Dynamic Lens Compensation for Active Color Imaging and Constant Magnification Focusing

Reg G. Willson and Steven A. Shafer

CMU-RI-TR-91-26

The Robotics Institute,
Carnegie Mellon University
Pittsburgh, Pennsylvania 15213

November 1991

© 1991 Carnegie Mellon University

This research was sponsored by the Avionics Lab, Wright Research and Development Center, Aeronautical Systems Division (AFSC), U. S. Air Force, Wright-Patterson AFB, OH 45433-6543 under Contract F33615-90-C-1465, ARPA Order No. 7597. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Government.

Contents

1 Introduction 1

2 Active Color Imaging 2
- 2.1 Chromatic Aberration 2
- 2.2 Compensation Approach 6
- 2.3 Experimental Results 9
- 2.4 Conclusions 13

3 Constant Magnification Focusing 14
- 3.1 Focus Magnification 14
- 3.2 Compensation Approach 18
- 3.3 Experimental Results 18
- 3.4 Conclusions 24

4 Summary 25

5 Acknowledgments 26

6 Appendix A - Chromatic aberration measurement 27

7 Appendix B - Relative image magnification model 34

8 Appendix C - Criterion function peak position detection 41

List of Figures

1 Chromatic aberration in a thin lens 3
2 Checkerboard test target 4
3 Image of grid pattern 4
4 Image profile near center 5
5 Image profile near edge 5
6 Active Color Imaging calibration target 8
7 Chromatic aberration correction process 8
8 Full field error - uncompensated lens 10
9 Full field error - compensated lens 11
10 Color histogram - uncompensated image 12
11 Color histogram - compensated image 12
12 Focus magnification 14
13 Relative magnification versus f_d and f_l motor positions 15
14 Focus magnification - feature clipping 16
15 Focus magnification - perimeter scaling 17
16 Focus magnification - gradient scaling 17
17 Box and bar focus targets 18
18 Lens settings for range from focus 20
19 Position of criterion function peak - Wide aperture (f/# $\approx$ 3.3) 21
20 Position of criterion function peak - Narrow aperture (f/# $\approx$ 22) 22
21 Criterion function peak position variation with aperture 23
22 Checkerboard image 28
23 Thresholded edge image 28
24 Test mask for horizontal reference edges 29
25 Edge image with highlighted reference edges 29
26 Blue-Red vertical misregistration 31
27 Blue-Green vertical misregistration 31
28 Blue-Red horizontal misregistration 32
29 Blue-Green horizontal misregistration 32
30 Blue-Red total misregistration 33
31 Blue-Green total misregistration 33
32 Calibration target for relative image magnification model 35
33 Relative image magnification model 37
34 Percentage error in relative image magnification model 38
35 Constant Magnification Focusing isomagnification contour 40
36 Peak detection in criterion function data 45

Abstract

Vision researchers base their work on assumptions that lenses behave ideally. In practice the image formation process is complex and camera lenses have many non-ideal behaviors that can cause problems in simple vision tasks like color image analysis and range from focus. For example, unlike the ideal lens the refractive power of a real lens is actually a function of wavelength. This means that images of a scene taken under different wavelengths will have slightly different focus and magnification. The resulting misregistration between the images when they are superimposed to form a composite color image can cause significant problems in color image analysis. Another property of real lenses is that focusing the lens also changes its magnification. In range from focus these changes in image magnification can bias the position of a sharpness criterion function's peak leading to errors in the range estimates.

By precisely controlling the lens during imaging it is possible to make lenses behave more ideally. We have developed a procedure we call Active Color Imaging in which lens magnification and focus are precisely compensated to produce color images that are over 10 times better registered than uncompensated color images. We have also developed a procedure we call Constant Magnification Focusing which allows a user to focus a lens without changing the image magnification, eliminating focus magnification induced errors in range from focus tasks.

1 Introduction

The ideal lens for machine vision would probably behave like a pinhole camera - with perfect perspective projection from every point in the object space down to the image plane and with every point in object space in perfect focus. Such a lens is not practical, first, because it would admit insufficient light for most imaging sensors, and second, because diffraction effects at the camera's pinhole would limit the image resolution that could be obtained. In real machine vision systems the pinholes are replaced by lenses to gather light from many points in the object space and project it down to the imaging plane. Such lenses exhibit many non-ideal behaviors that result from the fundamental physics involved in the optics and from manufacturing considerations. For the most part vision researchers make simplifying assumptions about the image formation process and ignore most non-ideal lens behaviors. In this report we examine how two low level vision tasks, color image analysis and focus ranging, are affected by two non-ideal lens behaviors. Subsequently for both cases we show how precise control of the lens during imaging can compensate for the non-ideal lens behaviors and provide significantly better results.

In color image analysis the information contained in three color bands is used to determine properties of the scene being imaged. Implicit in the color image analysis is the assumption that the per pixel information in each band corresponds to the same point, region, or volume in object space. In this report we show how chromatic aberration in lenses can result in significant misregistration between the bands of color images, and how this misregistration can cause problems for color image analysis. We then show how precise control of the lens' focal length, focus distance and XY position during imaging can reduce the misregistration between the color image bands by an order of magnitude. We call this method Active Color Imaging (ACI). As part of this work we present a method for measuring the amount of lateral chromatic aberration in a camera system.

In range from focus the distance to a feature in the camera's field of view is estimated by moving the camera lens through a series of focus positions while determining the sharpness of the feature's image at each position. The feature's sharpness is measured using a criterion function evaluated over the image. Given the focus position of the lens where the criterion function peaks, the distance between the camera and the feature in the image can be determined using a calibrated lens model. Implicit in the range from focus procedure is the assumption that the criterion function peaks at the point of sharpest focus. In this report we show how focus magnification can bias the position of the criterion function peak, yielding incorrect results in the focus ranging task. We then show how nulling out focus magnification by the precise control of the lens' focal length can eliminate the bias in the position of the criterion function's peak. We call this approach Constant Magnification Focusing (CMF). As part of this work we present a method to accurately and reliably estimate the position of unimodal criterion function peaks in the presence of noise.

This work was originally presented in Willson [17] at the 1991 IEEE International Conference on Robotics and Automation.

2 Active Color Imaging

Color image analysis uses the information contained in three spectral bands to determine properties of the scene being imaged. Implicit in any color image analysis is the assumption that the per pixel information in each band corresponds to the same point, region, or volume in object space. As we will demonstrate in the following section, this is not always so.

The starting point for any type of color image analysis is the acquisition of color images. To simplify the image alignment problem between the color bands virtually all color imaging processes use a single lens to image all of the bands. Variations between imaging processes occur in the methods used to digitize the bands. Three basic approaches are used. The simplest approach, used in most consumer color cameras, is to use a single sensor with every third pixel overlaid with one of the three bandpass filters. Spatial interpolation is then used fill in the gaps in the image bands.

The second approach, used in broadcast quality color cameras, is to split the image from the lens into three and then to project each of the three images through bandpass filters onto separate sensors.

The third approach, used in still color imaging, is to place three bandpass filters in the optical path of the camera system one after another and take the three images with the same sensor. This third approach is the one that we use in the Calibrated Imaging Lab, and is used for most machine vision research in color physics because of the superior registration between images and the improved SNR that is possible for each band.

The imaging system used for this work consists of a solid state video camera, a computer controlled filter wheel, a computer controlled lens and a computer controlled six degree of freedom camera platform. The camera is a General Imaging MOS-5300 connected to a Matrox frame grabber. The filter wheel houses six filters including Wratten #25 (red), 58 (green) and 47B (blue). The lens is a Cosmicar CCTV grade zoom lens with a focal length range of 12.5 to 75mm, a minimum focussed distance of 1.2m, and a maximum aperture ratio of 1:1.8. Lens automation is provided by three micro stepping motors providing 3900 steps of resolution for focus distance, 4000 steps for focal length and 2700 steps for aperture. The camera platform has a positioning accuracy of 0.025mm.

2.1 Chromatic Aberration

Chromatic aberration exists in camera lenses because the index of refraction of optical components varies as a function of wavelength. This difference in refractive index causes the different wavelengths of light to be refracted or bent different amounts by the components of the lens. For example, given a simple uncorrected thin lens with incident white light rays, the blue components of the incident rays will be brought to a focus closer to the lens than the red components, as shown in Figure 1. In [5] Funt and Ho make use of this property to extract spectral information to address the problem of color constancy.

Chromatic aberration is an intrinsic property of a camera lens. It can be partially compensated for in the lens optics by using pairs of lens elements with offsetting dispersion factors. Unfortunately such compensation is usually done for only two wavelengths (red and blue) and then only at two points in the image field: at the optical axis and at some specified radial distance from the optical axis [14]. For a more detailed discussion of lenses and chromatic aberration the reader is referred to [13].

To measure chromatic aberration we use black on white checker board targets like the target shown in Figure 2. The positions of the vertical and horizontal black to white step

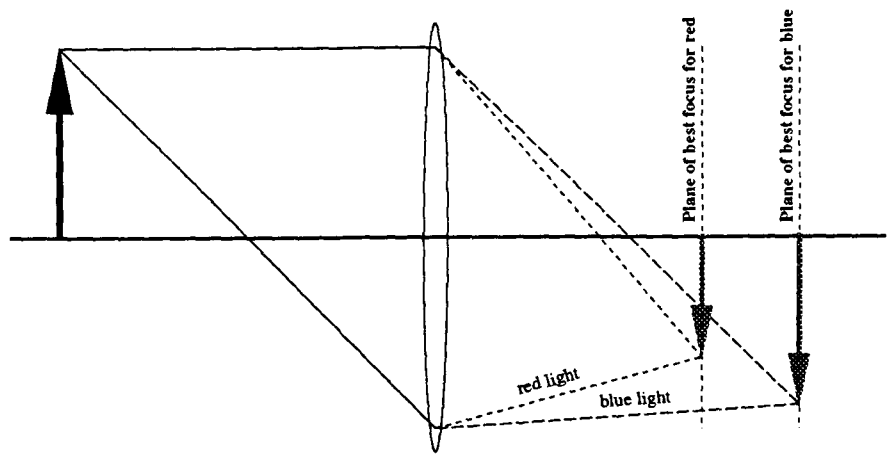


Figure 1: Chromatic aberration in a thin lens

edges, measured in each of the bands, provide a fairly precise measure of the lateral chromatic aberration in the camera system. Appendix A describes the measurement procedure in detail.

Effects of Chromatic Aberration on Images

Chromatic aberration has three effects on color images: magnification differences between bands, focus differences between bands, and decentering between the bands [11]. The first two of these effects are illustrated in Figure 1. In this figure the difference in the heights (magnification) of the two images is called lateral chromatic aberration. The difference in the positions of the plane of focus for the two images is called longitudinal chromatic aberration. These effects can also be seen by examining the color image of the black on white grid shown in Figure 3. Figure 4 is a graph of the red, green and blue pixel values versus pixel number for the scanline located in section 1 of the grid image in Figure 3. The vertical black line of the grid that falls within section 1 is near the optical axis of the lens and shows no noticeable misregistration between the red, green and blue color bands. Figure 5 is a graph of the red, green and blue pixels versus pixel number for section 2 of the same scanline. The vertical black line being examined in this figure is relatively far from the optical axis of the lens and clearly shows the misregistration of the three color bands due to lateral chromatic aberration. The slightly broader width of the red valley in both figures shows the defocus in the red image caused by longitudinal chromatic aberration (the lens was initially focussed with blue light). With our camera system the difference in the magnification between the uncorrected red and blue images is on the order of 0.5%, or 1.2 pixels near the outer edges of images that are 512 pixels wide (ie. a 1.2 pixel displacement in one direction at the left edge and a 1.2 pixel displacement in the other direction at the right edge).

While magnification and defocus result from intrinsic properties of the optical components, the third chromatic effect, decentering between the image bands, is likely the result of misalignment in the optical components of the lens. This causes light rays of different wavelengths to take slightly different optical paths through the lens. The decentering ef-

