Probabilistic Visual Recognition of Artificial Landmarks for Simultaneous Localization and Mapping

David Prasser and Gordon Wyeth
School of Information Technology and Electrical Engineering
University of Queensland
Brisbane, Australia
{prasserd, wyeth}@itee.uq.edu.au

Abstract

Probabilistic robotics, most often applied to the problem of simultaneous localisation and mapping (SLAM), requires measures of uncertainty to accompany observations of the environment. This paper describes how uncertainty can be characterised for a vision system that locates coloured landmarks in a typical laboratory environment. The paper describes a model of the uncertainty in segmentation, the internal camera model and the mounting of the camera on the robot. It explains the implementation of the system on a laboratory robot, and provides experimental results that show the coherence of the uncertainty model.

1 Introduction

Systems for simultaneous localisation and mapping (SLAM) use probabilistic approaches to account for the uncertainty in sensor data, for example [Thrun et al, 1998]. Robot perception systems are inherently uncertain, due to sensor limitations and noise. Probabilistic perception involves not only producing a good “guess” at an observation, but also providing the level of uncertainty associated with that “guess”.

This paper describes and characterises a visual perception system that can locate simple coloured landmarks in a cluttered laboratory environment, and can, most importantly, produce reasonable measures of the uncertainty in measurement of location of that landmark. Specifically the system measures range, bearing and an indicator of the type of landmark, as well as uncertainty measures for landmark. The system provides the measurements at 20 Hz, using 50% of an AMD 400 MHz processor. This allows CPU time for a SLAM algorithm to run in parallel with the landmark recognition system.

Importantly, the uncertainty in a measurement is characterised on a frame by frame basis. This provides the SLAM system with all of the necessary observation data for every frame. The paper shows that these instantaneous uncertainty measures are consistent with experimental results.

The following section describes the sources of uncertainty in the calculation of the landmark locations, and describes various models. Section 3 describes the robot platform that was used to test the models. Section 4 compares the model-based uncertainty with experimentally measured uncertainty to verify the suitability of the system as a sensor for SLAM.

2 Landmark Uncertainty

Coloured cylinders were chosen as the artificial landmarks because of their symmetry and simple shape when viewed from the side. A typical image of a cylinder in the test environment is shown in Figure 1. The visual recognition system identifies the cylinder in the scene, and reports the bearing and distance to the cylinder as well as the uncertainty in those measurements. Identification and measurement is performed by three components of the system: (i) a colour segmentation and blob finding algorithm; (ii) an estimator of segmentation noise; and (iii) a method for transforming from image coordinates to robot centred coordinates that accounts for uncertainties in the measurement of camera parameters.

Figure 1: An example image showing the coloured cylinder used as an artificial landmark.

2.1 Noise

One of the sources of error in measurement of centroid location is the noise in the image. Typically this error appears as the random changes in the location of the centroid over a series of frames. One way of assessing this noise is to measure the variance of the centroid’s location over a series of frames, but this causes a significant delay in the availability of data for mapping or localisation of a moving robot. This section investigates methods to estimate the variance in the visual centre of mass by examining the properties of the segmented objects to be classified and located.

The method used can be developed for rectangular targets by using the following line of reasoning for variance of the horizontal centre of mass.

The centre of mass of an object is found by finding the sum of moments ΣM and dividing by the area n.

$$m_x = \frac{\Sigma M}{n} = \frac{\sum_{i,j} iX_{i,j}}{\sum_{i,j} X_{i,j}} \tag{1}$$

where m_x is the centre of mass about the x axis and $X_{ij} \in \{0,1\}$ is the value of pixel at (i,j). If there is a perfect colour detection of a rectangle with no noise then its centroid will not move or change over a succession of frames. A random noise in the colour segmented image can be modelled as a probability ρ that a pixel is correctly segmented. The probability model for correct segmentation is a Bernoulli trial with probability ρ_{ij}, and therefore variance $\rho_{ij}(1-\rho_{ij})$. Propagation of variance can then be used to find the variance of the centre of mass, as below. The working is based on the x axis, but the same reasoning applies to the y axis.

$$\mathrm{VAR}(m_x) = \sum_{i,j} \left(\frac{\partial m_x}{\partial X_{i,j}} \right)^2 \mathrm{VAR}(X_{i,j}) \tag{2}$$

$$\frac{\partial m_x}{\partial X_{i,j}} = \frac{i}{\sum_{i,j} X_{i,j}} - \frac{\sum_{i,j} iX_{i,j}}{\left(\sum_{i,j} X_{i,j} \right)^2} = \frac{i}{n} - \frac{\sum_{i,j} iX_{i,j}}{n^2}$$

$$\frac{\partial m_x}{\partial X_{i,j}} = \frac{i - m_x}{n}$$

$$\mathrm{VAR}(m_x) = \sum_{i,j} \left(\frac{i - m_x}{n} \right)^2 \rho_{i,j} \left(1 - \rho_{i,j}\right) \tag{3}$$

In order to find the variance in the centre of mass, therefore, it is necessary to estimate the probability that a pixel is correctly segmented. The following subsections look at three approaches to estimating the noise in segmentation from analysis of the image.

