Computational approach for depth from defocus

Ovidiu Ghita*
Paul F. Whelan
John Mallon
Vision Systems Laboratory
School of Electronic Engineering
Dublin City University
Dublin 9, Ireland
E-mail: ghitao@eeng.dcu.ie

Abstract. *Active depth from defocus (DFD) eliminates the main limitation faced by passive DFD, namely its inability to recover depth when dealing with scenes defined by weakly textured (or texture-less) objects. This is achieved by projecting a dense illumination pattern onto the scene and depth can be recovered by measuring the local blurring of the projected pattern. Since the illumination pattern forces a strong dominant texture on imaged surfaces, the level of blurring is determined by applying a local operator (tuned on the frequency derived from the illumination pattern) as opposed to the case of window-based passive DFD where a large range of band pass operators are required. The choice of the local operator is a key issue in achieving precise and dense depth estimation. Consequently, in this paper we introduce a new focus operator and we propose refinements to compensate for the problems associated with a suboptimal local operator and a nonoptimized illumination pattern. The developed range sensor has been tested on real images and the results demonstrate that the performance of our range sensor compares well with those achieved by other implementations, where precise and computationally expensive optimization techniques are employed.* © 2005 SPIE and IS&T.
[DOI: 10.1117/1.1900743]

1 Introduction

Pentland[1] pointed out that the range information is not lost during the process of image formation as the objects are imaged according to their position in space. In this way, the objects situated along the surface where the image is in focus are accurately imaged, while others, not placed close to this surface are blurred. It is important to note that the level of blurring is in direct relation to the distance between the surface where the image is in focus and the actual spatial position of the object under investigation. Thus, by comparing several images captured with different focal levels (obtained by changing either the aperture of the lens or the internal parameters of the camera) we can estimate the depth for each point in the scene by analyzing the local blurring.

As opposed to depth from focus (DFF)[2–5] where the depth is estimated by taking a large number of images by incrementing the focal settings in small steps, depth from defocus (DFD) requires only two differently focused images to estimate the depth information.[6–10] This is a major advantage when dealing with dynamic scenes where the scene objects may change their spatial position during the image acquisition process. Furthermore, instead of searching for the best focused point in the image stack as is the case with DFF, the depth in DFD can be computed by evaluating the blurring difference between each point in the defocused images. Also it is worth noting that the ranging methods based on focus/defocus are less affected by occlusions or missing parts than the ranging techniques based on triangulation or stereo vision since the images to be analyzed are only differently focused.[11]

Historically, DFD methods have evolved as a passive range sensing strategy.[5,8,12,13] In general, passive DFD attempts to estimate the blurring level by applying a large range of narrow-band operators[5] since the image blurring varies with texture frequencies.[12] A different implementation has been proposed by Rajagopalan and Chaudhuri[14] where they applied a Markov random field model to improve the initial depth estimates obtained from a window-based DFD scheme. More recently Deschenes *et al.*[15] proposed a new algorithm to extract the blur difference between two defocused images by fitting the defocused images by Hermite polynomials. In this way the coefficients of the Hermite polynomial computed from the more blurred image are a function of the partial derivatives of the other image and the blur difference. Other recent contributions to passive DFD include the work of Bhasin and Chaudhuri[16] and Favaro *et al.*[17] However the main disadvantage of the passive DFD approaches is the fact that they are computationally intensive and they return unreliable depth estimates when dealing with weakly or nontextured image areas.

To address this limitation Pentland *et al.*[18] suggested projecting a structured light onto the scene and estimating the depth by analyzing the level of blurring associated with the projected pattern. The results proved to be accurate although obtained at a relatively coarse spatial resolution. Later, Nayar *et al.*[19] argued that optimizing the illumination pattern and the focus operator can lead to high density depth maps. They developed a symmetrical pattern organized as a rectangular grid optimized for a specific camera.

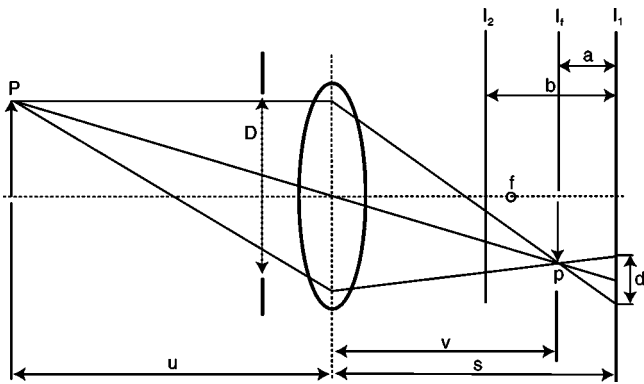*Author to whom correspondence should be addressed.