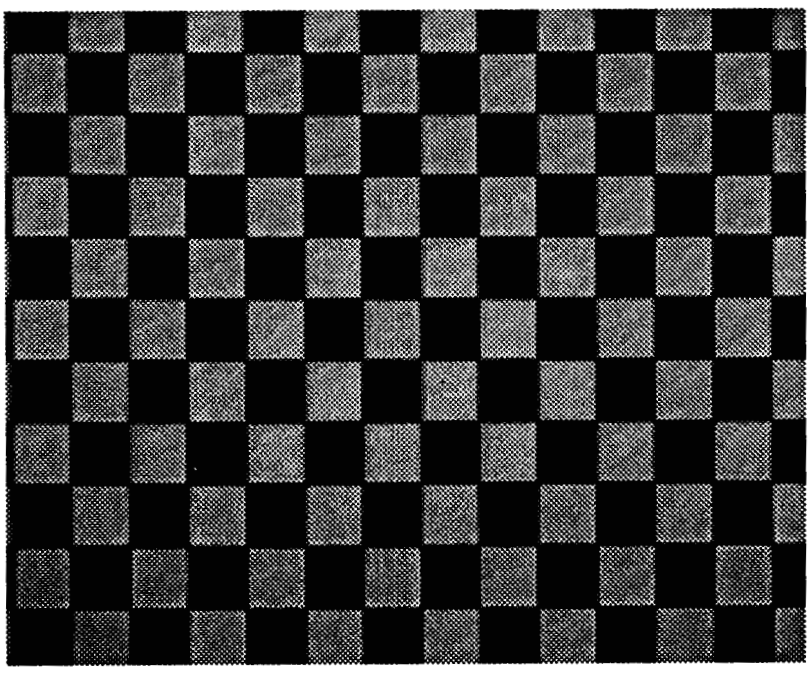


Figure 2: Checkerboard test target

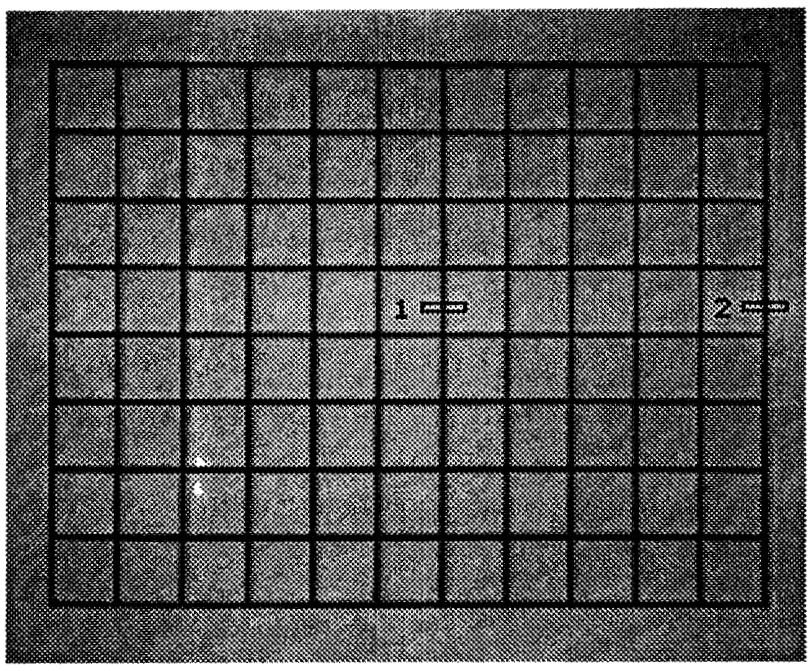


Figure 3: Image of grid pattern

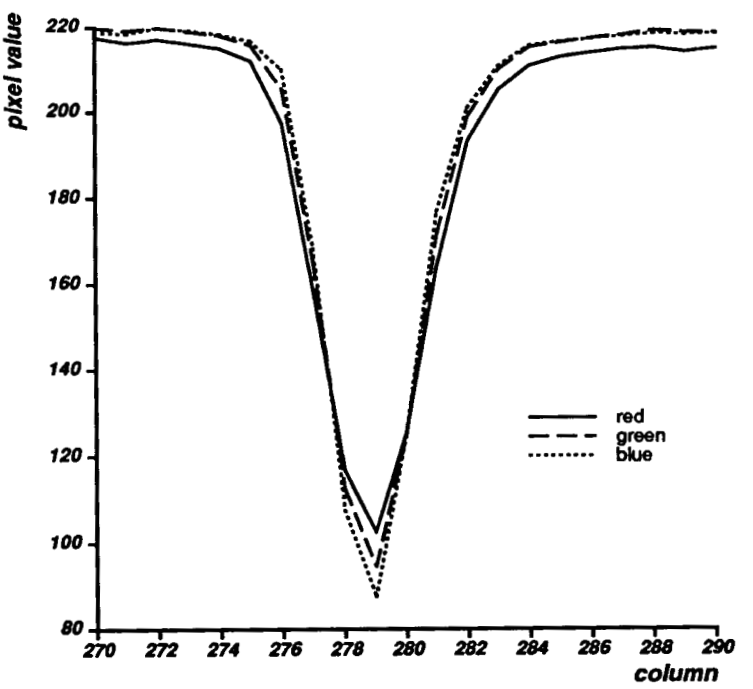


Figure 4: Image profile near center

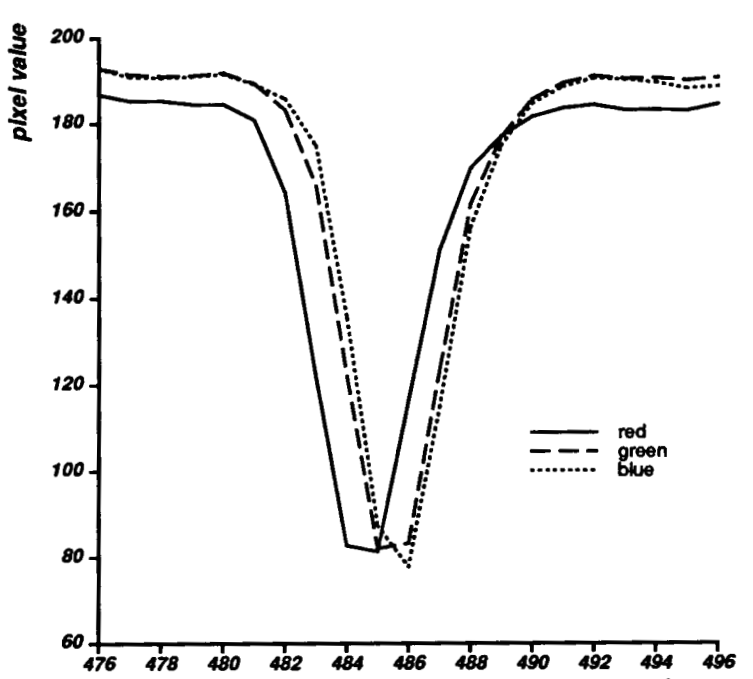


Figure 5: Image profile near edge

fect is compounded by optical roll out, where changing the lens' focussed distance and focal length settings changes the degree of misalignment, causing a drift in the image center. This effect has also been noted by Lenz [8], Nair [9] and Wiley [16]. Whereas the magnification differences are radially symmetric with respect to the center of the optical system and have a smaller magnitude nearer the center than at the edge, decentering differences affect the entire image uniformly. In our system the measured translation between the uncorrected red, green and blue bands ranges from 0.1 to 0.4 pixel widths, depending upon the image band and the lens settings.

Effects on color analysis

The demagnification, defocus and decentering between the color image bands means that the three band values for each pixel do not in fact correspond to the same point in object space. This has rather obvious impacts on color image analysis. For example, regions near rapid color changes (eg. near highlights or the edges of features), will show hue shifts. If we were to look closely at the pixels in section 2 of Figure 3 (shown in profile in Figure 5), going left to right, we would see the pixel colors change from white, to bluish, to black, to reddish, and back to white. As we will show later, these hue shifts show up as erroneous points in color histograms of the image which in turn can lead to color segmentation problems [6]. Edge detection and edge localization in color images are also confounded in obvious ways.

2.2 Compensation Approach

The most direct approach to compensating for chromatic aberration is to modify the magnification and focussed distance for each of the color bands to null out the magnification and defocus introduced by the chromatic aberration. Correction for image decentering can be accomplished by shifting the camera in the XY (sensor) plane. We call this compensation approach Active Color Imaging.

In our present color imaging system the test target shown in Figure 6 is first placed directly in front of the scene to be imaged so that the compensated lens settings can be determined. The black to white edges on the target are used to focus on and to determine the relative image magnification and shifting between the color bands. Once the correct focus, zoom and shift settings for the lens have been determined the calibration target is removed and a compensated color image of the scene is taken.

In our approach the blue image is used as ground truth. In the calibration procedure the lens settings for the red and the green bands are determined such that the errors between the blue and red images and the blue and green images are minimized. To determine the compensated lens settings for the red and green images we first find the position of best focus for each of the three bands using the automated focusing algorithm developed by Krotkov [7]. Next we take three focus corrected images and determine the magnification differences between the blue and red and the blue and green bands. Using a calibrated lens model (described in appendix B) that relates the image magnification to the positions of the focal length (zoom) and focussed distance motors, we then change the zoom appropriately for the red and the green bands to correct the magnification differences that result from the chromatic aberration and from the refocusing operation. Finally, to correct the color dependent image translation, we take three focus and magnification corrected images and then determine the amount of camera shifting that is needed in the red and green bands to compensate for translation introduced by the optics.

Calibration example

By examining the positions of the 12 edges on a scanline crossing the center of the calibration target in Figure 6 we can graphically show the effects of each of the steps of the compensation procedure. Plotting the difference in the edge positions (in pixel widths) between the uncorrected blue and red images versus the edge number, we obtain the solid line in Figure 7. For a lens with no chromatic aberration the line would be zero everywhere, indicating no difference in the positions of the edges in the blue and red images. Our plot shows two effects. The slope of the blue-red line is a measure of the relative magnification of the blue and red images. The displacement of the line up or down is a measure of the X component of the image translation. In this example the magnification difference between the uncompensated blue and red images is -0.45% and between the blue and green images is -0.14% [1]. The X shift between the blue and red images is 0.12 pixels and between the blue and green images is 0.08 pixels.

The dotted line shows the blue-red differences after refocusing the red image. The magnification difference has now shifted in the other direction as a result of the magnification change introduced by refocusing. After focus compensation the magnification difference measured between the blue and red images is 1.01% and between the blue and green images is 0.35%. The X shift between the blue and red images is 0.35 pixels and between the blue and green images is 0.14 pixels.

The dashed line shows the blue-red differences after rezooming the red image. While the slope of the line is approximately zero, indicating that the magnification differences between the bands have been eliminated, the image translation is now obvious. After focus and zoom compensating the images the magnification difference between the blue and red images is 0.07% and between the blue and green images is 0.01%. The X shift between the blue and red images is 0.46 pixels and between the green and blue images is 0.12 pixels.

The dash-dot line shows the final result, after the camera has been shifted. The plot is now close to zero everywhere, emulating ideal lens behavior. At the edge of the image the misregistration has dropped from about one pixel to about one-tenth of a pixel. In the final focus/zoom/shift compensated images the magnification difference between the blue and red images is 0.05% and between the blue and green images is -0.01%. The X shift between the blue and red images is 0.00 pixels and between the blue and green images is 0.01 pixels.

[1] For clarity the blue-green plot is not shown on this graph.

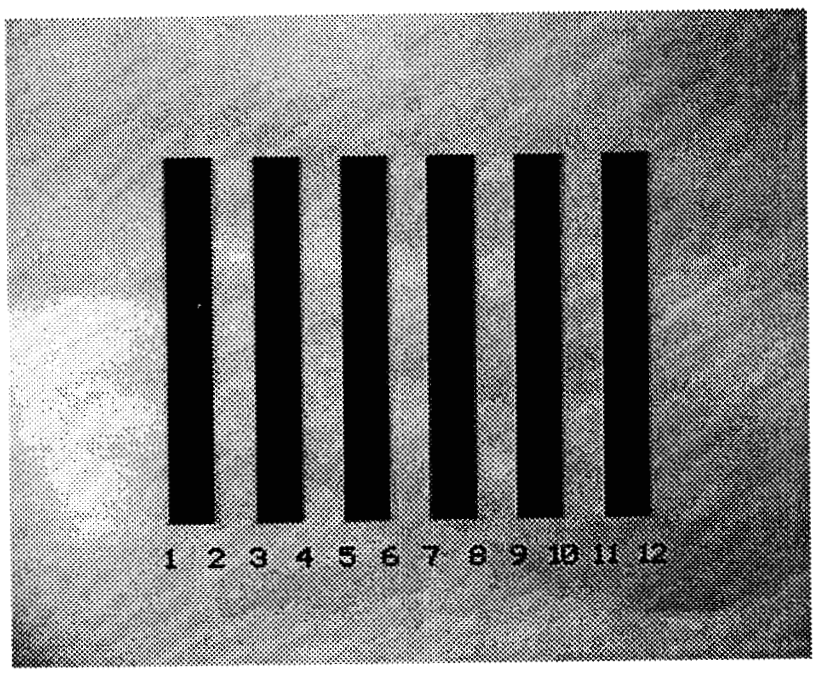


Figure 6: Active Color Imaging calibration target

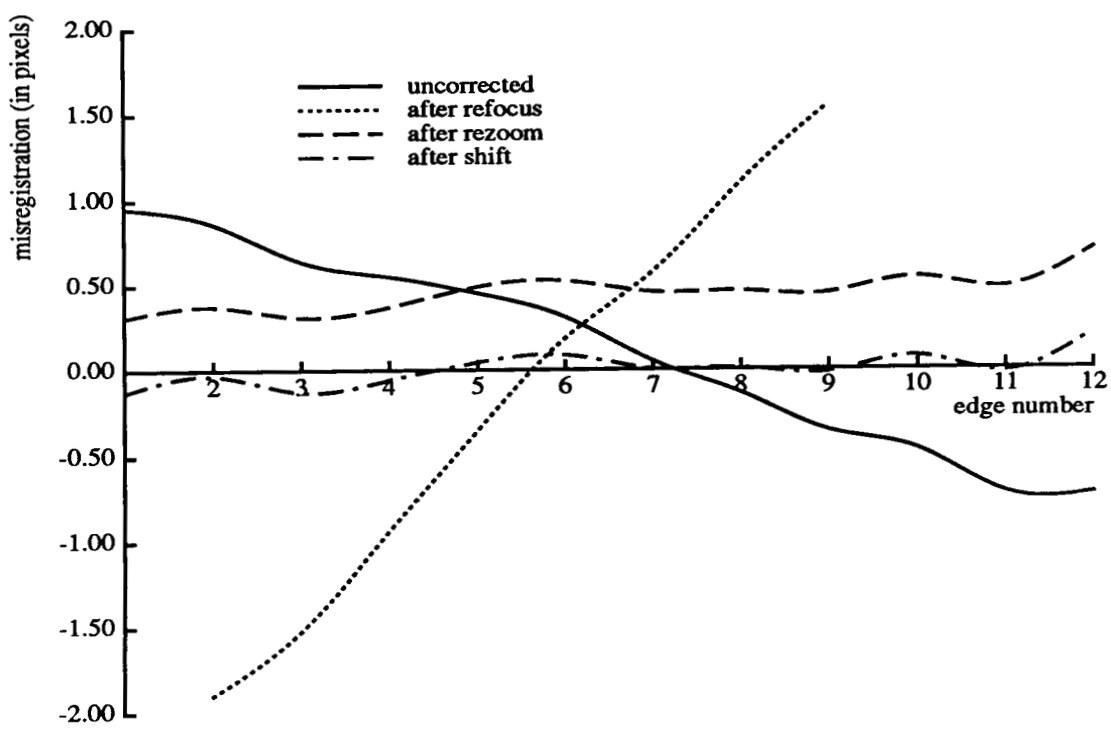


Figure 7: Chromatic aberration correction process

Active Color Imaging requires very fine control of the lens parameters. For this example the lens and camera parameters were changed by the amounts shown in the following table.

band	focus motor (out of 3900)	difference from blue setting	focal length motor (out of 4000)	difference from blue setting	camera X shift
red	1407	-6.36%	533	+0.83%	0.006 in
green	1592	-1.62%	508	+0.20%	0.002 in
blue	1655		500		

In this example the focal length motor was initially set to 500 units ($f_l \approx 67$mm) when the camera was set up. At the start of the Active Color Imaging procedure the lens was focussed with the blue, green and red filters. The final focus motor positions are listed in column 2, with the percentage deviation from the blue settings listed in column 3. Using the refocussed images, the compensating focal length motor settings were calculated for the green and red images. The final zoom motor settings are listed in column 4, with the percentage deviation from the blue settings in column 5. Finally, using the focus and zoom compensated images, the required X camera shift was calculated for the green and red images. These values are shown in column 6. Shifting in the Y direction was not performed for this example.

2.3 Experimental Results

To measure the effects of active lens compensation across the full field of view of the camera we use a checkerboard target and the procedure described in appendix A. Figures 8 and 9 show the magnitude of the image misregistration (in pixel widths) between the blue and red images coded in a gray scale ranging from 0 pixels (black) to 1.5 pixels (white). Figure 8 shows the blue-red misregistration for the target imaged without lens compensation. The misregistration across the image ranges from 0 to 1.2 pixel widths. For the blue-red case the zero error region appears to the lower right of the center of the image. In general the location of the zero error region for blue-green case will not be the same. Figure 9 shows the magnitude of the blue-red misregistration for the target imaged with lens compensation. The remaining misregistration is now less than 0.1 pixel widths over most of the image. So with active lens compensation it is possible to reduce the image band misregistration introduced by chromatic aberration by over an order of magnitude.

Another way of showing the effects of active lens compensation is to plot image pixels in three dimensional RGB space. Figures 10 and 11 contain color histograms of the region containing the rightmost vertical line of the grid image shown in Figure 3. Figure 3 is a color image of a black grid on a white background so ideally all of the pixels in the image should be black, white, or shades of gray. The plotted pixels from the region containing the grid line should form a straight line of grays between the lower left corner of the color cube (corresponding to black) and the upper right corner of the cube (corresponding to white). This is not the case in the uncorrected image. In Figure 10 the histogram of the pixels from the uncorrected image shows two distinct paths from black to white. The first path bends towards the red corner of the color cube while the second path bends towards the blue corner of the color cube. These deviations would pose a serious problem for any algorithm trying to segment the image by using regions in the color space. In Figure 11 the histogram of the pixels from the image taken with active lens compensation shows a much tighter grouping of pixel values along the gray line.