2.1.1 Uniformly Distributed Noise

Based on an assumption of uniformly distributed noise, $\rho_{i,j}$ may be estimated by dividing the total number of detected pixels by the number of pixels that should have been detected. In the case of the rectangle expected from a cylindrical marker, this is the width multiplied by the height. As this method assumes that the noise is uniformly distributed around the entire object so $\rho_{i,j}$ can be replaced by ρ. Consequently, based on (3),

$$\mathrm{VAR}(m_x) = h \sum_{i} \left(\frac{i - m_x}{n} \right)^2 \frac{n}{hw} \left(1 - \frac{n}{hw} \right) \tag{4}$$

where h is the height of the object, w is the width and n is the total number of detected pixels.

2.1.2 Column / Row Distributed Noise

The previous estimator applies uniform noise to all areas of the segmented object whereas segmentation noise is typically more prevalent near the perimeter of an object. A better method may be to consider each column to have its own noise ρ_i, which may be estimated by dividing the total number of detected pixels by the number of pixels that should be in that column. This means that a column in which each pixel is on will have no contribution to the variance (because the variance of a binomial sum with $\rho = 1$ is zero). The largest contributions to the variance in the number of moments will come from columns with values of ρ_i around 0.5 which is typically only the left and rightmost columns of the image. Based on the assumption of column distributed noise, the variance can be estimated as:

$$\mathrm{VAR}(m_x) = \sum_{i} \left(\frac{i - m_x}{n} \right)^2 n_i \left(1 - \frac{n_i}{h} \right) \tag{5}$$

The same argument can be applied to row based noise estimation.

2.1.3 Locally Distributed Noise

Another approach is to consider each pixel to have its own noise value ρ_{xy}. The local noise value at point (x,y) can be approximated as the number of detected pixels in a small window divided by the total area of the window. So for a three by three window:

$$\mathrm{VAR}(m_x) = \sum_{i,j} \left(\frac{i - m_x}{n} \right)^2 \frac{N(i,j)}{9} \left(1 - \frac{N(i,j)}{9} \right) \tag{6}$$

where $N(i,j)$ is the number of segmented pixels in the three by three region centred at (i,j).

2.2 Landmark Confidence

It is important to have some measure of whether or not an observed landmark is an actual landmark or some other object that has been accidentally segmented. From the measured distance and the physical dimensions of the landmark the expected height h_{est} and width w_{est} in pixels can be calculated. The ratio of the expected width and the width of the landmark in the image should be near unity if the object has the physical dimensions of a landmark. Inverting this number if it is greater than one gives a confidence measure for the width, ranging from zero to unity. A similar ratio based measure can be found for the height. The height and width confidence measures multiplied together will give an overall confidence measure varying from no confidence at zero to complete confidence at unity:

$$c_{LM} = r\left(\frac{w}{w_{est}} \right) r\left(\frac{h}{h_{est}} \right) \tag{7}$$

$$r(x) = \begin{cases} x, x \le 1 \\ 1/x, x > 1 \end{cases} \tag{8}$$

2.3 Camera Model

A model of the camera is developed for both normalisation of the image to account for imperfections in the camera, and for the coordinate transform to move from pixel centred to robot centred coordinates. The procedure used for calibrating the camera is based on [Zhang 1999] and uses software developed by [Bouguet 1998]. The procedure uses a set of calibration images to calculate the cameras focal length in pixels, $\mathbf{f_l}$, and the location of the lenses centre, $\mathbf{p_p}$. The software developed by Bouguet is attractive in that it provides estimates of the error in all of the parameters it estimates.