Fig. 1 The image formation process. The depth *u* is a function of the sensor position *s*, lens aperture *D*, focal length *f* and the blur patch *d* (See Refs. 9 and 19).

Fig. 2 The focus operator. (a) Standard Laplacian. (b) Four peak Laplacian operator.

Then they optimized the Laplacian operator in order to obtain a narrow band operator. The reported results indicate the efficiency of this approach but it is worth noting that in their implementation the illumination pattern has to be registered with the sensing elements at a subpixel resolution, a fact that makes this approach difficult to apply in practice.

In this paper we describe the implementation of a real-time active DFD range sensor, where special emphasis is placed on the focus operator and the image refinements employed in order to alleviate the problems caused by arbitrary object textures and a nonoptimized illumination pattern.

2 Active Depth from Defocus. Related Research

A range sensor based on focus error and structured light has been proposed by Pentland *et al.*[18] and Girod and Scherock.[20] This approach extends the passive range sensor developed by Pentland[1] where the large aperture camera was replaced by a structured light source (for more details see also Ref. 21). Since the camera's lens has a small aperture, its depth of field is significantly larger than the depth of field of the structured light. They employed an illumination pattern consisting of evenly spaced vertical lines. Since the position of the pattern is known *a priori* and using the fact that the width of the light stripe gets larger when defocused, the depth can be easily estimated by measuring the spread of the defocused line. In spite of simplicity this approach proved to be relatively accurate. The major limitation of this approach is the coarse-spaced illumination pattern and as a direct consequence the resulting depth map is low resolution.

In order to address this limitation Nayar *et al.*[19] developed an active DFD range sensor consisting of two sensing elements separated by a known distance *b* used in conjunction with a dense optimized illumination pattern.[22] In this way, one of the sensing elements will capture a near focused image while the other will capture the far focused image (see Fig. 1). The illumination pattern was projected onto the scene in order to force an artificial texture on all imaged areas. The depth is in direct relation to the relative level of blurring present in both images which is measured by filtering the near and far focused images with a local operator such as the Laplacian.[18,19] Since our goal is
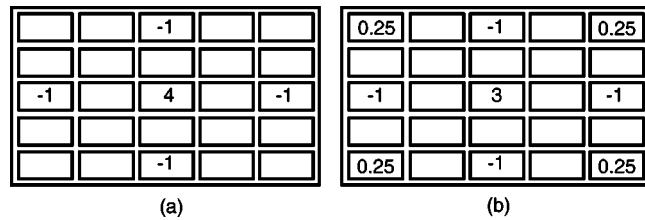
achieving dense depth maps in our implementation we used the latter approach.

3 The Blur Function

If the object to be imaged is placed in or very close to the surface of best focus (object point is in the position *P* and the sensing element is placed at I_f, see Fig. 1), the image formed on the sensing element is sharp while each object point is imaged by the lens into a point on the sensor plane. Conversely, if the object is shifted from the surface where the image is in focus, the object points are distributed over a patch on the surface of the sensing element, where the diameter of the patch indicates the level of blurring.

The blurring effect can be thought of as a convolution between the perfectly focused image and a blurring function called point spread function (PSF). In vision literature various models have been proposed to approximate the blurring function[10,23,24] but in practice the two-dimensional Gaussian[1,9,25] has been widely employed to approximate the PSF when paraxial geometric optics are used and diffraction effects are negligible.

The standard deviation (or the spread parameter) of the Gaussian operator is the parameter of interest as it indicates the level of blurring contained in the defocused images (the larger the level of blurring the larger the value of the standard deviation). Since the PSF function approximates a low pass filter, to extract the level of local blurring (i.e., determine the standard deviation of the PSF) it would be necessary to extract the high frequency information derived from the scene. This is achieved by convolving the near and far focused images with a local focus operator, where the output indicates the local blurring level.

However, the earlier-mentioned approach returns reliable results only if the scene under investigation is highly textured. To eliminate this restriction a solution is to project a structured light onto the scene, thus forcing a dominant artificial texture on all visible surfaces.