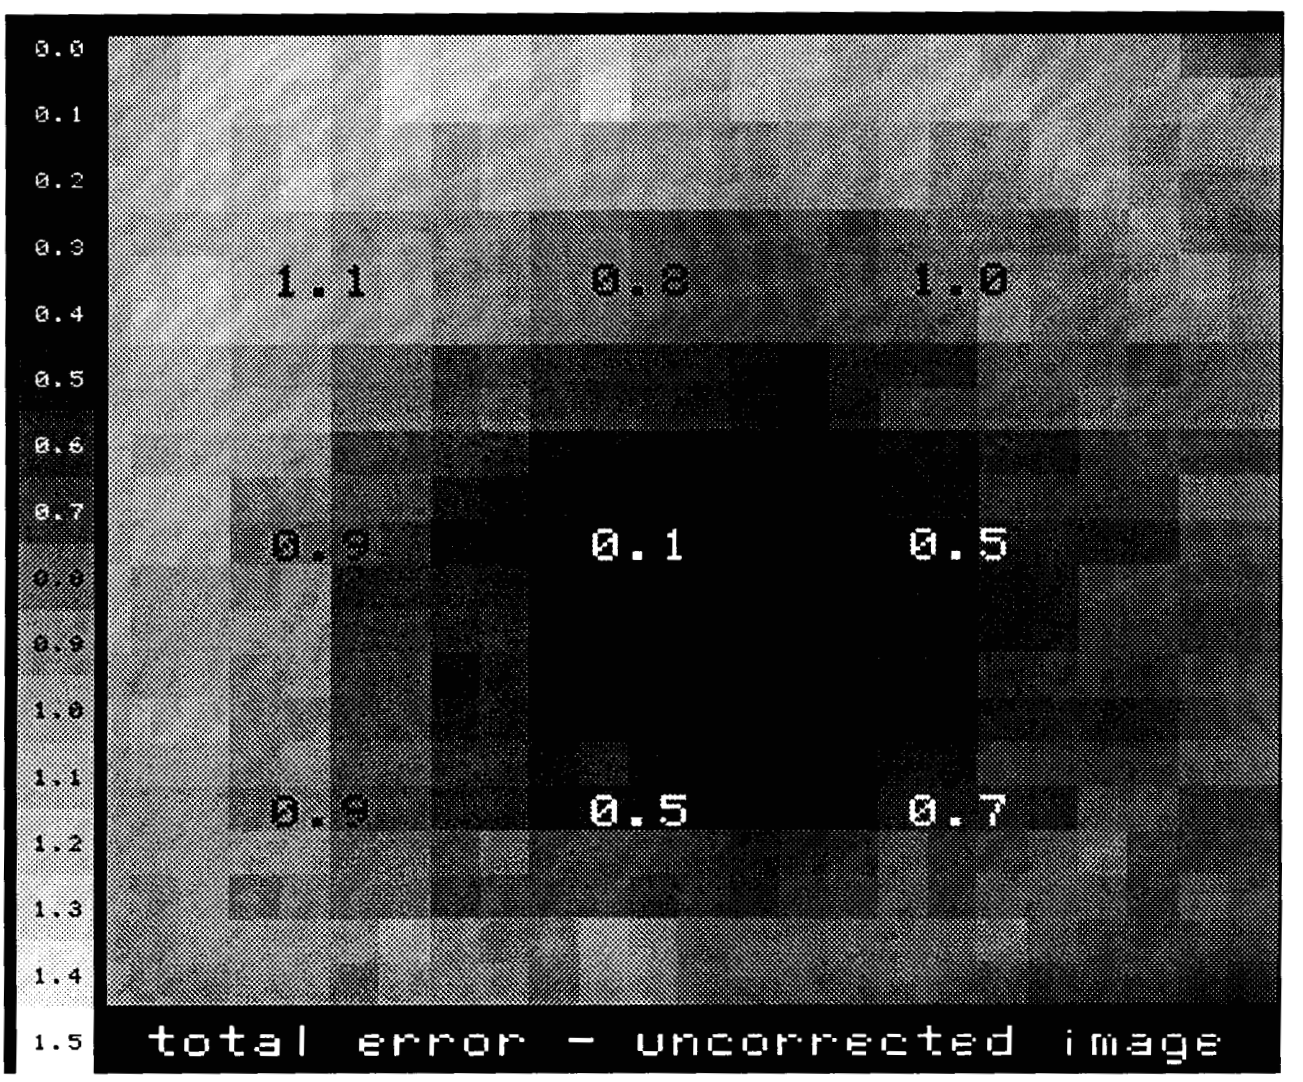


Figure 8: Full field error - uncompensated lens

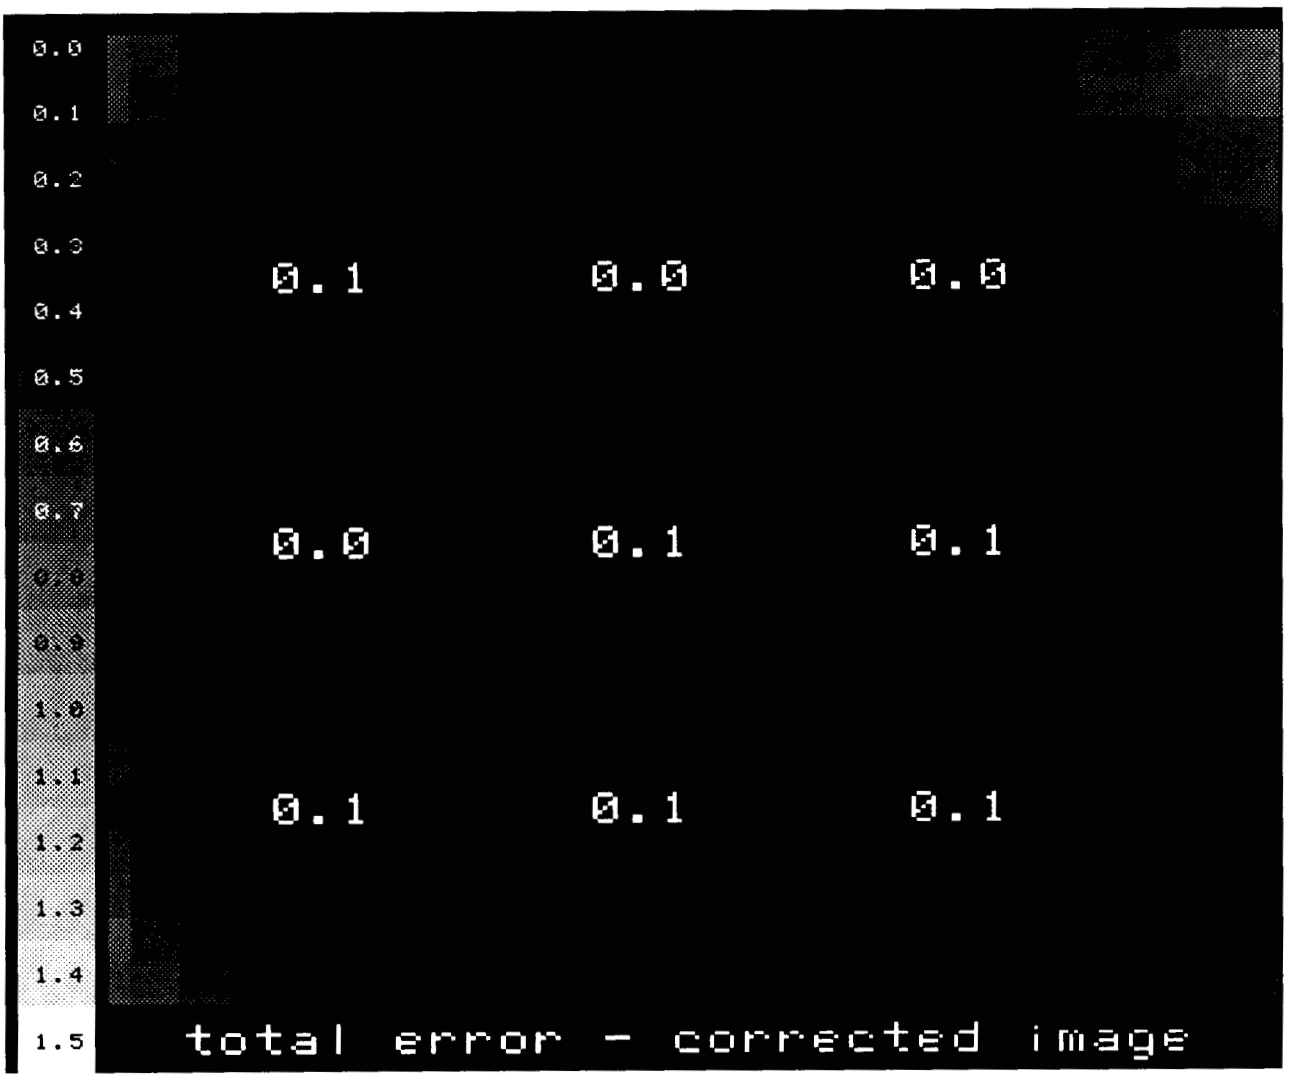


Figure 9: Full field error - compensated lens

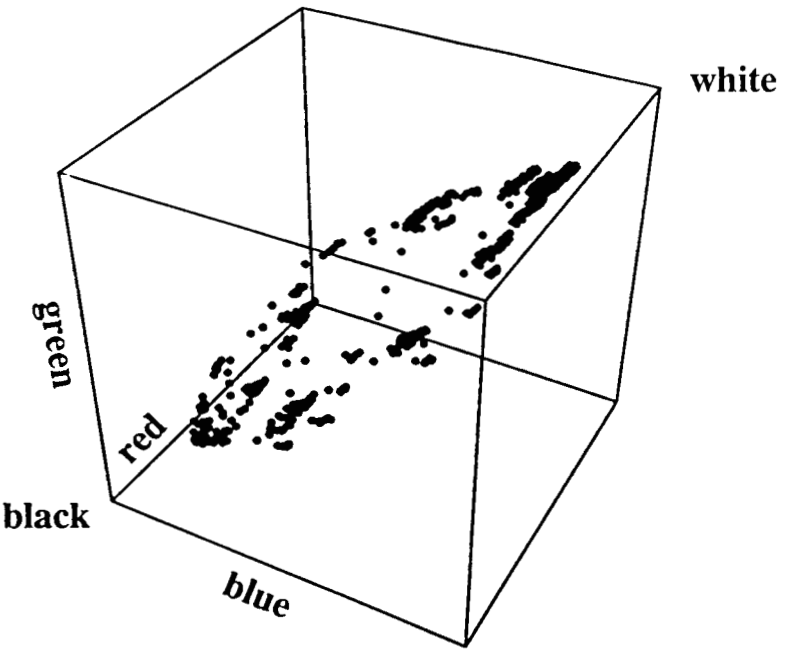


Figure 10: Color histogram - uncompensated image

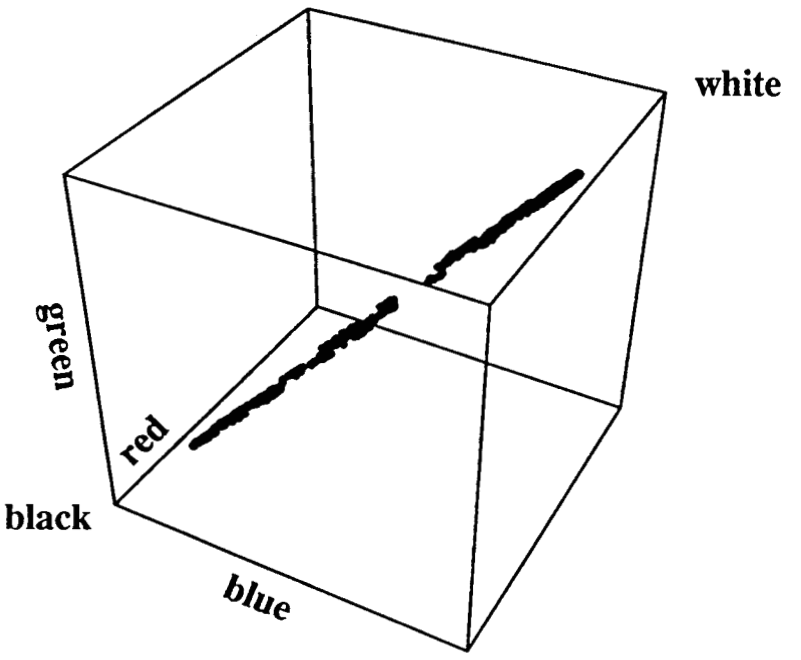


Figure 11: Color histogram - compensated image

2.4 Conclusions

With the lenses typically used in computer vision[2], the processes by which color images are taken will generally produce images that have magnification, focus and image center differences between the color bands. These effects are due both to chromatic aberration and to optical misalignment in the camera lens. The end result is significant levels of misregistration and defocus in the final composite image. With active compensation of the lens settings and the camera position the misregistration between the color bands of an image can be reduced by an order of magnitude.

One alternative to active lens control for color image correction is to post process the image bands to remove the chromatic aberration effects. In [2] Boult compares our active lens control approach to their image warping approach for chromatic aberration correction. The major advantages of image warping are that it can be applied to cameras that take RGB images in parallel (eg. color video cameras), it does not require special lens control hardware, and it can handle chromatically varying geometric distortions that cannot be corrected by focusing or zooming. The disadvantages of the image warping approach are its inability to compensate for the defocus between image bands and the relatively large amount of calibration data that it can require. Boult found that while image warping compared reasonably well both qualitatively and qualitatively with the active lens compensation approach, the active lens compensation approach could produce better overall results for simple uncorrected lenses because of its ability correct for focus differences between bands.

In our present color imaging procedure we determine the control settings for the compensated images by imaging a test target to measure the amounts of defocus, magnification change and decentering between the color bands. We are now proceeding to develop calibrated lens models that relate the focussed distance, magnification and optical center of the lens to the camera's control motors and the imaging color. With these models it will be possible, given the lens settings for one color band, to calculate the settings for the remaining two bands without the necessity of a test target.

[2]We have measured comparable levels of chromatic aberration in a wide variety of lenses, including CCTV lenses, 35mm SLR lenses and ENG/EFP color TV lenses, including lenses that are advertised as being corrected to eliminate chromatic aberration.

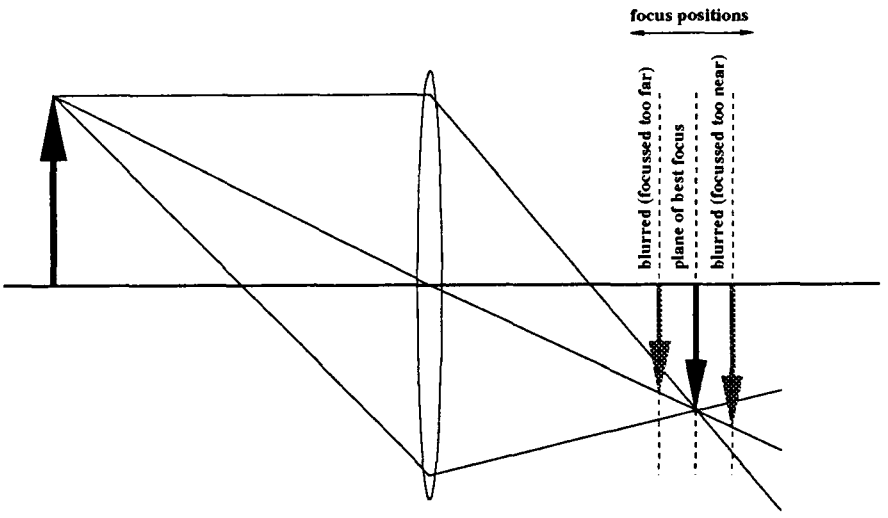


Figure 12: Focus magnification

3 Constant Magnification Focusing

In range from focus the distance to a target is estimated by moving the camera lens through a series of focus positions while determining the sharpness of the target's image at each position. The target's sharpness is determined by evaluating a sharpness criterion function over the image. Given the focus position of the lens where the criterion function peaks, the distance between the camera and the target can be determined using a calibrated lens model. Implicit in the range from focus procedure is the assumption that the criterion function peaks at the point of sharpest focus[3]. There is also the implicit assumption that the value of the criterion function varies only with the focussed distance of the lens. In the following section we demonstrate how, for the most widely used class of criterion functions, the criterion function value (and the position of its peak) is also dependent on the lens' aperture and on the content of the image.

Many sharpness criterion functions have been suggested for focus ranging [7]. In our experiments we make use of a sum of squared image gradients function, also called the Tenengrad function. Most of our discussion transcends a specific choice of criterion function.

3.1 Focus Magnification

For a fixed focal length lens, focus magnification is the change in image magnification that results as the camera's sensing plane is moved along the optical axis in order to vary the lens' focussed distance. This is illustrated in Figure 12. Effectively focus magnification causes the image to scale up as the lens is focussed from far to near. For a variable focal length (zoom) lens focus magnification is conceptually the same.

If one considers focus magnification then the net magnification of a zoom lens is actually a function of two lens variables: the focal length and the focussed distance. By measuring the dimensions of a target across a range of lens settings we have produced a calibrated lens model relating the focal length and focussed distance motor settings to the relative

[3]The point of best focus is usually defined as the point at which the region of interest has the highest spatial frequency content.