2.3.1 Camera Normalization

The vision system accounts for imperfections in the camera by correcting the pixel coordinates of blob centroids to normalized 'pinhole' coordinates using a standard normalization method [Zhang, 1999]. The camera calibration data is used to convert the pixel coordinates to camera relative angles. At this stage the uncertainty in the principle point and focal length is introduced [Heikkilä 1997]:

$$\mathbf{x_n}' = \frac{\mathbf{x_p} - \mathbf{p_p}}{\mathbf{f_l}} \tag{9}$$

$$\sigma_{\mathbf{x_n}'} \approx \left(\frac{\partial \mathbf{x_n}'}{\partial \mathbf{f_l}}\right)^2 \sigma_{\mathbf{f_l}} + \left(\frac{\partial \mathbf{x_n}'}{\partial \mathbf{p_p}}\right)^2 \sigma_{\mathbf{p_p}} + \left(\frac{\partial \mathbf{x_n}'}{\partial \mathbf{x_p}}\right)^2 \sigma_{\mathbf{x_p}} \tag{10}$$

where $\mathbf{x_p}$ is the pixel coordinates of the centre of mass. The uncertainties in the camera parameters ($\sigma_{\mathbf{xp}}$ and $\sigma_{\mathbf{fl}}$) are reported by the calibration software, while $\sigma_{\mathbf{xp}}$ is the output of the noise estimator described above. After this, the distorted normalized coordinates $\mathbf{x_n}$' can be converted to the undistorted normalized coordinates $\mathbf{x_n}$ using a sixth order polynomial [Heikkilä 1997]. The uncertainty of this stage is not considered and it is assumed that because $\mathbf{x_n}' \approx \mathbf{x_n}$ then the uncertainty of $\mathbf{x_n}$ is the same as that of $\mathbf{x_n}$'.

2.3.2 Coordinate Transforms

The normalized camera coordinates can be converted into a pair of angles from the cameras axis. The vertical angle, ω_c, is increased by the camera's tilt angle, α_{tilt}. The vertical angle can then be used to calculate the distance to the landmark, d_c, assuming a constant height difference between the landmark and the camera h. The bearing to the landmark, θ_c, is derived from the centre of mass of the landmark in the x axis and increased by the cameras yaw angle, α_{yaw}.

$$\omega_c = \tan^{-1}(y_n) + \alpha_{tilt} \tag{11}$$

$$d_c = \tan(\omega_c)h \tag{12}$$

$$\theta_c = \tan^{-1}(x_n) \tag{13}$$

Propagation of variance can be used to calculate the uncertainty in range and bearing as well.

$$\sigma_{\omega_c}^2 = \left(\frac{\partial \omega_c}{\partial y_n}\right)^2 \sigma_{y_n}^2 + \sigma_{\alpha_{tilt}}^2 \tag{14}$$

$$\sigma_{d_c}^2 = \left(\frac{\partial d_c}{\partial \omega_c}\right)^2 \sigma_{\omega_c}^2 + \left(\frac{\partial d_c}{\partial h}\right)^2 \sigma_h^2 \tag{15}$$

$$\sigma_{\theta_c}^2 = \left(\frac{\partial \theta_c}{\partial x_n}\right)^2 \sigma_{x_n}^2 + \sigma_{\alpha_{yaw}}^2 \tag{16}$$

This gives the range and bearing, along with the uncertainties, relative to the *camera*. However, on the robot, the camera is mounted forward of the robot's centre which is the usual reference point for navigation. From the cosine rule, the transformations can be developed to move range and bearing from the camera to the robot's centre. Again the uncertainty is also calculated:

$$d_r = \sqrt{d_c^2 + s^2 + 2ds\cos\theta_c} \tag{17}$$

$$\theta_r = \tan^{-1}\left(\frac{d_c \sin\theta_c}{s + d_c \cos\theta_c}\right) \tag{18}$$

$$\sigma_{dr}^2 = \left(\frac{\partial d_r}{\partial d_c}\right)^2 \sigma_{dc}^2 + \left(\frac{\partial d_r}{\partial \theta_c}\right)^2 \sigma_{\theta c}^2 + \left(\frac{\partial d_r}{\partial s}\right)^2 \sigma_s^2 \tag{19}$$

$$\sigma_{\theta_r}^2 = \left(\frac{\partial \theta_r}{\partial d_c}\right)^2 \sigma_{d_c}^2 + \left(\frac{\partial \theta_r}{\partial \theta_c}\right)^2 \sigma_{\theta_c}^2 + \left(\frac{\partial \theta_r}{\partial s}\right)^2 \sigma_s^2 \tag{20}$$

where d_r and θ_r are the range and bearing to the landmark from the robot centre and s is the separation distance between the camera and robot centre. In this way each reasonably sized coloured object gives a range and bearing pair with uncertainties, colour type and a value which indicates the likelihood it is a landmark.

