pixel(impointer, table);
133         while (impointer < end_addr) {
134             next_pixel = impointer+CD_SUBSAMPLE;
135             output = lookup_pixel(next_pixel, table);
136             if (output == last_output) {
137                 for (i = 0; i < CD_SUBSAMPLE; i++) {
138                     *(++outpointer) = output;
139                 }
140                 //printf("s");
141                 impointer= next_pixel;
142             } else {
143                 for (i = 0; i < CD_SUBSAMPLE-1; i++) {
144                     impointer++;
145                     outpointer++;
146                     *outpointer = lookup_pixel(impointer, table);
147                 };
148                 // then do the last one which has already been looked up
149                 outpointer++;
150                 impointer= next_pixel;
151                 *(outpointer) = output;
152                 last_output = output;
153             }
154             // assert(outpointer < end_out);
155         };
156     };
157
158
159     // This code is slower than the quick lookup version by about 4 fps at 240*180
160     void lookup_bytable(unsigned int* the_image, UCHAR* table,
161                       UCHAR* out_image) { //, int* profile) {
162         const unsigned int* end_addr = the_image + IMWIDTH * IMHEIGHT;
163         UCHAR* outpointer = out_image;
164         unsigned int* impointer = the_image;
165         unsigned int mean = ymean(the_image)<<16;
166         ythigh = mean + Y_POSOFFSET;
167         ytlow = mean - Y_NEGOFFSET;
168         //int i;
169         while (impointer < end_addr) {
170             *outpointer = lookup_pixel(impointer++, table);
171             outpointer++;
172         };
173     };
174
175
176 #endif
177
178 // Slow Erosion
179 /*void erode13(UCHAR* input, UCHAR* output, int width, int height) {
180     register int x,y;
181     register UCHAR* inpointer = input + 1;
182     register UCHAR* outpointer = output + 1;
183     UCHAR* end_addr = input+IMSIZE -1;
184     while (inpointer < end_addr) {
185         *outpointer = *(inpointer-1) & *(inpointer) & *(inpointer+1);
186         outpointer++;
187         inpointer++;
188     }
189 }*/
190
191
192 //erode13 only works for a specific image width
193 #if IMWIDTH != 240
194     #error "Incorrect width"
195 #endif
196
197
198 void erode13(UCHAR* input, UCHAR* output, int width, int height) {
199     int x,y;
200     UCHAR* inpointer = input;
201     UCHAR* outpointer = output;

```