The structured light should have a symmetrical or semi-symmetrical arrangement in order to achieve rotational invariance. We can recall that the near and far focused images are captured with different focal settings and as a consequence a variation in magnification between these images will be noticed. As in our implementation the magnification changes between the defocused images cannot be alleviated on an optical basis (for details refer to Sec. 6) this issue introduces a new challenge as we cannot perform a registration between the illumination pattern and the pixel elements of the complementary metal–oxide–semiconductor (CMOS) cameras. Perfect registration between the illumination pattern and camera's pixels is quite difficult to
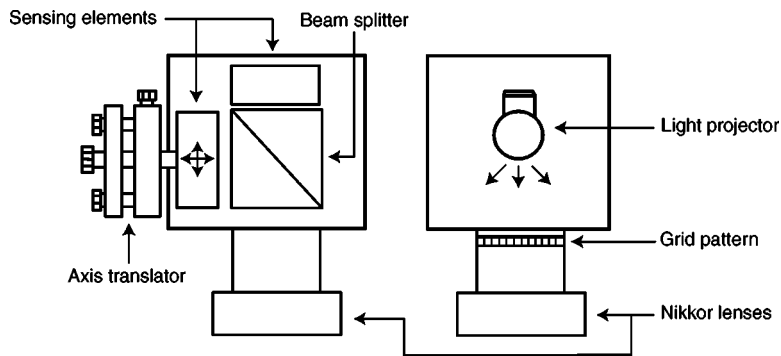
Fig. 3 The diagram of the developed range sensor.

achieve in practice as it would require specialized equipment to construct a custom grating filter, and in addition this solution would be effective only if the magnification were maintained at a constant level for both near and far focused images.

Fortunately, the depth errors caused by misregistrations between the illumination pattern and the camera's pixels are very small when compared with errors introduced by the focus operator, magnification changes, and the imperfections of the optical and sensing equipment (the procedure employed to compensate for the nonlinear response of the CMOS sensors is detailed in Sec. 6). Thus in our implementation we relaxed the requirement for an optimized illumination pattern. To achieve high resolution depth estimation, in our implementation we have used a simple illumination pattern defined by a sequence of horizontal stripes with a density of 10 lines per millimeter. Our efforts were concentrated on the development of a new focus operator that can be easily tuned on the spatial arrangement of the illumination pattern.

4 Focus Operator

The problem of recovering the local blurring is greatly simplified in active DFD since the scene has a dominant frequency, namely the frequency associated with the illumination pattern. Thus, the focus operator has to be designed in order to respond strongly to this frequency. When the illumination pattern is projected onto a blank sheet of white paper the projected pattern consists of evenly spaced dark and bright horizontal lines, where the period is 6 pixels (projector elevation 71 cm from the base line, fitted with a 60 mm lens). Since the illumination pattern has a symmetric arrangement, the focus operator also has to be symmetric and must be immune to direct current (dc) components. The most common focus operator is the Laplacian where the size of the kernel is dependent on the spatial arrangement of the illumination pattern (5×5 for the present illumination pattern). Although the Laplacian has sharp peaks at the frequency derived from the illumination pattern, it also enhances the features associated with the scene's tex-
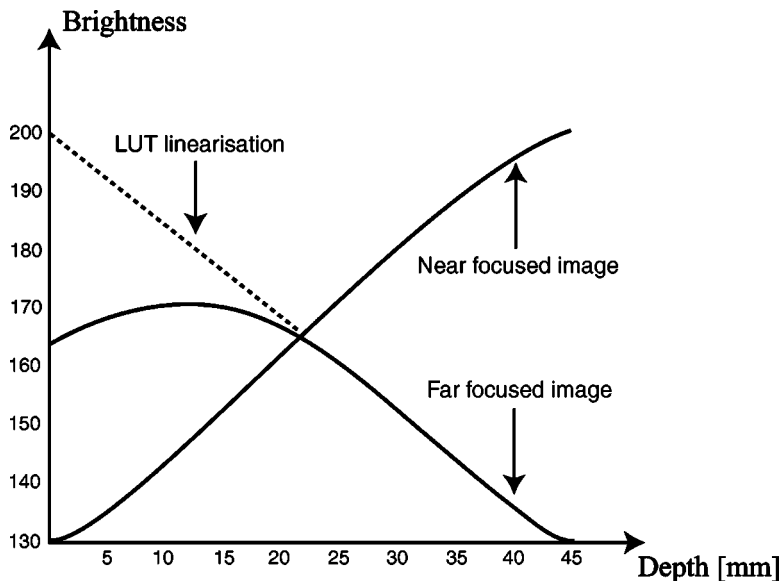