3 Experimental System

The noise estimator was tested on the robot that is the intended vehicle for the SLAM research. A simple colour detection and segmentation system was developed to test it out properly.

3.1 Test Platform

The system was tested on a Pioneer DX2E robot from ActivMedia. The robot is equipped with a single pan / tilt / zoom camera with PAL video output, a frame grabber, and an AMD 400 MHz computer running Windows™ XP. The generic nature of the computer system and input means that the whole software package can be quickly adapted to running on other robots or desktop systems for evaluation. As the processor in a SLAM system would be shared between vision, navigation and mapping components the system should make as little use of the processor as possible while still operating at a rate acceptable to the navigation module.

The current landmarks are cylinders made from coloured paper. Several colours were tried: orange and green having the best results and low rates of natural occurrence in the test environment. The location of the camera on the top of the robot at a fixed height and the assumption that landmarks lie on a uniformly flat floor makes the calculation of range straightforward.

3.2 Colour Segmentation

The system uses the YUV / YC_bC_r output format found on most frame-grabbers. Colour detection operates with a lookup table approach. The lookup table has two dimensions which correspond to the U and V dimensions in colour space. Each cell in the table contains a result code which represents either one of the colours of interest or zero for colours that do not need to be detected. Colour detection can then be accomplished by looking up each pixels U and V values and reading the result from the table.

The U and V components are not sufficient for robust colour segmentation, as the brightness (Y) is also needed to distinguish between say yellow and white. To avoid errors caused by the incorrect identification of bright or dark colours an extra lookup vector for the Y dimension is included. The result from this lookup vector is logically ANDed with the result from the main lookup table. Zeros in this vector force the output of the colour detection to zero for extreme values of Y.

After colour detection, following an approach similar to [Bruce 2000], the data is converted to a run length encoded format (RLE). Colour detection and run length encoding consume the majority of the processor time. Later stages do not need to be particularly fast as they are operating on much less data than these first two stages. The RLE stage is also used to perform a two dimensional binary opening on the image by rejecting RLE groups less than three pixels in length

After RLE compression, the horizontal runs representing orange and green pixels are examined for four connectivity and grouped into blobs. Blobs of the same colour that are nearby are then grouped together into one larger blob. The output of this segmentation scheme is a set of blobs with bounding rectangles and pixel counts. Associated with these blobs are the RLE runs so the shape of a blob can be recovered by examining the RLE elements.

3.3 Colour Lookup Table Calibration

The lookup table is a 256 × 256 array with axes that represent the U and V coordinates. Each cell then corresponds to a particular combination of red and blue chrominance. Calibration then becomes a problem of determining which cells in the table correspond to different coloured objects. The solution is for a human operator to identify a uniformly coloured training region in a sample image. A two dimensional histogram is computed using the U and V values of the pixels in the area that the operator selected. Any non-zero bin in the histogram corresponds to a point in the lookup table that should be set to represent the colour of the object in the sample image.

Several improvements have been made to this approach. The first is to threshold the histogram before using it to add data to the lookup table. This reduces the chance of errors in calibration caused by including non-target pixels in the training region. Using this technique the operator can build the complete lookup table by repeating the process described above several times to account for minor changes in illumination. The number of sample images can be reduced by increasing the size of the area that is filled in on the lookup table. The approach currently used in the software is to fill in a rectangular region of the lookup table around the histogrammed area but extended by a small amount.

4 Results

The results illustrate the performance of the noise estimation, landmark confidence and camera modelling techniques. These results are combined in experimental tests on the Pioneer robot that illustrate the relative importance of the various sources of uncertainty in the landmark measurement process.

4.1 Noise Estimators

The noise estimators can be compared to the variance of the centroid position when it is measured over a series of frames. Table 1 shows the measured variance of the vertical centroid over four hundred frames at four different ranges. Also shown are the average outputs of each estimator. The noise estimators are not particularly accurate, but it should be remembered that these estimates come from a single frame with no prior knowledge about the noise function.

Table 1: Variances (in pixels^2) of the measurement of the vertical centroid. Comparison of the uniform, column and local noise models against actual noise measurements over four trials using 400 images.