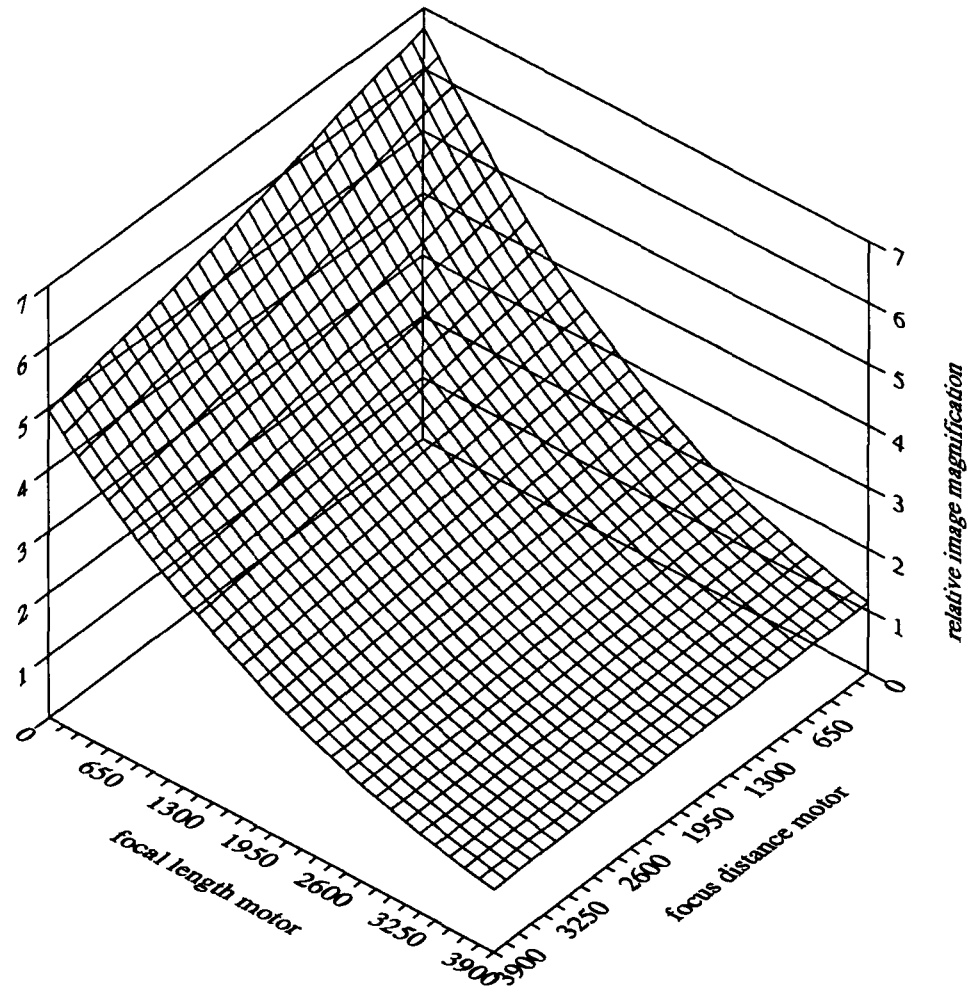


Figure 13: Relative magnification versus f_d and f_l motor positions

image magnification ($M = g(f_l, f_d)$). Appendix B describes the relative magnification model in detail. Figure 13 shows a plot of the relative image magnification versus the focussed distance motor position and focal length motor position for the lens used in these experiments.

Image effects

Focus magnification causes three significant problems for range from focus: feature clipping, perimeter scaling, and gradient scaling. The first of these effects, illustrated in Figure 14, involves the movement of features into or out of the window of interest as the image magnification changes. This effect is pretty straight forward and has been noted by several researchers, including Krotkov [7], Darrell [4], Nair [9] and Nayar [10]. The second effect, perimeter scaling, illustrated in Figure 15, involves the proportional scaling of the length of feature perimeters as the image magnification is changed. Effectively this is image scaling done perpendicular to the direction of image intensity gradients. The final effect, gradient scaling, illustrated in Figure 16, involves the inverse scaling of the width and slope of inten-

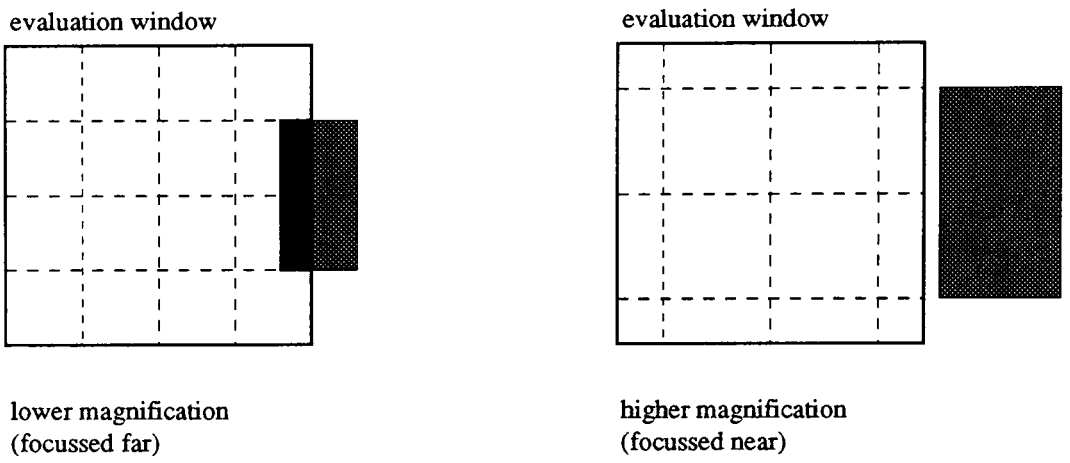


Figure 14: Focus magnification - feature clipping

sity gradients as the magnification changes. Effectively this is image scaling done parallel to the direction of the image intensity gradients.

Effects on focus ranging

For accurate focus ranging we would like the position of the peak of the criterion function to depend solely on the focussed distance of the lens. As we will demonstrate, focus magnification can make the position of the peak sensitive to additional factors.

The first focus magnification effect, feature clipping, causes abrupt changes in the criterion function value as features leave or enter the evaluation window. These changes in turn may result in multiple or false criterion function peaks which can lead to erroneous range estimates. The last two focus magnification effects, perimeter scaling and gradient scaling, cause more subtle biases in the position of the criterion function peak.

With focus magnification the size of an image scales up as the lens is focussed from far to near. As the image scales up feature perimeters will also be proportionately scaled up causing the value of the criterion function to increase. This increase in turn causes a bias in the position of the criterion function peak, making the point of best focus appear too near. However, as the image scales up with focus magnification, the intensity gradients on features will also spread out laterally. If the width b of the gradient shown in Figure 16 is significantly larger than the width of the point spread function of the focussed lens, then the blur introduced by the defocusing lens will not significantly alter the value of the gradient across most of the width b. In the one dimensional case for a sum of squared gradients criterion function, the value of the criterion function is given by

$$value = \frac{a^2}{b}.$$

Thus as the image (and b) are scaled up, the value of the criterion function is proportionately scaled down. This decrease causes a bias in the position of the criterion function peak making the point of best focus appear too far away.

Both the perimeter and gradient scaling effects are functions of the image magnification and the rate of blurring as the lens is moved away from the point of best focus. The blur

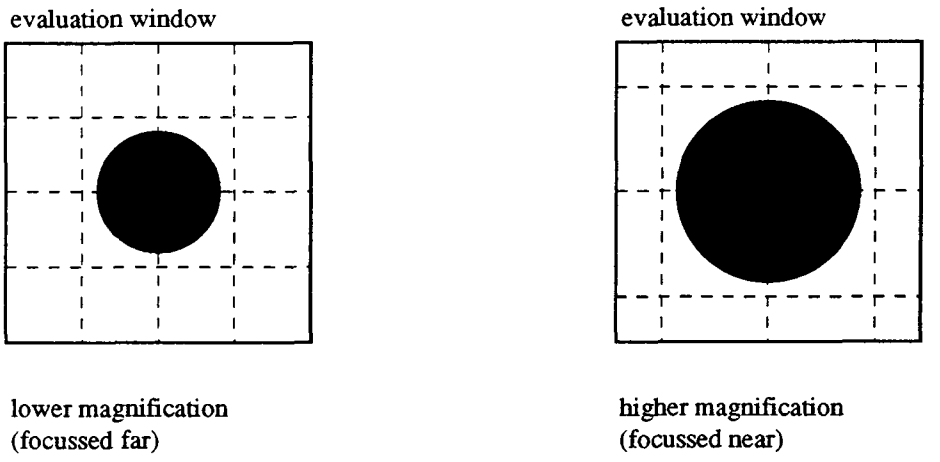


Figure 15: Focus magnification - perimeter scaling

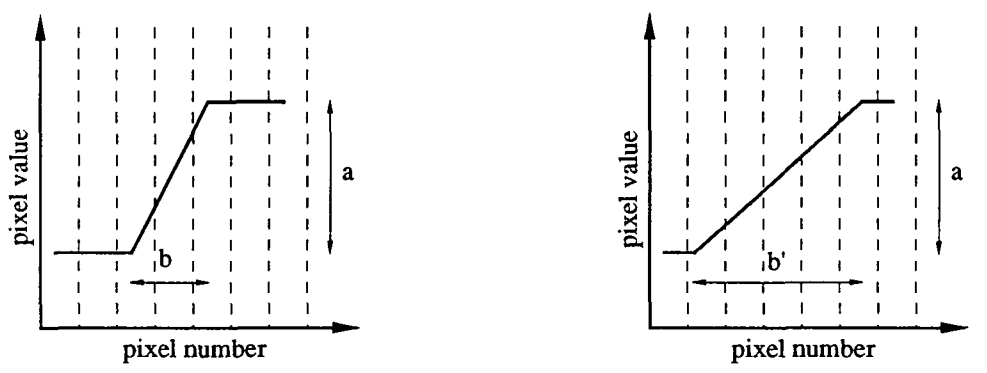


Figure 16: Focus magnification - gradient scaling

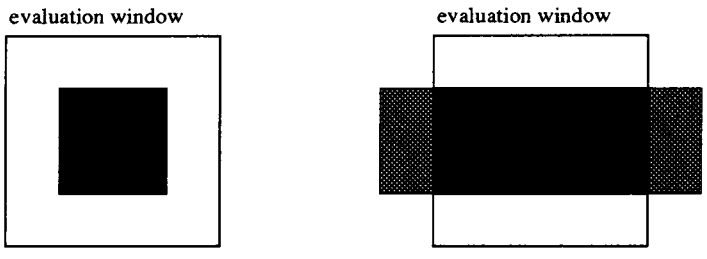


Figure 17: Box and bar focus targets

rate is effectively a function of the lens' aperture. The net result of the two competing effects is an additive bias to the position of the criterion function peak which is dependent on the lens' aperture and on the content of the image.

3.2 Compensation Approach

To avoid focus magnification effects in focus ranging researchers have suggested either keeping the image magnification constant [4] or keeping the evaluation window large enough to completely encompass the features being ranged [7]. As suggested above however, simply keeping the features being ranged entirely within the evaluation window will not overcome all of the focus magnification effects. Only a constant image magnification approach can avoid the problems that result from focus magnification.

In [4] Darrell suggests two approaches for dealing with focus magnification. The first approach, and the one implemented by Darrell, is to cancel out the magnification change by scaling the images in software after they are taken. This approach is relatively simple and works even for fixed focal length lenses. One major drawback with this approach is that as the image is scaled, the spatial interpolation and resampling process will smooth the image and influence the value of the criterion function. An additional drawback is the computational effort required to scale the images.

The second approach suggested by Darrell is to null out the focus magnification by changing the lens' focal length before the images are taken. While being more efficient and effective than scaling images, this approach requires precise active lens control as the focus distance of the lens is varied. For our system, given the current focussed distance and focal length settings and a new focussed distance setting, we use the relative magnification model described in appendix B to find a new focal length setting which maintains the same relative magnification. We call this compensation approach Constant Magnification Focusing.

3.3 Experimental Results

To illustrate the effects of focus magnification on the position of the Tenengrad criterion function peak we use two targets, illustrated in Figure 17. Target 1 is a black square on a white background, completely enclosed in the evaluation window. Target 2 is a black bar on a white background with only the center region of the bar contained in the evaluation window. For both targets, the majority of the value of the criterion function results from the black to white edge, with the totally white and totally black regions contributing insignificant amounts. For target 1 the length of the perimeter of the black region changes with the focus magnification, while for target 2 the perimeter length remains constant. For

both targets the width of the gradient at the black to white edge is not large enough (relative to the lens' point spread function) to produce gradient scaling effects. The box and bar targets are both located on the same plane exactly 1.5m ±5mm away from the camera's sensor plane.

For these experiments we use a sum of squared image gradients criterion function with the image gradients being calculated using a 3x3 Sobel operator. The position of the criterion function peaks are determined by least squares fitting the data to a quadratic curve, and then solving the quadratic for the position of the peak. Appendix C describes the peak locating procedure in detail.

To show the effects of focus magnification two different blur rates are selected by using two different aperture values. At higher blur rates (wider apertures, narrower depths of field) the value of the Tenengrad function falls off too quickly for focus magnification effects to significantly bias the position of the criterion function peak. At lower blur rates (narrower apertures, wider depths of field) the value of the Tenengrad function falls off more slowly and focus magnification has a more pronounced effect.

The experiments are done first without and then with Constant Magnification Focusing. For the uncompensated focusing the focussed distance motor is varied while the focal length motor is held constant at 500 motor units ($f_l \approx 67$mm). The image magnification over the range of focus changes by a factor of 0.816. For Constant Magnification Focusing the focussed distance motor is varied while the focal length motor is concurrently varied to keep the magnification constant. Figure 18 shows the focal length motor settings versus the focussed distance motor settings for both the compensated and the uncompensated focusing. For compensated focusing the focal length motor settings are varied from 747 motor units ($f_l \approx 63$mm) to 277 motor units ($f_l \approx 71$mm). The Constant Magnification Focusing curve is essentially an isomagnification contour from the magnification model shown in Figure 13.

Figure 19 shows that for a wider aperture the positions of the criterion function peaks are the same for both the box and bar targets, both with and without magnification compensation. In contrast Figure 20 shows that for a narrower aperture with no compensation there is a significant bias in the peak position for the box target towards near focus, as predicted earlier. When the focus magnification is compensated for, the bias in the peak position for the box target is eliminated. With an actual range of 1.5m, the bias in the uncompensated box target's peak position corresponds to a 6% error in the range measurement determined from the calibrated lens model.

Even without focus magnification effects the position of the criterion function peak varies with the size of the aperture. This can be seen more clearly in Figure 21 which shows similar variations in the peak positions for a range of aperture positions for both the box and bar targets. The error bars in this graph are the estimated standard deviations of the criterion function peak position, as calculated by the peak finding procedure described in appendix C. One potential explanation for this variation is spherical aberration in the optics. As the aperture is opened the addition of marginal rays to the image changes the distribution of light in the 3 dimensional point spread function of the lens, potentially altering the position of the point spread function's centroid.