```

202     unsigned int a,b,c;
203
204     for (y = 0; y < height; y++) {
205         *outpointer++ = 0;
206         inpointer++;
207         a = *(inpointer - 1);
208         b = *(inpointer);
209         c = *(inpointer + 1);
210
211         for (x = 1; x < width - 4; x+=3) {
212             *outpointer++ = a & b & c;
213             inpointer++;
214
215             a = *(inpointer+1);
216             *outpointer++ = b & c & a;
217             inpointer++;
218
219             b = *(inpointer+1);
220             *outpointer++ = c & a & b;
221             inpointer++;
222
223             c = *(inpointer+1);
224
225         }
226         //inpointer and outpointer now have finished column 178
227         //and still point to it.
228         *outpointer++ = 0;
229         *outpointer++ = 0;
230         inpointer+=2;
231     }
232 }
233
234
235 void rlecode(UCHAR* segmented_image, rle* output) {
236     // watch for runs that end at the rightmost column
237     register UCHAR* impointer = segmented_image;
238     rle* rlepointer = output;
239     #ifndef NDEBUG
240         rle* spacecheck = output + ENCODEDSPACE;
241     #endif
242     int tagcounter = 1;
243     register int x,y;
244     register int data;
245     int rowcount;
246     int color, oldend;
247     //int height = segmented_image->height;
248     //int width = segmented_image->width;
249     int inrun = 0;
250     for (y = 0; y < IMHEIGHT; y++) {
251         inrun = 0;
252         rowcount = 0;
253         oldend = 0;
254         for (x = 0; x < IMWIDTH; x++) {
255             data = *impointer;
256             if (inrun && (data != color)) {
257                 rlepointer->end = x-1;
258                 oldend = rlepointer->color;
259                 #ifdef USEPRINTF
260                     printf("[tag: %d b: %d ", rlepointer->tag, rlepointer->begin);
261                     printf("e: %d c: %d]\n", rlepointer->end, rlepointer->color);
262                 #endif
263                 rlepointer++;
264                 inrun = 0;
265
266                 #ifdef USEPRINTF
267                     rowcount++;
268                     printf("rowcount %d\n", rowcount);
269                 #endif
270             }
271
272             //logic this line
273             if (!data && !inrun) {
274                 rlepointer->begin = x;
275                 color = data;
276                 rlepointer->color=color;
277                 rlepointer->tag = tagcounter++;
278                 rlepointer->blobpointer = 0;
279                 rlepointer->y = y;
280                 inrun = 1;
281                 #ifdef USEPRINTF
282                     printf("{ Start Run x:%d y:%d}", x, y);
283                 #endif
284             };
285
286             impointer++;
287
288             //assert(rlepointer < spacecheck);
289             //assert(rlepointer >= output);
290         };
291         if (inrun) {
292             rlepointer->end = IMWIDTH - 1;
293             #ifdef USEPRINTF
294                 printf("[tag: %d b: %d ", rlepointer->tag, rlepointer->begin);
295                 printf("e: %d c: %d]\n", rlepointer->end, rlepointer->color);
296             #endif
297             rlepointer++;
298             //Set the next one to zero
299             inrun = 0;
300         }
301     };
302     rlepointer->tag = 0; //null terminate the run array