Trial #	Actual	Uniform	Column	Local
1	0.023	0.037	0.035	0.016
2	0.018	0.041	0.039	0.023
3	0.009	0.033	0.034	0.010
4	0.013	0.028	0.024	0.015

Clearly the local estimator provides the closest estimate of the noise in each frame. This is because the other two methods consider any pixel in the rectangular region that is not detected to be noise. The local method will declare any pixel which is “off” and surrounded by eight other “off” pixels to have a noise value $\rho_{i,j}$ of 0 and consequently a variance of zero. In other words the local method has implicit support for slightly non-rectangular shapes while in the other two methods non-rectangular shapes cause larger amounts of noise.

When the noise estimators are used on an object that has been distorted by occlusion the uniform and row / column based estimators report an increase in the noise, while the local system still reports the actual variance. In the case of occlusion it would be preferable for the estimated variance to increase representing the lower confidence in the output.

Table 2: Variances (in pixels^2) of the measurement of the vertical centroid for a partially occluded landmark. Comparison of the uniform, column and local noise models against actual noise measurements using 400 images.

Actual	Uniform	Column	Local
0.0267	0.1103	0.0877	0.0198

The estimates of the horizontal centroid show another problem (Table 3). The measured variance for trial 1 is significantly larger than expected or predicted. The cause of this is that while measuring this data the width of the target was alternating between two values. A histogram of the centroid position reveals bimodal distribution, whereas the estimators only produce estimates for normal distributions as they all assume that the objects width does not change.

Table 3: Variances (in pixels^2 × 10^{-4}) of the measurement of the horizontal centroid. Comparison of the uniform, column and local noise models against actual noise measurements over four trials using 400 images.

Trial #	Actual	Uniform	Column	Local
1	323	22	10	15
2	5.1	22	9.2	16
3	8.9	21	17	11
4	3.1	17	8.0	15

The noise estimators, and particularly the local noise estimator, produce reasonable values for uncertainty in the measurement of centroid location of the artificial landmarks. The local noise estimator comes the closest to experimental values, but suffers from the significant computational expense of a nine point calculation for every pixel in the blob.

4.2 Landmark Confidence

Landmarks tended to have high confidence values attached to them 0.80 – 0.95, while random non landmark objects that were found by the colour detection process had values ranging from 0.10 to 0.30 usually. In practice, these values could reasonably be used as is; or perhaps used to derive a more meaningful confidence value where landmark identification is in doubt.

4.3 Intrinsic Camera Model

The camera model based on [Heikkilä 1997] and [Bouguet 1998] produces the results shown in Table 4 for the camera mounted on the Pioneer robot. The estimated errors in the focal length and principal point are passed into the vision system and included when calculating the uncertainty in position of a landmark.

Table 4: The camera model and uncertainty found for the Pioneer camera setup.

Parameter	Value	Uncertainty
Horizontal Principal Point	386.3	0.3216
Vertical Principal Point	278.8	0.2826
Horizontal Focal Length	875.2	0.4565
Vertical Focal Length	872.4	0.4399

4.4 Extrinsic Calibration

In order to calculate the distance and bearing to the landmark, the extrinsic parameters must be determined. The camera itself has three significant external parameters: (i) the height of the focal point, (ii) the angle the camera makes with the ground plane, and (iii) the angle between the camera and the forward direction of the robot. These parameters can be found by using the intrinsic parameters of the camera to calculate the rotation and translation vectors from a calibration target which is located on the ground. This process was repeated ten times to acquire a range for the angle of depression and height of the camera. Each time the camera's position was reset – mimicking the process the camera goes through during power on. The calculated variance of the height and angle of depression was used as the overall variance as illustrated in Table 5. The camera is assumed to be pointing directly forward and the measured error caused by this assumption is used as the uncertainty in the camera yaw angle.

Table 5: Extrinsic parameters of the camera as determined by an experiment of 10 camera start-up sequences.

Parameter	Value	σ
Camera Tilt Angle (°)	1.656	0.575
Camera Yaw Angle(°)	0.000	0.443
Camera Height (m)	0.326	3.04×10^{-4}

The height of the robot and the displacement from the camera to the robot's centre of rotation can be measured directly from the robot. The robots height is dependent upon tire pressure but can always be measured to ±0.5 mm. The displacement from the camera reference point to the robot centre of rotation is more difficult to measure so a 95% confidence of ±5 mm is used.

4.5 Accuracy of System

The uncertainty in the range measurements increases as the distance to the target becomes larger. The uncertainty also increases as the target approaches the left or right edges of the field of view. This is caused by the translation from camera centred to robot centred coordinates becoming more sensitive to the angular error as the angle increases.

Table 6: Measurements of the distance (in m) from the robot centre to a landmark. Comparison of the distance measured manually with the value calculated from the image. The calculated variance is shown, and compared to the error in measurement.