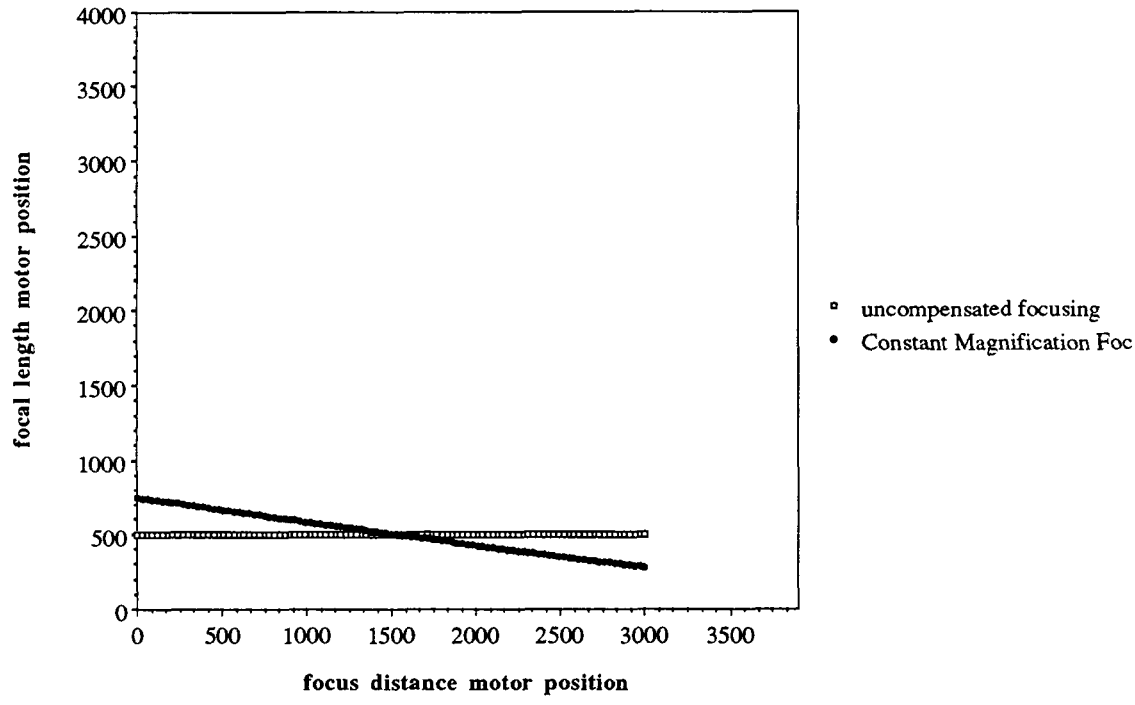


Figure 18: Lens settings for range from focus

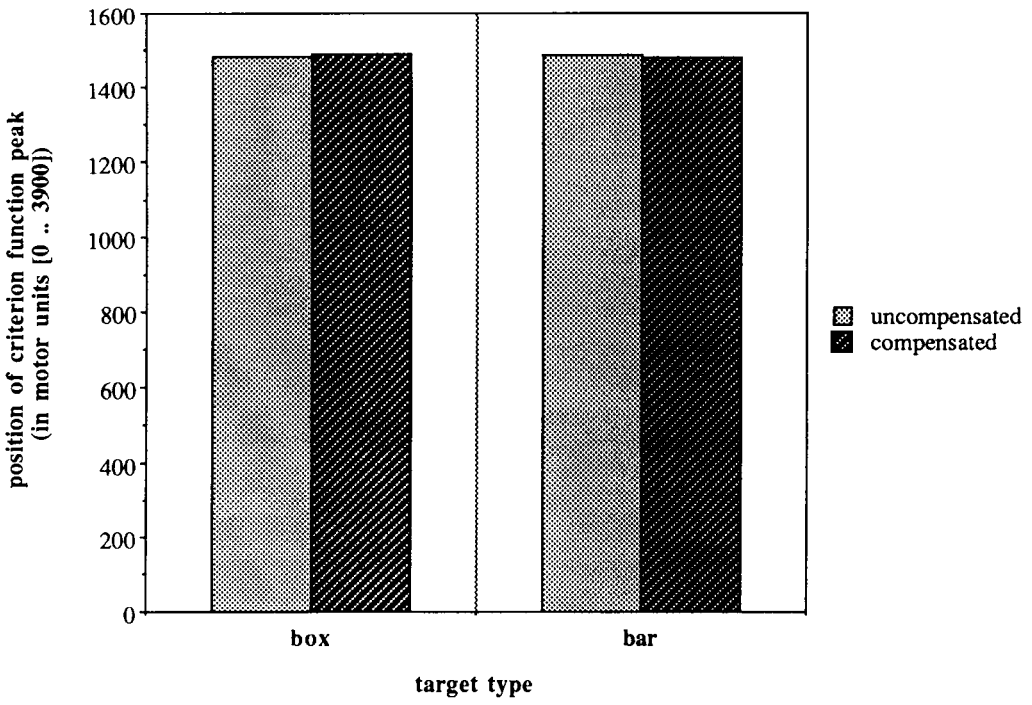


Figure 19: Position of criterion function peak - Wide aperture (f/# $\approx$ 3.3)

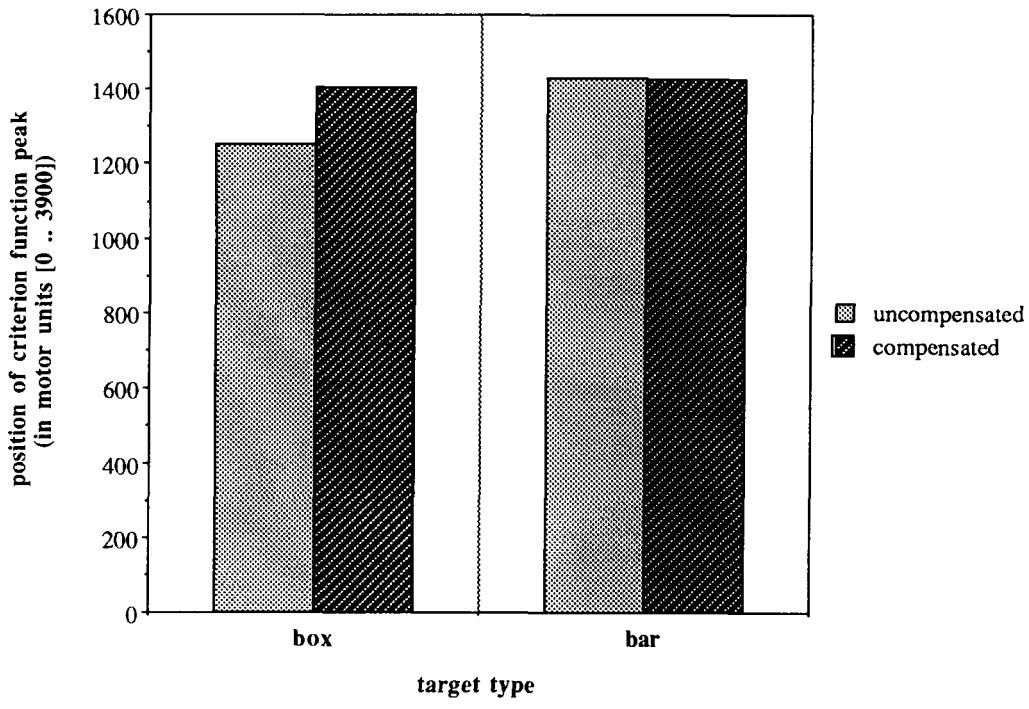


Figure 20: Position of criterion function peak - Narrow aperture (f/# $\approx$ 22)

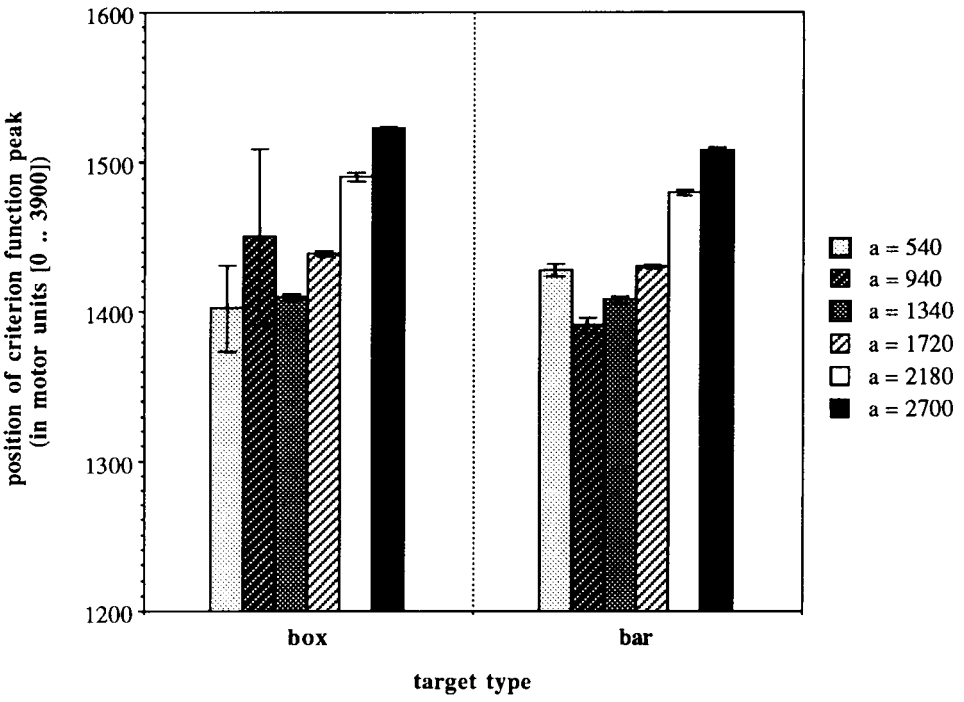


Figure 21: Criterion function peak position variation with aperture

3.4 Conclusions

In range from focus tasks where the lens' focus is varied without magnification compensation the position of the sharpness criterion function's peak can potentially be biased away from the point of best focus by several image scaling effects. The amount of bias depends on the size of the aperture and on the content of the image region being evaluated. With active lens compensation, focus magnification induced bias problems can be efficiently and effectively eliminated.

Our experiments have shown that the position of criterion function peaks also varies with lens aperture, possibly due to spherical aberration within the lens. Whatever the cause this behavior implies that a calibrated lens model that relates lens settings to focussed range will have to include a parameter for the aperture position.

Our work with chromatic aberration has shown the strong dependence of a lens' focussed distance on the wavelength of the light involved. This suggests that to eliminate chromatic aberration as a source of error in range from focus, all lens model calibration and range from focus measurements should all be done from within a single color band.

4 Summary

In this report we have demonstrated how non-ideal lens behaviors can affect low level imaging tasks and have introduced precise active lens compensation as a way to overcome some of the resulting problems. For color image analysis we have shown how chromatic aberration can cause image band misregistration and defocus, which in turn, can significantly degrade the performance of color image analysis algorithms. To compensate for chromatic aberration we have developed a new color imaging procedure we call Active Color Imaging which reduces the inter band misregistration in color images by an order of magnitude. We have also developed an approach to measuring lateral chromatic aberration across the full field of view of a camera system.

For focus ranging we have shown how focus magnification can cause bias in the position of criterion function peaks, which in turn, causes incorrect estimates of range. To compensate for focus magnification we have implemented a method (first proposed by Darrell), that we call Constant Magnification Focusing to eliminate the effects of image scaling directly. As part of this work we have also developed an approach to accurately and reliably estimate the position of unimodal criterion function peaks in the presence of noise.

5 Acknowledgments

This research was sponsored by the Avionics Lab, Wright Research and Development Center, Aeronautical Systems Division (AFSC), U. S. Air Force, Wright-Patterson AFB, OH 45433-6543 under Contract F33615-90-C-1465, ARPA Order No. 7597. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Government.

I would also like to thank Terry Boult at Columbia University for drafts of his paper.

6 Appendix A - Chromatic aberration measurement

Our basic approach to measuring lateral chromatic aberration involves accurately measuring the positions of black to white step edges in each of the bands of a color image. The relative positions of the edges in each band provide a measure of the chromatic aberration induced misregistration at that position in the image. By using a dense set of black to white edges (eg. a checkerboard pattern) we can approximate the misregistration caused by lateral chromatic aberration at any point in the image by using the nearest pairs of horizontal and vertical step edges to determine the orthogonal components of the misregistration.

Our measurement procedure involves two steps. In the first step one band of a checkerboard image is used to determine a set of reference edges who's positions can accurately and reliably be measured. In the second step the positions of the reference edges are calculated to subpixel accuracy in each of the image bands.

Reference edge selection

Reference edge selection begins with one band of the checkerboard image, shown in Figure 22. The first part of this step is to perform thresholded edge detection using a 3x3 sobel operator. The resulting image is shown in Figure 23. Following the edge detection suitable reference edges are selected from the image. Reference edges must meet four criteria. First, the edges must be fairly close to horizontal or vertical. Second, the edges should have sufficient length so that any shifting component in the direction parallel to the edge will still leave enough of the edge to measure the cross edge shifting component. Third, the edges should be sharp enough for the edge position to be accurately calculated. And finally, all of the horizontal and all of the vertical reference edges should be chosen far enough apart so that there is no ambiguity in tracking them between image bands.

To pick out reference edges we test the edge image with an 21x21 pixel test mask. The test mask for finding horizontal edges (used for detecting vertical shifting), shown in Figure 24, contains an *edge area* and two *non edge areas*. If the region of the image under the test mask contains an edge along the mask's edge area, and no edge in the mask's non edge area, and if the mask does not overlap any previously successful masks, then the section of the edge under the mask is labeled as a reference edge. Given a reference edge, the image column (or row) running across the center of the edge is used to measure the position of the edge to sub-pixel accuracy in the three image bands. Figure 25 shows the edge image with the vertical and horizontal test masks superimposed. The tick marks across the edges identify the columns and rows used to measure the positions of the edges. In this image there are 415 vertical reference edges and 275 horizontal reference edges.

Edge position measurement

The second step in measuring lateral chromatic aberration is to accurately determine the position of each of the reference edges identified in the reference edge selection step. To accomplish this we use a standard Laplacian-of-Gaussian edge localization technique. The LoG technique is nearly optimal for edge localization [3], is relatively easy to implement and is computationally very efficient. To locate the position of an edge a 21 pixel segment of the row or column crossing the reference edge is first convolved with the second derivative of a Gaussian convolution kernel[4]. The location of the convolution's zero crossing is then

[4]We use a standard deviation of 4.0 pixels for the Gaussian convolution kernel.