303 }
304
305 // Converts the rle objects back to colour codes

```

```

306 void reconstruct(rle* input, image* output)
307 {
308     int x,y;
309     UCHAR* data = output->data;
310     rle* rlepointer = input;
311
312     //for each row of the image
313     y = 0;
314     //for all the rle runs on this row
315     while(rlepointer->tag) {
316         for(x = rlepointer->begin; x < rlepointer->end + 1; x++) {
317             *(data + x + output->width*rlepointer->y) = rlepointer->tag;
318         }
319         rlepointer++;
320     };
321
322 };
323
324 //Colour codes to RGB
325 void code2rgb(UCHAR* input, UCHAR* output, int size)
326 {
327     UCHAR* impointer = input;
328     UCHAR* out_pointer = output;
329     const UCHAR* end_addr = input + size;
330     UCHAR r,g,b;
331     while(impointer < end_addr) {
332         switch(*(impointer++)) {
333             case FIELDCODE:
334                 r = 0;           g = 255;       b = 0;
335                 break;
336             case WALLCODE:
337                 r = 255;        g = 255;       b = 255;
338                 break;
339             case BALLCODE:
340                 r = 255;        g = 0;         b = 0;
341                 break;
342             case BLACKCODE:
343                 r = 0;          g = 0;         b = 0;
344                 break;
345             case BLUECODE:
346                 r = 0;          g = 0;         b = 255;
347                 break;
348             case YELLOWCODE:
349                 r = 255;        g = 255;       b = 0;
350                 break;
351             default:
352                 r = 64;         g = 64;       b = 64;
353                 break;
354         }
355         #ifdef WIN32
356             *(out_pointer++) = b;
357             *(out_pointer++) = g;
358             *(out_pointer++) = r;
359         #else
360             // assert(0);
361             *(out_pointer++) = r;
362             *(out_pointer++) = g;
363             *(out_pointer++) = b;
364         #endif
365     };
366
367 };
368
369 };
370
371 #ifdef WIN32
372 //FIX optimize
373 void color_reconstruct(rle* input, UCHAR* output)
374 {
375     int x,y;
376     UCHAR r,g,b;
377     UCHAR* data = output;
378     UCHAR* impointer;
379     rle* rlepointer = input;
380
381
382
383     for (impointer=data;impointer<data+IMSIZE;impointer++) *impointer = 96;
384     //for all the rle runs on this row
385     while(rlepointer->tag) {
386         y = rlepointer->y;
387         switch(rlepointer->color) {
388             case FIELDCODE:
389                 r = 0;   g = 255;       b = 0;
390                 break;
391             case WALLCODE:
392                 r = 255;g = 255;       b = 255;
393                 break;
394             case BALLCODE:
395                 r = 255;g = 0;         b = 0;
396                 break;
397             case BLACKCODE:
398                 r = 0;   g = 0;         b = 0;
399                 break;
400             case BLUECODE:
401                 r = 0;   g = 0;         b = 255;
402                 break;
403             case YELLOWCODE:
404                 r = 255;g = 255;       b = 0;
405                 break;
406         }
407         for(x = rlepointer->begin; x < rlepointer->end + 1; x++) {
408             #ifdef WIN32
409                 *(data + 3 * x + 3 * IMWIDTH * y) = b;

```