Actual	Calculated	σ	Error / σ
2.97	2.86	0.34	0.33
2.89	2.84	0.33	0.15
1.83	1.92	0.14	0.63
1.62	1.68	0.11	0.54
1.45	1.44	0.079	0.074
1.17	1.19	0.052	0.51

The angular uncertainty is approximately the same as the original uncertainty in camera heading as described section 4.4.

Table 7: Measurements of the bearing (in degrees) from the robot centre to a landmark. Comparison of the angle measured manually with the value calculated from the image. The calculated variance is shown, and compared to the error in measurement.

Actual	Calculated	σ	Error / σ
0	0.17	0.41	0.40
8.1	8.42	0.41	0.71
11.4	11.9	0.40	1.69
18.4	18.5	0.42	0.036

4.6 Significance of Error Sources

The sensitivity of the outputs to each of the various sources of error can be measured by setting the uncertainty of all of the other parameters to zero and examining the output uncertainty. Table 8 shows the percentages contributed by each source of error, as a contribution to the total variance in the measurement of angle and distance.

Table 8: Percentages contributed by each source of error, as a contribution to the total variance.

Source of Error	Angle	Distance
Estimated Noise	> 1%	> 1%
Principal Point	> 1%	> 1%
Focal Length	> 1%	>> 1%
Camera Height	>> 1%	>> 1%
Camera Tilt Angle	3 %	99%
Camera Yaw Angle	96 %	> 1%
Camera Position	> 1%	> 1%
Total Variance	0.0012	0.0020

These results indicate which components of the error are most significant. Clearly the dominant source is the uncertainty in cameras orientation. Obtaining a more accurate result for these parameters would greatly reduce the overall uncertainty of the system and possibly highlight other sources of error.

Examining the sensitivity of each source of error by setting $\sigma^2 = 1$ for each error at a time gives the results shown in Table 9. The sensitivity is dependent upon the measured range and bearing which were 1.36 m and 13.2° respectively. These results show that the system is most sensitive to the external parameters of the camera and then to the pixel location of the landmark. This indicates that any operations involving movement of the camera must be carefully modelled for their potential impact on the uncertainty in the measurement of landmark location.

Table 9: Sensitivity analysis of the sources of error.

Source of Error	Angle (deg^2)	Distance (m^2)
Camera Tilt Angle	1.46×10^{-2}	2.08×10^{-2}
Camera Yaw Angle	0.80	3.62×10^{-7}
Camera Height	30.4	43.22
Camera Position	93.4	0.95
Principal Point	6.0×10^{-5}	3.06×10^{-3}
Focal Length	2.08×10^{-4}	1.35×10^{-6}
Segmentation Noise	1.2×10^{-3}	1.0×10^{-4}

5 Conclusion

SLAM systems require a measure of the uncertainty inherent in a sensor reading. The visual perception system described in this paper provides not only useful landmark location information, but also reasonable measures of the uncertainty in the location information. The paper illustrates the various sources of uncertainty in the measurement and how they can be characterised on a frame-by-frame basis. Implementation on a laboratory robot has been described, and results presented that show (i) that segmentation noise is best characterised using a local noise model, (ii) that the uncertainty in the measurement of distance and bearing to a landmark is well characterised, (iii) that greater confidence in calibration parameters is needed, and (iv) that the system is most sensitive to the external parameters of the camera.

References

[Bouguet 1998] J.Y. Bouguet and P. Perona, “Closed form camera calibration in dual–space geometry”, Technical Note, 1998.

[Bruce 2000] J. Bruce, T. Balch and M. Veloso, “Fast and Inexpensive Color Image Segmentation for Interactive Robots”, *Proc. of International Conference on Intelligent Robots and Systems*, vol. 3, pp. 2061-2066, 2000.

[Heikkilä 1997] J. Heikkilä and O. Silvén, “A four-step camera calibration procedure with implicit image correction”, *Proc. IEEE conference on Computer Vision and Pattern Recognition*, pp. 1106-1112, 1997.

[Thrun et al, 1998] S. Thrun, D. Fox and W. Burgard, “Probabilistic Mapping of an Environment by a Mobile Robot”, *Proc. IEEE of International Conference on Robotics and Automation,* vol 2., pp. 1546-1551, 1998.

[Zhang 1999] Z. Zhang, “Flexible camera calibration by viewing a plane from unknown orientations”, *Proc. IEEE International Conference on Computer Vision*, vol. 1, pp. 666-673, 1999.