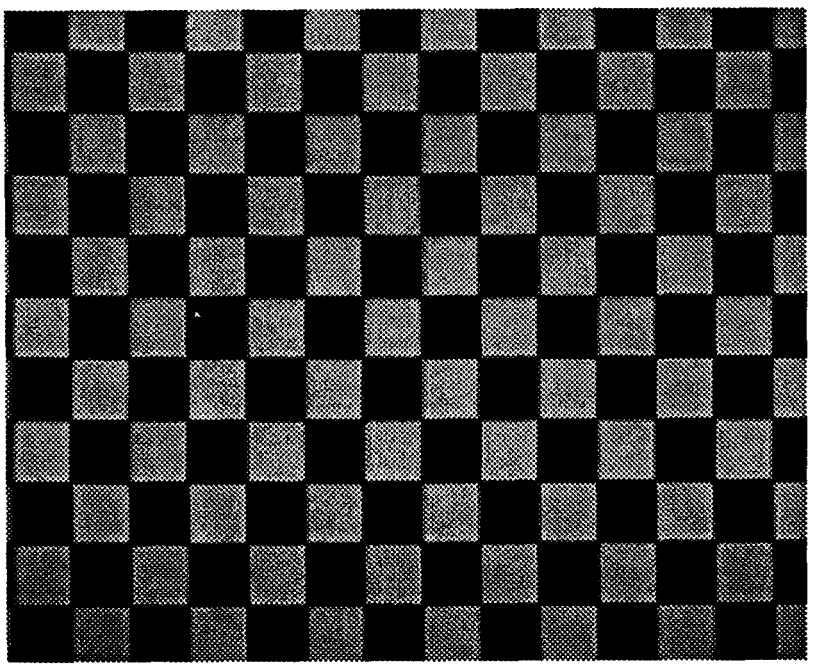


Figure 22: Checkerboard image

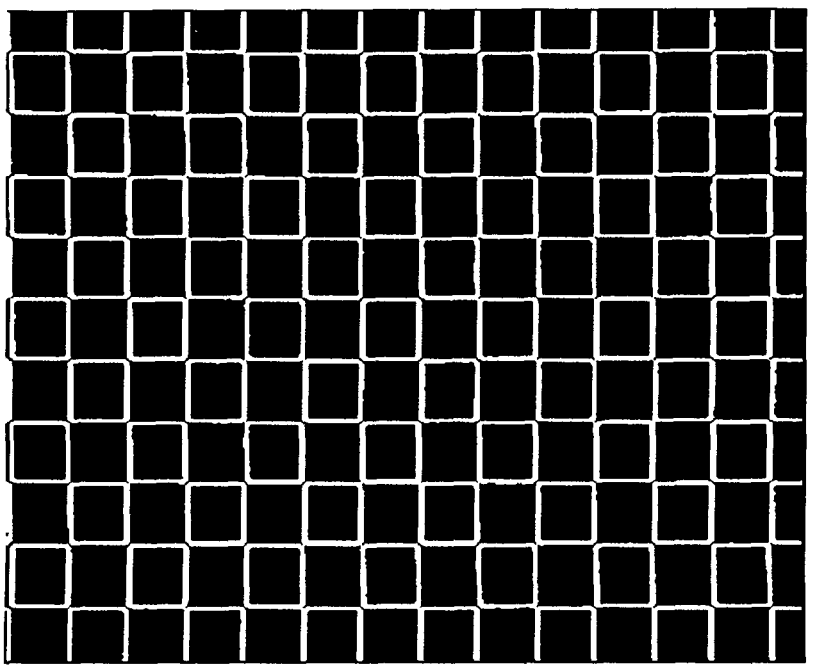


Figure 23: Thresholded edge image

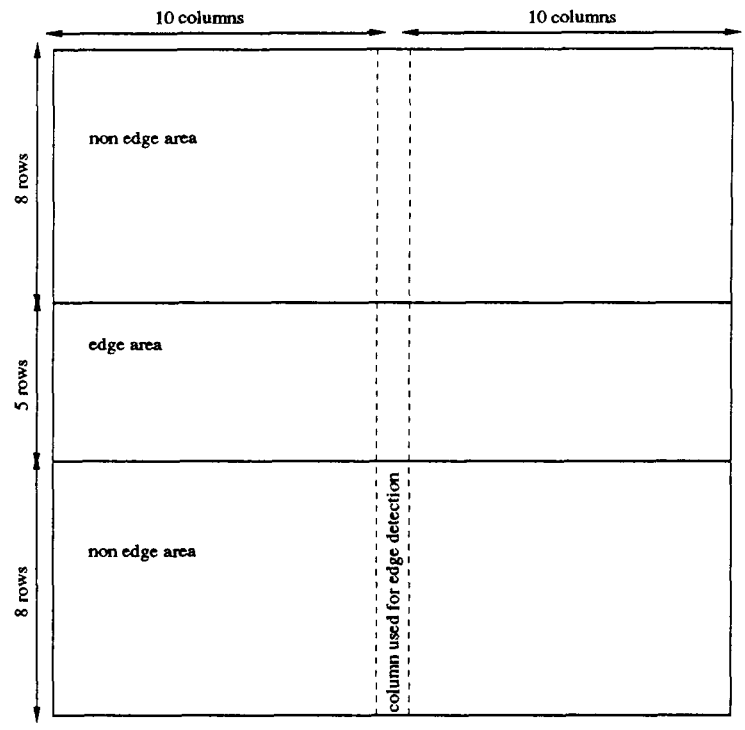


Figure 24: Test mask for horizontal reference edges

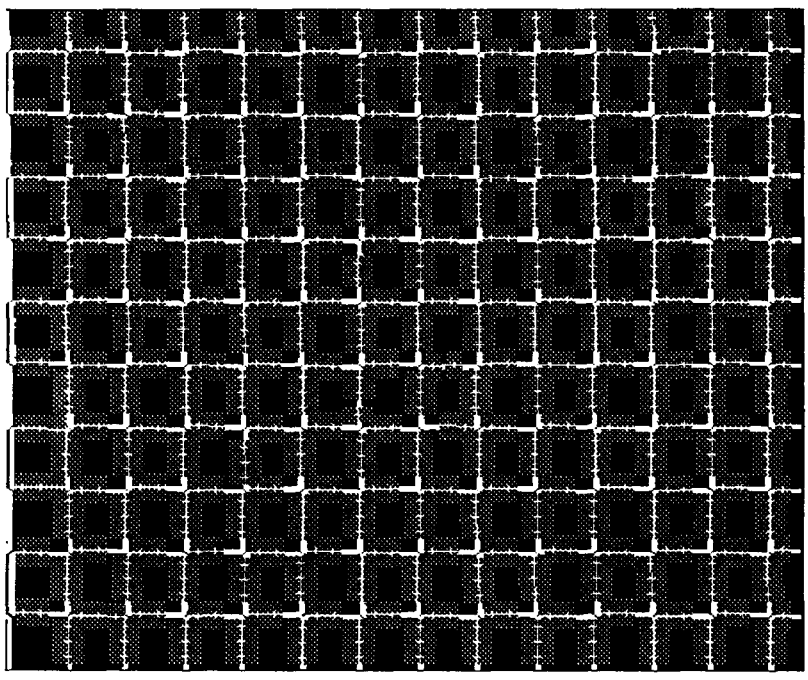


Figure 25: Edge image with highlighted reference edges

interpolated to give the exact position of the edge.

Display of chromatic aberration

To display the misregistration due to lateral chromatic aberration we plot the difference between the positions of the edges in the blue and red images and the blue and green images. This can be done for both the vertical edges and the horizontal edges. If the vertical and horizontal reference edges are close enough we can approximate the total magnitude of the misregistration at any point in the image by using the misregistration at nearest pairs of vertical and horizontal reference edges as the orthogonal components.

Figures 26 and 27 show the vertical misregistration between the blue and red and the blue and green images respectively. Figures 28 and 29 show the horizontal misregistration between the blue and red and the blue and green images respectively. Finally, Figures 30 and 31 show the total misregistration between the blue and red and the blue and green images respectively.

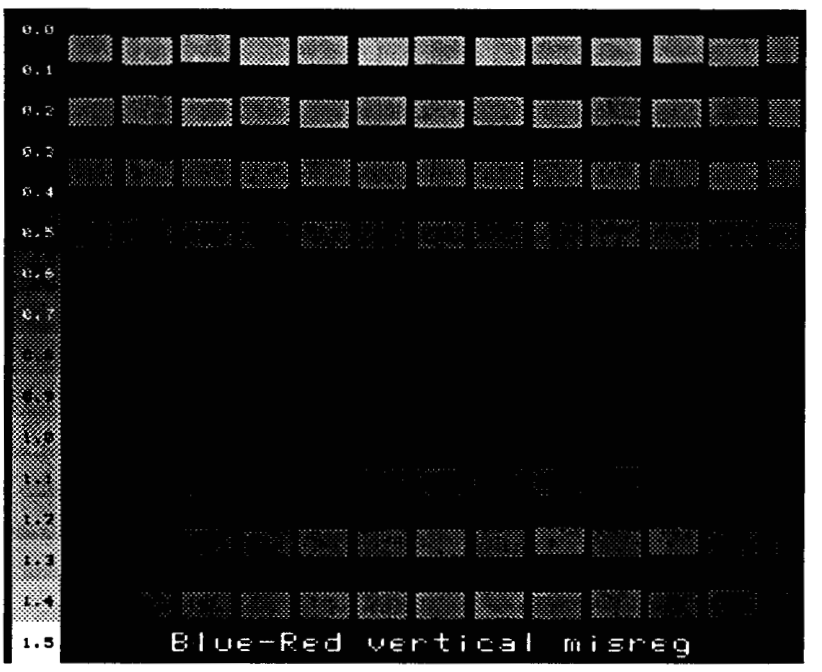


Figure 26: Blue-Red vertical misregistration

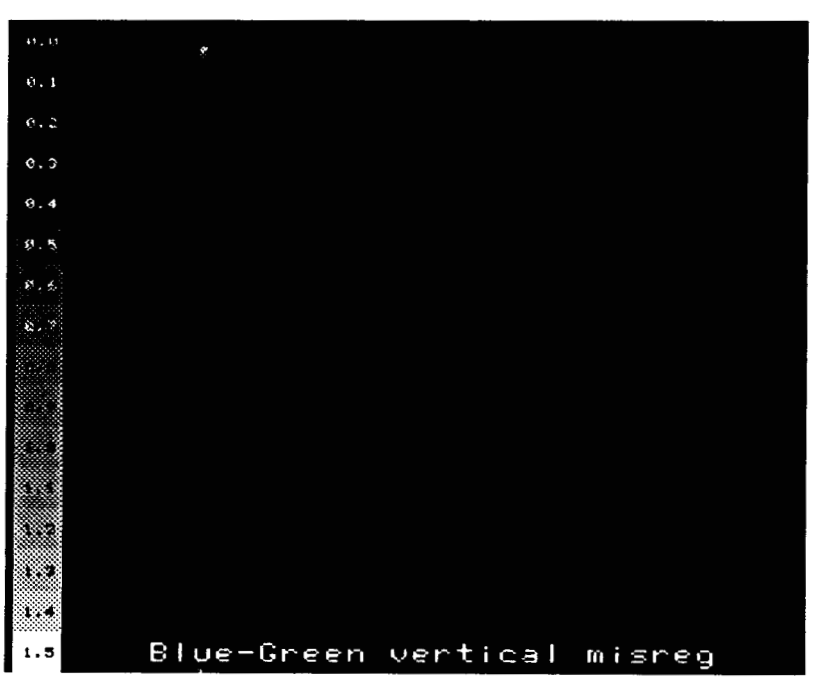


Figure 27: Blue-Green vertical misregistration

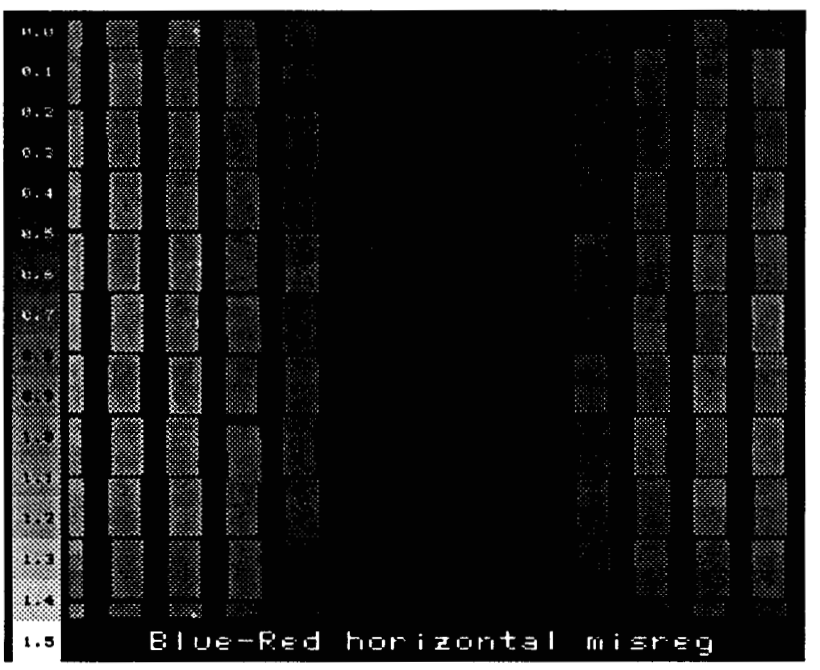


Figure 28: Blue-Red horizontal misregistration



Figure 29: Blue-Green horizontal misregistration

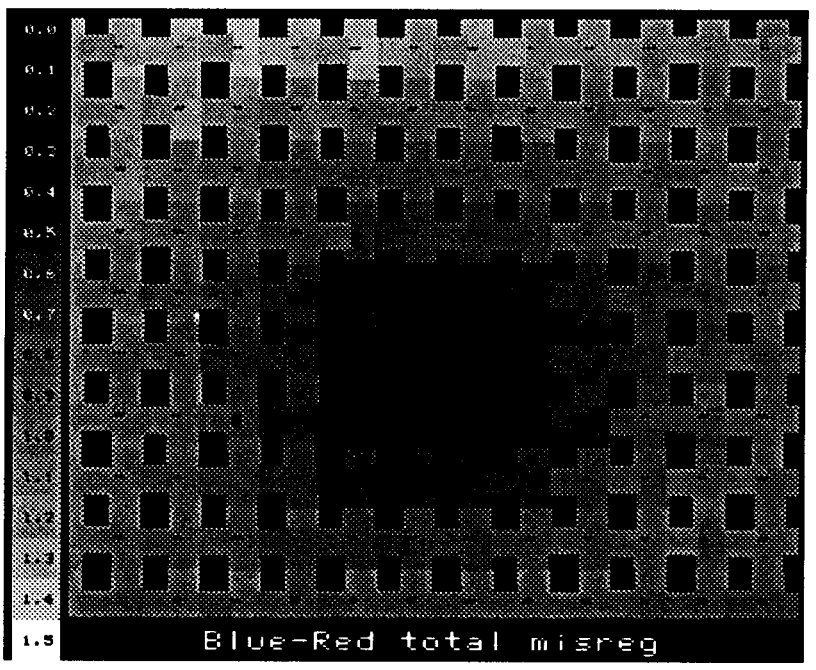


Figure 30: Blue-Red total misregistration

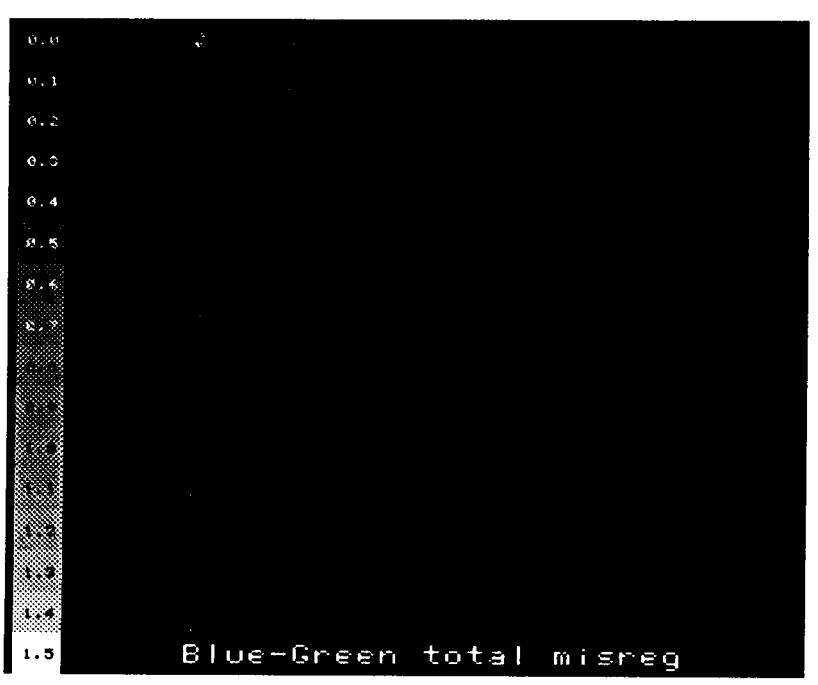


Figure 31: Blue-Green total misregistration

7 Appendix B - Relative image magnification model

Image magnification is primarily a function of the focal length of a lens. However, varying the focussed distance of the lens will also cause changes in the image magnification. While the change in image magnification introduced by varying the focus distance is somewhat smaller than the change introduced by varying the focal length, focus induced magnification changes are still significant in color imaging tasks and in range from focus tasks. For these tasks we often need to know how the image magnification changes in response to changes in both the lens' focal length and focus distance motors. We call this relationship the relative image magnification model for the lens. The model can easily be extended to include relative magnification differences that result from imaging in different color bands.

In this report we use the relative magnification model primarily for magnification equalization across two or more images. In Active Color Imaging the model is used to determine the changes in the focal length required to compensate for the change in magnification introduced by the different focus distance motor settings used for each band and by chromatic aberration. In Constant Magnification Focusing the model is used to determine the changes in the focal length required to compensate for the magnification changes resulting from variations in the focussed distance of the lens.

Magnification model

The general form of our relative magnification model is

$$M = g(f_d, f_l)$$

where f_d and f_l are the focal length and focussed distance of the lens respectively. In our model the f_d and f_l values are specified in motor units. The relative image magnification between two lens positions is found by taking the ratio of the model at those two lens positions, ie.

$$relative\ image\ magnification = \frac{g(f_{d_1}, f_{l_1})}{g(f_{d_2}, f_{l_2})}$$

To mathematically model the relative image magnification we use a bivariate cubic polynomial [5]

$$u = a_{00} + a_{01}y + a_{02}y^2 + a_{03}y^3 + a_{10}x + a_{11}xy + a_{12}xy^2 + a_{20}x^2 + a_{21}x^2y + a_{30}x^3$$

where

u is the relative image magnification
x is the focus distance motor position
y is the focal length motor position, and
a_{ij} are polynomial coefficients determined by a calibration procedure.