```

410                                     *(data + 3 * x + 3 * IMWIDTH * y + 1) = g;
411                                     *(data + 3 * x + 3 * IMWIDTH * y + 2) = r;
412                                     #else
413                                     *(data + 3 * x + 3 * IMWIDTH * y) = r;
414                                     *(data + 3 * x + 3 * IMWIDTH * y + 1) = g;
415                                     *(data + 3 * x + 3 * IMWIDTH * y + 2) = b;
416                                     #endif
417                                     }
418                                     rlepointer++;
419                                     );
420                                     #endif
421                                     };
422                                     #endif
423
424                                     // Sub part of group
425                                     inline rle* matchrows(rle* rlepointer, rle* nextrow, blob** nextblob, int y) {
426                                     rle* rowbelow = nextrow;
427                                     int begin = rlepointer->begin; // optimize
428                                     int end = rlepointer->end; // optimize
429                                     int color = rlepointer->color;
430                                     blob* belowblob;
431                                     int abovearea;
432                                     blob* aboveblob = rlepointer->blobpointer;
433                                     int belowarea;
434
435                                     while(rowbelow->y==y+1) {
436                                     //Now check the for connectivity.
437                                     #ifdef USEPRINTF
438                                     printf("%X(%d,%d)", rowbelow->tag, rowbelow->y, rowbelow->begin);
439                                     printf("-%d ", rowbelow->end);
440                                     #endif
441                                     if (((rowbelow->begin >= begin)&&(rowbelow->begin <= end))||
442                                     ((rowbelow->end >= begin) && (rowbelow->end <= end))||
443                                     ((rowbelow->begin <= begin) && (rowbelow->end >= end))) &&
444                                     (rowbelow->color == color)) {
445                                     if (rowbelow->tag > rlepointer->tag) {
446                                     //Subsume below into above
447                                     //Create a blob for the above
448                                     if (!aboveblob) {
449                                     rlepointer->blobpointer = (*nextblob)++;
450                                     aboveblob = rlepointer->blobpointer;
451                                     aboveblob->color = rlepointer->color;
452                                     abovearea = end - begin + 1;
453                                     aboveblob->area = abovearea;
454                                     aboveblob->ymin = y;
455                                     aboveblob->ymin = y;
456                                     aboveblob->xmax = end;
457                                     aboveblob->xmin = begin;
458                                     }
459
460                                     if(rowbelow->blobpointer) {
461                                     // assert(aboveblob);
462                                     //row below is a blob get info
463                                     belowblob = rowbelow->blobpointer;
464                                     //this test could be better done with tag numbers
465                                     if (aboveblob != belowblob) {
466                                     aboveblob->area += belowblob->area;
467                                     belowblob->area = 0;
468                                     };
469                                     aboveblob->ymin = MAX(aboveblob->ymin, belowblob->ymin);
470                                     aboveblob->ymin = MIN(aboveblob->ymin, belowblob->ymin);
471                                     aboveblob->xmax = MAX(aboveblob->xmax, belowblob->xmax);
472                                     aboveblob->xmin = MIN(aboveblob->xmin, belowblob->xmin);
473                                     // assert(belowblob);
474                                     rowbelow->blobpointer=rlepointer->blobpointer;
475                                     rowbelow->tag = rlepointer->tag;
476                                     } else {
477                                     // assert(aboveblob);
478                                     //just a row so calc stats
479                                     belowarea = rowbelow->end-rowbelow->begin + 1;
480                                     aboveblob->area += belowarea;
481                                     //aboveblob->ymin not need updating?
482                                     // assert(aboveblob->ymin<y+1);
483                                     aboveblob->ymin = MAX(aboveblob->ymin, y+1); // unlikely
484                                     aboveblob->xmin = MIN(aboveblob->xmin, rowbelow->begin);
485                                     aboveblob->xmax = MAX(aboveblob->xmax, rowbelow->end);
486                                     rowbelow->blobpointer=aboveblob;
487                                     rowbelow->tag = rlepointer->tag;
488                                     };
489                                     } else {
490                                     //subsume above blob into below blob.
491                                     belowblob = rowbelow->blobpointer;
492                                     if (aboveblob) {
493                                     // rlepointer has a blob
494                                     if (aboveblob != belowblob) {
495                                     belowblob->area += aboveblob->area;
496                                     aboveblob->area = 0;
497                                     }
498
499                                     belowblob->ymin=MIN(belowblob->ymin, aboveblob->ymin);
500                                     belowblob->ymin=MAX(belowblob->ymin, aboveblob->ymin);
501                                     belowblob->xmin=MIN(belowblob->xmin, aboveblob->xmin);
502                                     belowblob->xmax=MAX(belowblob->xmax, aboveblob->xmax);
503
504                                     // assert(aboveblob);
505                                     rlepointer->blobpointer = rowbelow->blobpointer;
506                                     } else {
507                                     //rlepointer just a single row
508                                     // assert(belowblob);
509                                     abovearea = end - begin + 1;
510                                     belowblob->area+=abovearea;
511                                     belowblob->ymin=MIN(belowblob->ymin, y);
512                                     belowblob->ymin=MAX(belowblob->ymin, y);
513                                     belowblob->xmin=MIN(belowblob->xmin, begin);

```