Calibration of the bivariate cubic polynomial model is done by making a series of magnification measurements over a range of focus distance and focal length motor positions and then using these measurements as control points in a least squares fitting procedure. The magnification measurements are made by measuring the width of the bar on the calibration target shown in Figure 32. To calculate the width of the bar the positions of the white to

[5] We had initially tried using a closed-form optics model but we found that the calibration data could be more accurately followed by using a bivariate cubic polynomial.

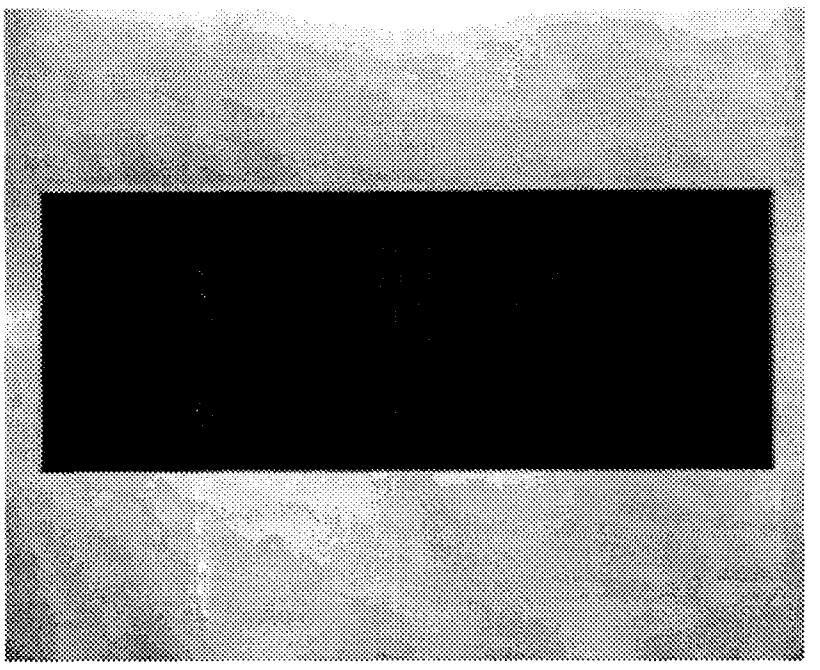


Figure 32: Calibration target for relative image magnification model

to black and the black to white edges are determined using the edge detection procedure described in appendix A. The target's width is then calculated as the difference in the two edge positions. To avoid the effects of scanline timing jitter and pixel resampling [1] width measurements are made along vertical image columns. The dimensions of the calibration target are chosen so that the horizontal bar occupies approximately 67% of the camera's vertical field of view at the point of highest magnification $(f_d, f_l) = (0,0)$, and approximately 10% of the vertical field of view at the point of lowest magnification $(f_d, f_l) = (3900, 3900)$.

One drawback in calculating image magnification from the width of a single fixed calibration target is the fact that as the lens parameters are changed the target's edges can move quite far across the image. As a result the positions of the target's edges and thus the target's width measurement can potentially be affected by local geometric distortions in the image. We make the assumption that the errors introduced by local geometric distortions are small compared to the changes in the global image magnification. We also make the assumption that any biases caused by the variation in the camera's point spread function across the sensor will also be relatively small.

The final calibration data set consists of target width measurements u taken at focus distance motor and focal length motor positions x and y. For our model we make measurements at 31 focal length positions and 31 focus distance positions for a total of 961 measurement points. Focal length motor positions are varied from 0 ($f_l \approx 75mm$) to 3900 ($f_l \approx 12.5mm$). Focussed distance motor positions are varied from 0 ($f_d = 1.2m$) to 3900 ($f_d = \infty$).

Using the pseudo inverse procedure described in [18] the least squares coefficients of best fit for the model are

$$
\begin{aligned}
a_{00} &= 6.7052 \\
a_{01} &= -2.8650 \times 10^{-3} \\
a_{02} &= 4.6365 \times 10^{-7} \\
a_{03} &= -2.4321 \times 10^{-11} \\
a_{10} &= -5.2330 \times 10^{-4} \\
a_{11} &= 2.4753 \times 10^{-7} \\
a_{12} &= -3.1141 \times 10^{-11} \\
a_{20} &= 1.6786 \times 10^{-8} \\
a_{21} &= -5.7644 \times 10^{-12} \\
a_{30} &= 6.3530 \times 10^{-13}
\end{aligned}
$$

Figure 33 contains a plot of the lens magnification relative to the position (0,0). Figure 34 contains a plot of the difference between the calibrated model and the calibration data, expressed as a percentage error relative to the calibrated data. The maximum error between the model and the calibration data is -3.1% at point $(3900, 3900)$, with the mean and standard deviations of the error being 0.0016% and 0.48% respectively.

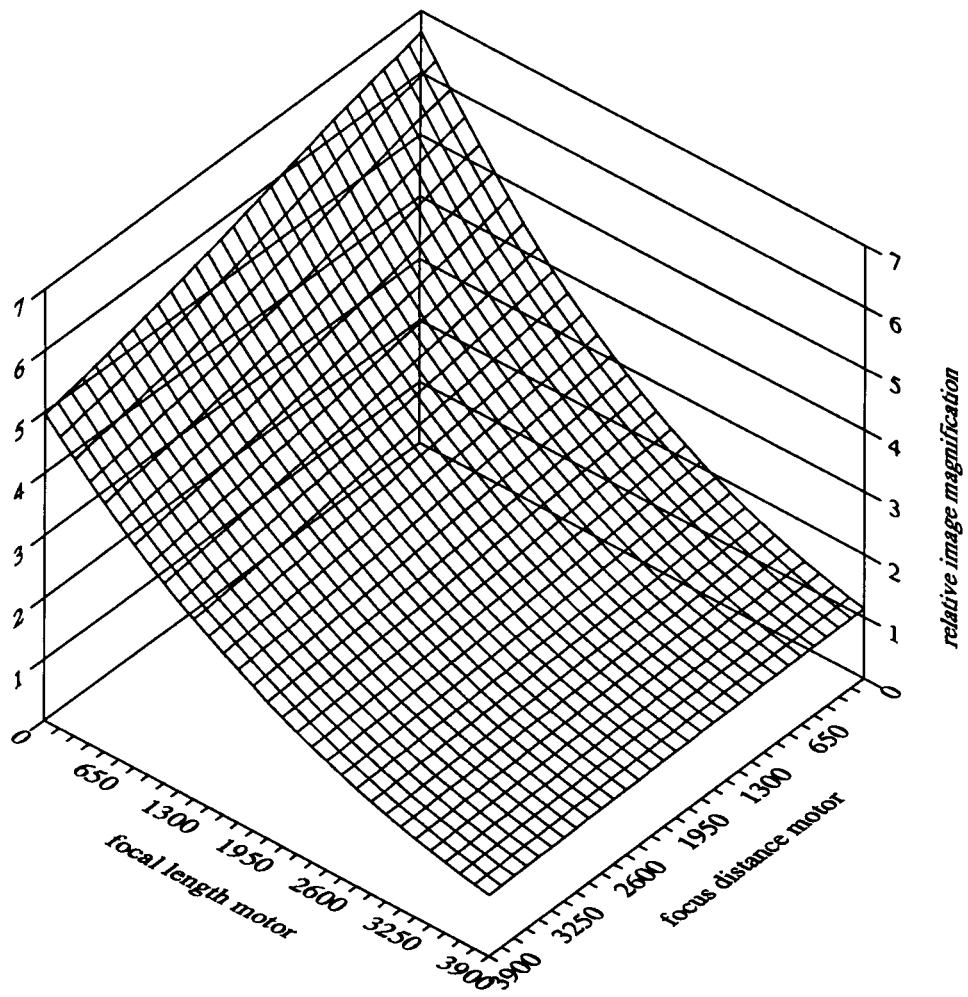


Figure 33: Relative image magnification model

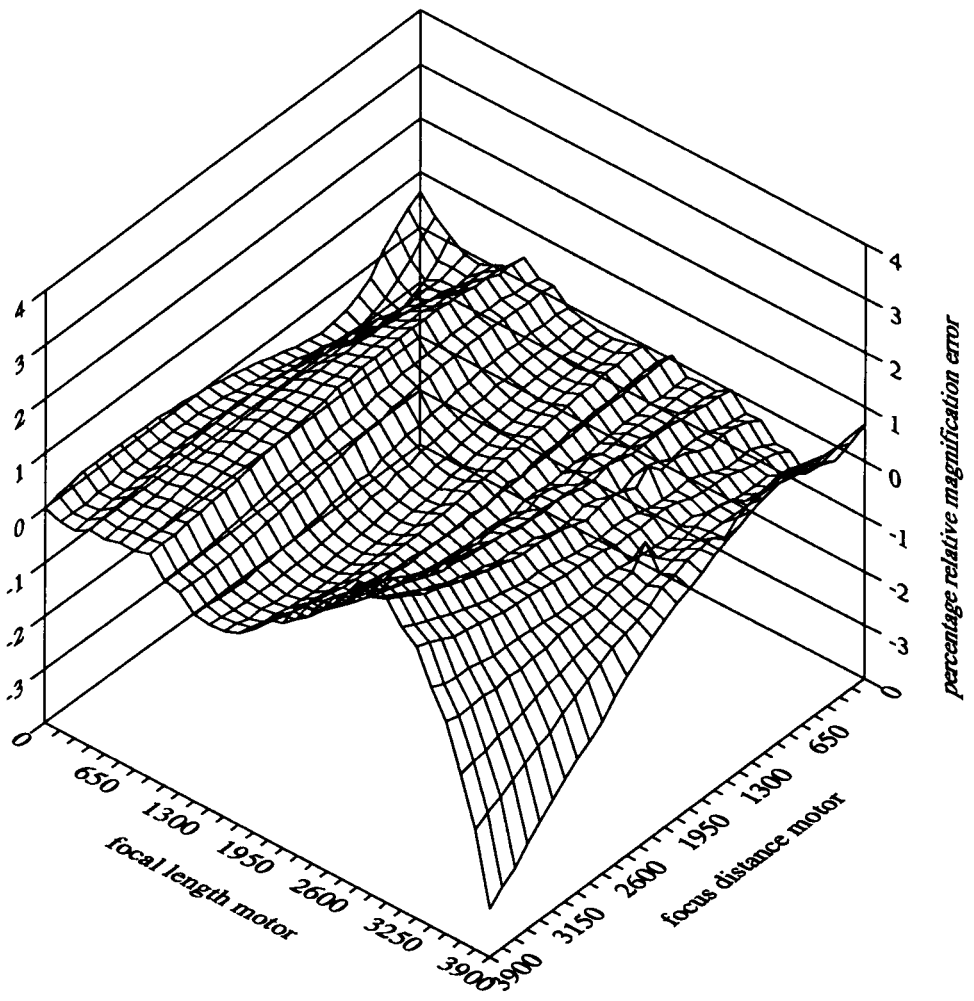


Figure 34: Percentage error in relative image magnification model

Using the model

In color imaging our objective is to take red, green and blue images with the same relative image magnification, ie.

$$g_r(f_{d_r}, f_{l_r}) = g_g(f_{d_g}, f_{l_g}) = g_b(f_{d_b}, f_{l_b})$$

With ACI f_{d_r}, f_{d_g} and f_{d_b} are determined by using an auto focusing routine or by using a calibrated focussed distance model. Generally f_{d_r}, f_{d_g} and f_{d_b} will not be the same because of longitudinal chromatic aberration. Fixing one of the focal length parameters (f_{l_b} in our implementation), the remaining two focal length parameters can be determined by searching for the settings that most closely satisfy the equality constraint on the relative image magnification.

In range from focus our objective is to take a series of images at varying focus distances while maintaining the same relative image magnification, ie.

$$g(f_{d_1}, f_{l_1}) = g(f_{d_2}, f_{l_2}) = ... = g(f_{d_n}, f_{l_n})$$

With CMF the focus distance parameters f_{d_1} ... f_{d_n} are determined by the application performing the focusing. After fixing one of the focal length parameters (say f_{l_1}) to establish a base image magnification, the remaining focal length parameters are then determined by searching for the parameter values that most closely satisfy the equality constraint on the relative image magnification. Effectively this causes the camera lens to operate on an isomagnification contour on the surface shown in Figure 33. This can be more clearly seen in Figure 35 where the lens motor positions have been plotted for one such contour.

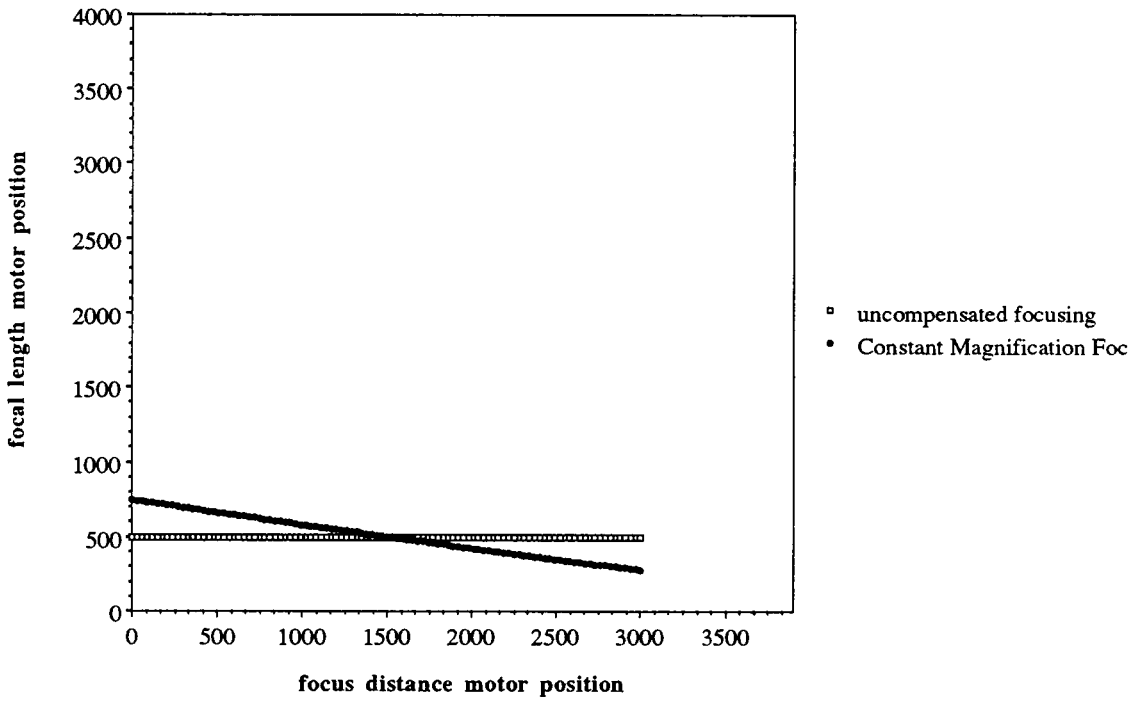


Figure 35: Constant Magnification Focusing isomagnification contour

8 Appendix C - Criterion function peak position detection

In both range from focus and autofocus we need to determine the position of the peak in a set of sharpness criterion function measurements taken over a range of focus distance positions. Typically the presence of measurement noise makes simple gradient searches over the criterion function values unreliable and inaccurate. One approach to determine the position of the peak of the criterion function reliably and accurately is to fit a model to the data and then find the position of the model's peak. Unfortunately the processes underlying the data are scene dependent and generally very difficult to model. For the calibration work described in this report we can simplify the problem somewhat by using targets that produce unimodal criterion function responses. Unfortunately the form of the unimodal curve still varies with the imaging parameters and is still difficult to determine. To get around this problem instead of trying to fit all of the data to a complex model we fit a subset of the data from focus positions immediately surrounding the peak to a simple quadratic model. The position of the peak in this locally fit model is then taken as the position of the criterion function peak.