```

514                                     belowblob->xmax=MAX(belowblob->xmax,end);
515                                     rlepointer->blobpointer = rowbelow->blobpointer;
516                                     };
517                                     rlepointer->tag = rowbelow->tag;
518     }
519     #ifdef USEPRINTF
520         printf("m%X ",rowbelow->tag);
521     #endif
522     }
523     rowbelow++; //and now is here
524 }
525     return rowbelow+1; //Should be pointing to the start of the next row;
526 }
527
528
529
530 blob* group(rle* input, blob* blobs)
531 {
532     int y = input->y;
533     rle* rlepointer = input;
534     rle* rowbelow;
535     rle* nextrow = input;
536     rle* temprow = nextrow;
537     blob* nextblob = blobs;
538     int begin, end;
539     blob* spacecheck = blobs+NUMBERBLOBS;
540     while(++temprow->y==y) {};
541
542     //for each row of the image
543     while(y < IMHEIGHT-1) { // HEIGHT MINUS ONE DON'T DO LAST ROW
544         rlepointer=nextrow;
545         y = rlepointer->y;
546         // now find next row after this one if it is not y++
547         // then skip this row too
548         while(++nextrow->y==y) {}; //could be optimized
549         //nextrow = temprow;
550         if (nextrow->y == y+1) {
551             //for all the rle runs on this row
552             while(rlepointer < nextrow) {
553                 #ifdef USEPRINTF
554                     begin = rlepointer->begin;
555                     end = rlepointer->end;
556                     printf("Tag %X{y %d - ", rlepointer->tag, rlepointer->y);
557                     printf("%d %d} [", begin, end);
558                 #endif
559                 //rowbelow = input + ENCODEDWIDTH * (y+1);
560                 rowbelow = nextrow;
561                 temprow = matchrows(rlepointer, nextrow, &nextblob, y);
562                 //assert(nextblob<spacecheck);
563                 //call match here returns last rle on the row
564                 //next run along the current row
565                 #ifdef USEPRINTF
566                     printf("]\n");
567                 #endif
568                 rlepointer++;
569                 if (rlepointer->tag == 0) return nextblob;
570             }
571         };
572     };
573     #ifdef USEPRINTF
574         printf("END GROUP\n");
575     #endif
576     return nextblob;
577 }
578
579 inline int xcentroid(blob* target) {
580     return (target->xmin + target->xmax)>>1;
581 };
582
583 inline int ycentroid(blob* target) {
584     return (target->ymin + target->ymax)>>1;
585 };
586
587 #define MAX_ERROR 4
588 object* analyse(blob* blobdata, blob* endblob, object* results,
589               object* pball, object* pygoal, object* pbgoal)
590 {
591     blob* pblob = blobdata;
592     blob* wall = 0;
593     blob* ball = 0;
594     blob* bgoal = 0;
595     blob* ygoal = 0;
596     int bluesize = MINGOALSIZE;
597     int yellowsize = MINGOALSIZE;
598     object* nextresult = results;
599     int expected_height, error;
600
601     int topline = 0;
602     int wallsize = MINWALLSIZE;
603     int ballsize = 0;
604
605     for (pblob = blobdata; pblob < endblob; pblob++) {
606         switch (pblob->color) {
607             case WALLCODE:
608                 if (pblob->area > wallsize) {
609                     wallsize = pblob->area;
610                     wall = pblob;
611                 };
612                 break;
613             }
614     }
615
616     if (wall) {
617         topline = wall->ymin - TOPWALLOFFSET;

```

```

618
619
620     #ifndef USEPRINTF
621         printf("walldimension %d x %d y - ", wall->xmin, wall->ymin);
622         printf(" %d x %d y\n", wall->xmax, wall->ymax);
623         printf("horizon line %d\n", topline);
624         printf("top of interest %d\n", topline);
625     #endif
626 };
627 // topline = 0;
628 for (pblob = blobdata; pblob < endblob; pblob++)
629     if ((pblob->area)&&(pblob->ymin>topline)) {
630         #ifndef USEPRINTF
631             printf("%d %d ", pblob->color, pblob->area);
632             printf("%d %d ", pblob->xmin, pblob->ymin);
633             printf("%d %d \n", pblob->xmax, pblob->ymax);
634         #endif
635
636         switch (pblob->color) {
637             case BALLCODE:
638                 expected_height = (pblob->xmax-pblob->xmin+1)*SQUARERATIO;
639                 error = expected_height - (pblob->ymin - pblob->ymin + 1);
640                 if ((error < 6) && (error > -3)) {
641                     if (pblob->area > ballsize) {
642                         ballsize = pblob->area;
643                         ball = pblob;
644                     }
645                 }
646             break;
647             case BLACKCODE:
648                 if (pblob->area > MINOBSTACLESIZE) {
649                     nextresult->distance = ycentroid(pblob);
650                     nextresult->bearing = xcentroid(pblob);
651                     nextresult->color = BLACKCODE;
652                     nextresult++;
653                 }
654             break;
655             case BLUECODE:
656                 if (pblob->area > bluesize) {
657                     bgoal = pblob;
658                     bluesize = pblob->area;
659                 }
660             break;
661             case YELLOWCODE:
662                 if (pblob->area > yellowsize) {
663                     ygoal = pblob;
664                     yellowsize = pblob->area;
665                 }
666             break;
667         }
668     }
669 };
670 if (ball) {
671     // pball->distance = angletable[ycentroid(ball)];
672     pball->distance = ycentroid(ball);
673     pball->bearing = xcentroid(ball);
674     pball->color = BALLCODE;
675 } else {
676     pball->distance = 0;
677     pball->bearing = -40;
678     pball->color = 0;
679 };
680 if (bgoal) {
681     pbgoal->distance = ycentroid(bgoal);
682     pbgoal->bearing = xcentroid(bgoal);
683     pball->color = BLUECODE;
684 } else {
685     pbgoal->distance = 0;
686     pbgoal->bearing = -40;
687     pball->color = 0;
688 }
689 if (ygoal) {
690     pygoal->distance = ycentroid(ygoal);
691     pygoal->bearing = xcentroid(ygoal);
692     pygoal->color = YELLOWCODE;
693 } else {
694     pygoal->distance = 0;
695     pygoal->bearing = -40;
696     pygoal->color = 0;
697 }
698
699     return (nextresult);
700 };
701
702
703 int findedgeline(int* profile) {
704     int y;
705     int sum;
706     int max=MINEDGELINE;
707     int oldmax = 0;
708     int besty=0;
709     for (y=0;y<IMHEIGHT; y++) {
710         sum = profile[y];
711         if(sum > oldmax) {
712             if (sum > max) {
713                 oldmax = max;
714                 max = sum;
715                 besty = y;
716             } else {
717                 oldmax = sum;
718             }
719         }
720     }
721     if (max - oldmax > MINEDGEDIFF)

```

```

722         return besty;
723     else
724         return 0;
725 };
726
727 /*void dilate22(UCHAR* the_image, UCHAR* out_image) {
728     UCHAR mask;
729     int x,y;
730
731     for (y = 1; y<IMHEIGHT - 1; y++)
732         for (x = 0; x<IMWIDTH - 1; x++) {
733             mask = *(the_image+x+IMWIDTH*y);
734             if (mask != 0) {
735                 *(out_image+(x+1)+IMWIDTH*y) = mask;
736                 *(out_image+(x+2)+IMWIDTH*y) = mask;
737                 *(out_image+(x+3)+IMWIDTH*y) = mask;
738             };
739         }
740 }*/
741
742 //
743 // uchar for histogram
744 void histogram_region(UCHAR* the_image, int sx, int sy, int ex, int ey, int* hist) {
745     int x, y;
746     UCHAR u,v;
747     for (y = sy; y <= ey; y++)
748         for (x = sx; x <= ex; x++) {
749             u = *(the_image + 3 * x + y * IMWIDTH * 3 + 1);
750             v = *(the_image + 3 * x + y * IMWIDTH * 3 + 2);
751             (*(hist + v + TABLEWIDTH * u))++;
752         }
753 }
754
755 // Threshold the the histogram into a UCHAR lookup table
756 void thresh(int* hist, int code, int level, UCHAR* table) {
757     int x, y;
758     for (x = 0; x < TABLEWIDTH; x++)
759         for (y = 0; y < TABLEWIDTH; y++)
760             if (*(hist + TABLEWIDTH * y + x) > level)
761                 *(table + TABLEWIDTH * y + x) = code;
762 };
763
764 void edgeprofile(UCHAR* target, int* profile) {
765     int y,x;
766     UCHAR* impointer = target;
767     for (y = 0; y < IMHEIGHT - 1; y++)
768         for (x = 0; x < IMWIDTH; x++) {
769             profile[y] += (*impointer==WALLCODE) && (*(impointer+IMWIDTH) != WALLCODE);
770             impointer++;
771         }
772 }
773
774 }

```