The data and model

Our data consists of criterion function measurements y taken at known focus positions x. At each focus position a number of criterion function evaluations are made providing a mean and standard deviation for the criterion function at that focus position. So the data consists of N chosen focus positions

$$x = \{x_1, x_2, ..., x_N\}$$

having criterion function means

$$y = \{\bar{y}_1, \bar{y}_2, ..., \bar{y}_N\}$$

and standard deviations

$$\sigma = \{\sigma_1, \sigma_2, ..., \sigma_N\}$$

To model the relationship between x and y we use the second order polynomial

$$y = a_1 + a_2 x + a_3 x^2$$

Fitting a quadratic model using least squares linear regression

To find the coefficients of best fit for the quadratic model using least squares linear regression we need to solve the following matrix equation [12, 15]:

$$\begin{bmatrix} \Sigma\frac{1}{\sigma^2} & \Sigma\frac{x}{\sigma^2} & \Sigma\frac{x^2}{\sigma^2} \\ \Sigma\frac{x}{\sigma^2} & \Sigma\frac{x^2}{\sigma^2} & \Sigma\frac{x^3}{\sigma^2} \\ \Sigma\frac{x^2}{\sigma^2} & \Sigma\frac{x^3}{\sigma^2} & \Sigma\frac{x^4}{\sigma^2} \end{bmatrix} \begin{bmatrix} a_1 \\ a_2 \\ a_3 \end{bmatrix} = \begin{bmatrix} \Sigma\frac{y}{\sigma^2} \\ \Sigma\frac{xy}{\sigma^2} \\ \Sigma\frac{x^2y}{\sigma^2} \end{bmatrix}$$

Defining

$$S = \begin{bmatrix} \Sigma\frac{1}{\sigma^2} & \Sigma\frac{x}{\sigma^2} & \Sigma\frac{x^2}{\sigma^2} \\ \Sigma\frac{x}{\sigma^2} & \Sigma\frac{x^2}{\sigma^2} & \Sigma\frac{x^3}{\sigma^2} \\ \Sigma\frac{x^2}{\sigma^2} & \Sigma\frac{x^3}{\sigma^2} & \Sigma\frac{x^4}{\sigma^2} \end{bmatrix}$$

we can solve for the coefficients

$$
\begin{aligned}
a = \begin{bmatrix} a_1 \\ a_2 \\ a_3 \end{bmatrix} &= S^{-1} \begin{bmatrix} \Sigma \frac{y}{\sigma^2} \\ \Sigma \frac{xy}{\sigma^2} \\ \Sigma \frac{x^2 y}{\sigma^2} \end{bmatrix} \\
&= S^{-1} \begin{bmatrix} \frac{1}{\sigma_1^2} & \frac{1}{\sigma_2^2} & \cdots & \frac{1}{\sigma_N^2} \\ \frac{x_1}{\sigma_1^2} & \frac{x_2}{\sigma_2^2} & \cdots & \frac{x_N}{\sigma_N^2} \\ \frac{x_1^2}{\sigma_1^2} & \frac{x_2^2}{\sigma_2^2} & \cdots & \frac{x_N^2}{\sigma_N^2} \end{bmatrix} \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_N \end{bmatrix} \\
&= My \qquad (1)
\end{aligned}
$$

Finding x_{max} and its variance $\sigma^2_{x_{max}}$

To find the position of the peak of the fitted quadratic curve we set the first derivative of the model to zero,

$$\frac{\partial y}{\partial x} = a_2 + 2a_3 x = 0$$

and then directly solve for x_{max},

$$x_{max} = -\frac{a_2}{2a_3}$$

To determine the variance of the position of the peak of the data we first calculate the Jacobian of x_{max} wrt a:

$$j = \frac{\partial x_{max}}{\partial a} = \begin{bmatrix} \frac{\partial x_{max}}{\partial a_1} \\ \frac{\partial x_{max}}{\partial a_2} \\ \frac{\partial x_{max}}{\partial a_3} \end{bmatrix} = \begin{bmatrix} 0 \\ -\frac{1}{2a_3} \\ \frac{a_2}{2a_3^2} \end{bmatrix} \qquad (2)$$

From equation 1 we also have the linear transformation between a and y

$$\frac{\partial a}{\partial y} = M \qquad (3)$$

Thus, from equations 2 and 3 we can define a linear transformation between our calculated peak position x_{max} and y

$$
\begin{aligned}
d^T = \frac{\partial x_{max}}{\partial y} &= \frac{\partial x_{max}}{\partial a_1}\frac{\partial a_1}{\partial y} + \frac{\partial x_{max}}{\partial a_2}\frac{\partial a_2}{\partial y} + \frac{\partial x_{max}}{\partial a_3}\frac{\partial a_3}{\partial y} \\
&= j^T M \qquad (4)
\end{aligned}
$$

If we assume the measurements y_i to be independent, with covariance matrix P_{yy}

$$P_{yy} = \begin{bmatrix} \sigma_1\sigma_1 & 0 & \cdots & 0 \\ 0 & \sigma_2\sigma_2 & & \\ \vdots & & \ddots & \vdots \\ 0 & & \cdots & \sigma_N\sigma_N \end{bmatrix}$$

then $\sigma^2_{x_{max}}$ can be approximately determined using the covariance matrix and the linear transformation defined in equation 4

$$\begin{aligned}\sigma^2_{x_{max}} &\approx d^T P_{yy} d \\ &\approx \sum_{i=1}^{N} d_i^2 \sigma_i^2\end{aligned}$$

As j (and thus d) represents a linearization of $\frac{\partial x}{\partial a}$ at x_{max}, the value of $\sigma^2_{x_{max}}$ will only be an approximation.

Choosing a subset of data

In our approach instead of trying to fit data from the full range of focus positions to a complex model of the criterion function, we fit the data from a smaller subrange of focus positions immediately surrounding the apparent peak to a simpler quadratic function. The assumption here is that at least locally the criterion function behaves as a symmetric unimodal function corrupted by measurement noise. Modeling this with a quadratic function provides a computationally efficient means of estimating the position of the criterion function's peak and its variance.

Naturally the more data that is used in fitting the model, the more accurate our estimation of the peak's position will be. To determine just how much data we can use from around the peak we use a chi-squared figure of merit Q [12]. Q gives a measure of the goodness-of-fit between a set of data and a model. The test is defined as:

$$\chi^2 = \sum_{i=1}^{N} \left(\frac{y_i - y(x_i; a_1, a_2, a_3)}{\sigma_i} \right)^2$$

$$Q = Q(\chi^2 \mid \nu) = \Gamma(\frac{\nu}{2}, \frac{\chi^2}{2})$$

where Γ is the incomplete gamma function and ν is the number of degrees of freedom ($N-3$ for our model). Q is the probability that the observed chi-square will exceed the value χ^2 by chance, *even* for a correct model. Values of $Q < 0.001$ indicate a poor fit between the model and the data, with the most probable reason for the poor model fit being that the model is wrong for the data. Values of $Q > 0.001$ indicate a reasonable fit between the data and the model.

To determine the subset of data to be used in fitting the quadratic model we start by fitting the full data set to the model using the least squares linear regression procedure. This gives us an estimated peak position x_{max} and a value for Q. If Q is > 0.01 we're done. If Q is < 0.01 we drop the data point that is furthest from the estimated peak position and repeat the process.

This data fitting approach is easily adapted to an auto focus algorithm by having the camera system replace the farthest data point in the data set by a new measurement that is taken at a focus position that is closer to the estimated peak. Again the approach is continued until the fit between the working set of data and the quadratic model is acceptable.

In an automated focus algorithm the data fitting approach has several advantages over more conventional gradient search schemes. For one, the peak position is estimated from a set of measurements and their variances taken over a range of focus positions, rather than

from a sequence of decisions based on measurements made at single focus positions. This greatly improves the accuracy and robustness of the peak position detection in the presence of measurement noise. Another advantage of the data fitting approach is the confidence interval that it gives for the estimated peak position.

Peak detection example

Figure 36 shows the detected peak in a set of Tenengrad criterion function data collected for a fairly narrow aperture (540 motor units). For each data point the mean and 1 σ error tolerances are shown. The full data set contains 101 points, of which only 29 were used to fit the final curve. The final figure of merit Q for the curve was 0.125. The estimated variance of the peak position of 1426 motor units was 4 motor units.

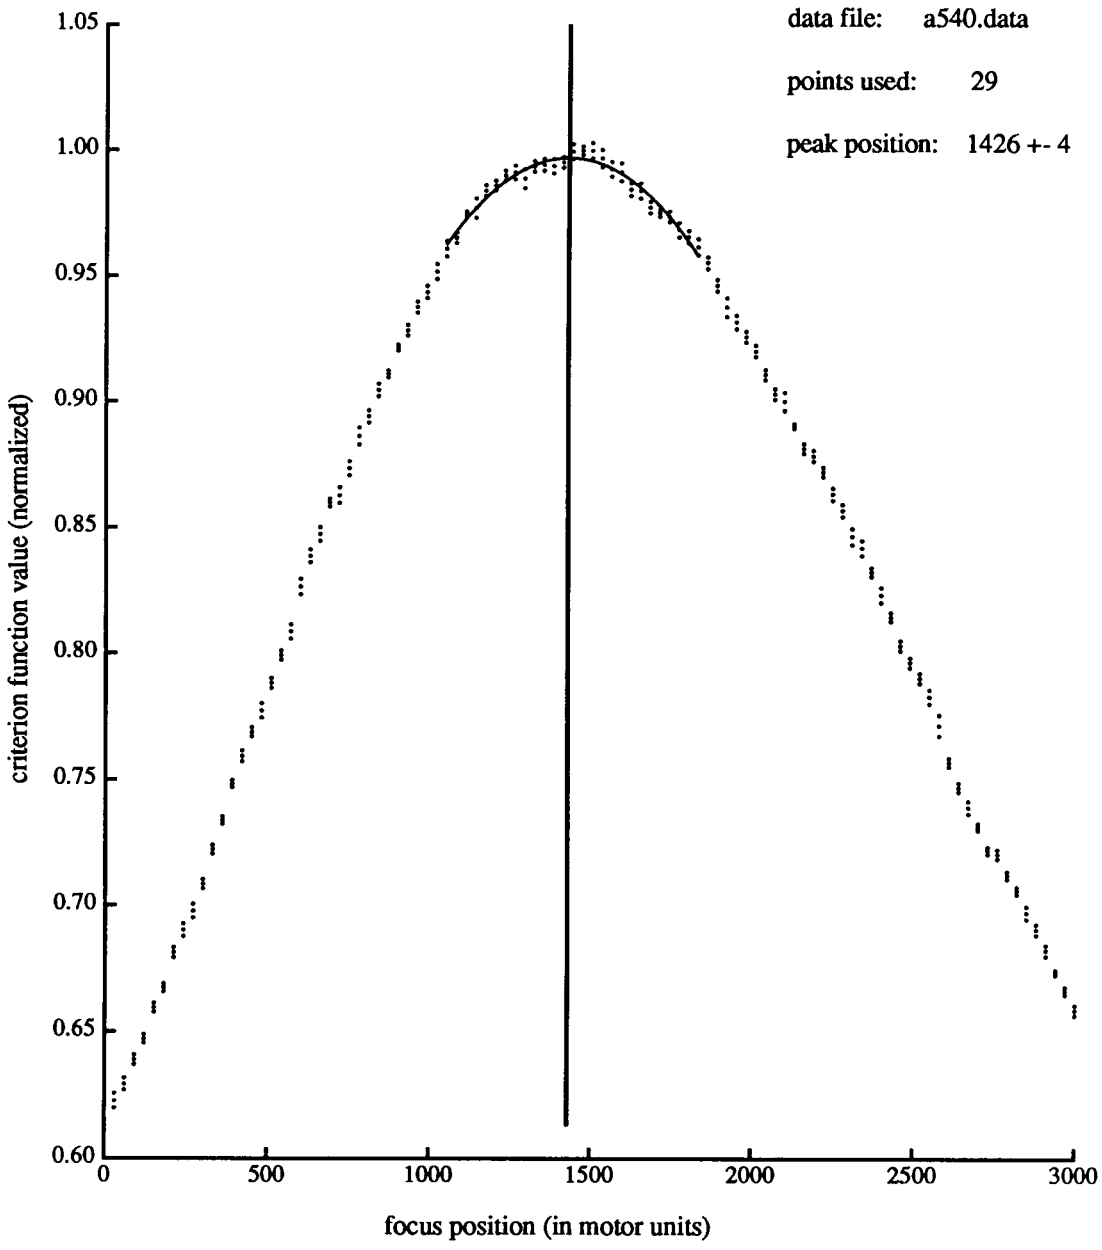


Figure 36: Peak detection in criterion function data

References

[1] H. A. Beyer. Linejitter and geometric calibration of ccd-cameras. *ISPRS Journal of Photogrammetry and Remote Sensing*, 45(1):17–32, 1990.

[2] T. E. Boult and G. Wolberg. Correcting chromatic aberrations using image warping. In *Proceedings of Image Understanding Workshop*, San Diego, California, January 1992. DARPA.

[3] J. S. Chen and G. Medioni. Detection, localization, and estimation of edges. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(2):191–198, February 1989.

[4] T. Darrell and K. Wohn. Pyramid based depth from focus. In *Proceedings of Computer Vision and Pattern Recognition*, pages 504–509, Ann Arbor, Michigan, June 1988. Computer Society, IEEE Computer Society.

[5] B. Funt and J. Ho. Color from black and white. *International Journal of Computer Vision*, 3(2):109–117, June 1989.

[6] G. J. Klinker. *A Physical Approach to Color Image Understanding.* PhD thesis, Carnegie Mellon University, May 1988.

[7] E. P. Krotkov. Exploratory visual sensing for determining spatial layout with an agile stereo camera system. PhD Thesis MS-CIS-87-29, University of Pennsylvania, April 1987.

[8] R. K. Lenz and R. Y. Tsai. Techniques for calibration of the scale factor and image center for high accuracy 3d machine vision metrology. In *Proceedings of IEEE International Conference on Robotics and Automation*, pages 68–75, Raleigh, NC, March 1987.

[9] H. N. Nair and C. V. Stewart. Depth from focus. Technical Report 90-10, Rensselaer Polytechnic Institute, 1990.

[10] S. K. Nayar and Y. Nakagawa. Shape from focus: An effective approach for rough surfaces. In *Proceedings of IEEE International Conference on Robotics and Automation*, pages 218–225, Cincinnati, OH, May 1990.

[11] C. L. Novak, S. A. Shafer, and R. G. Willson. Obtaining accurate color images for machine vision research. In *Proceedings of Conference on Perceiving, Measuring, and Using Color*, Santa Clara, February 1990. SPIE.

[12] W. H. Press, B. P. Flannery, S. A. Teukolsky, and W. T. Vetterling. *Numerical Recipes in C.* Cambridge University Press, New York, New York, 1986.

[13] C. C. Slama, C. Theurer, and S. W. Henriksen. *Manual of Photogrammetry.* American Society of Photogrammetry, Falls Church, VA, fourth edition, 1980.

[14] W. J. Smith. *Modern Optical Engineering, The Design of Optical Systems.* Optical and Electro-Optical Engineering Series. McGraw-Hill, New York, 1966.

[15] G. Strang. *Linear Algebra and Its Applications.* Harcourt Brace Jovanovich, San Diego, 1988.

[16] A. G. Wiley. *Metric Aspects of Zoom Vision.* PhD thesis, University of Illinois at Urbana-Champaign, June 1991.

[17] R. G. Willson and S. A. Shafer. Active lens control for high precision computer imaging. In *Proceedings of IEEE International Conference on Robotics and Automation*, volume 3, pages 2063–2070, Sacramento, CA, April 1991.

[18] G. Wolberg. *Digital Image Warping.* IEEE Computer Society Press, Los Alamitos, California, 1990.