Fig. 4 The effect of the supplementary blur introduced by the lens of the light projector. The errors are compensated by using a look-up table linearization. Numerical values are obtained when the simple cell was employed as focus operator.
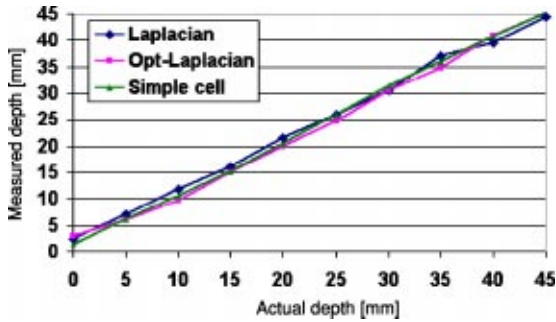
Fig. 5 Recovered depth for a textureless planar object placed at different elevations from the base line of the workspace.
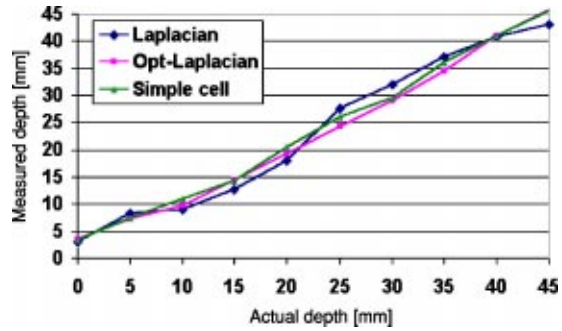

Fig. 6 Recovered depth for a randomly textured planar object placed at different elevations from the base line of the workspace.

ture which alters the local blurring measurements. To alleviate this problem Nayar *et al.*[19] employed a frequency analysis approach to develop a narrow-band Laplacian with four sharp peaks at the frequency derived from an optimized illumination pattern. In Fig. 2 the kernels of the 5×5 Laplacian operator and the 5×5 four-peak narrow band operator are depicted.

Taking into account that the illumination pattern forced on the scene is organized as a sequence of evenly spaced light stripes, this motivates us to introduce a new focus operator to estimate the local blurring, namely the simple cell.[26]

The relationships that implement the simple cell operator are illustrated in Eqs. (1)–(3):

$$s(x,y) = e^{-(x'^2 + y'^2)/2\sigma^2} \cos\left(\frac{2\pi}{T} x' + \varphi\right), \qquad (1)$$

$$x' = x \cos(\theta) - y \sin(\theta), \qquad (2)$$

$$y' = x \sin(\theta) + y \cos(\theta), \qquad (3)$$

where T represents the period, σ is the standard deviation of the Gaussian filter, θ specifies the orientation of the normal to the illumination pattern, and φ is the phase offset. There are various psychophysical experiments which indicate that the simple cell operator acts as a line or edge detector, by responding to lines or edges with a specific orientation and spatial frequency.[26,27] For other texture orientations the simple cell will respond weakly, and this will result in a decreased sensitivity as compared to the Laplacian operator when applied to arbitrary object textures.

Therefore, the properties of this operator are very attractive for our application since the illumination pattern is defined by a periodic arrangement with a well defined orientation. In our implementation, the following values are used to tune the simple cell operator on the projected illumination pattern: $2\pi/T = 1.5$, $\sigma^2 = 2$, $\theta = \pi/2$, and $\varphi = \pi/2$. The resulting filter implements an antisymmetric oriented derivative operator and the elements of the kernel are adjusted in order to ensure that their sum is equal to zero (to achieve insensitivity to dc components).

In order to assess the efficiency of this new focus operator we evaluated its performance compared to that offered by the Laplacian operator and the narrow-band operator.[19]

5 Image Refinements

Since the focus operator has a finite support (defined by 5×5 masks) it will generate windowing errors when it is applied to the near and far focused images. As expected, the image distortions inserted by the focus operator are more severe around the transitory regions between the dark and bright light stripes. This is caused by the imperfections in construction of the filter employed to generate the illumination pattern, i.e., the transparent and opaque regions of the projection filter are not perfectly defined. Given that the central part of the illumination stripe is less affected by the errors introduced by the focus operator and the illumination pattern (and has the highest intensity values), we normalized each stripe by vertically propagating the value of the pixel positioned on the center of the stripe. It is important to note that this stripe normalization procedure does not affect the local blurring level since the illumination pattern is dense and the resulting stripes are only 3–4 pixels wide and the blurring is assumed to be constant in small neighborhoods.

However, the focus operator and the imperfections of the illumination pattern were not the only source of errors. Given that the near and far focused images are captured with different camera settings, a variation in image magnification (which is dependent on the spatial position of the imaged object) occurs and as a direct consequence the stripes contained in the near and far focused images do not match perfectly together. This forced researchers to either implement computationally intensive techniques such as image registration and warping[28] or to address this problem on an optical basis.[29] In our implementation we compensate for this issue by employing image interpolation. While the dark stripes of the illumination pattern do not reveal any useful information and the spatial shift induced by magnification changes is smaller than half of the period of the illumination pattern, we propose to map them by vertically interpolating the adjacent bright illumination stripes. Taking into consideration that the illumination pattern is very dense, linear interpolation proved to be sufficient. The experiments indicate that the performance of the sensor significantly improved after the application of these image refinements.
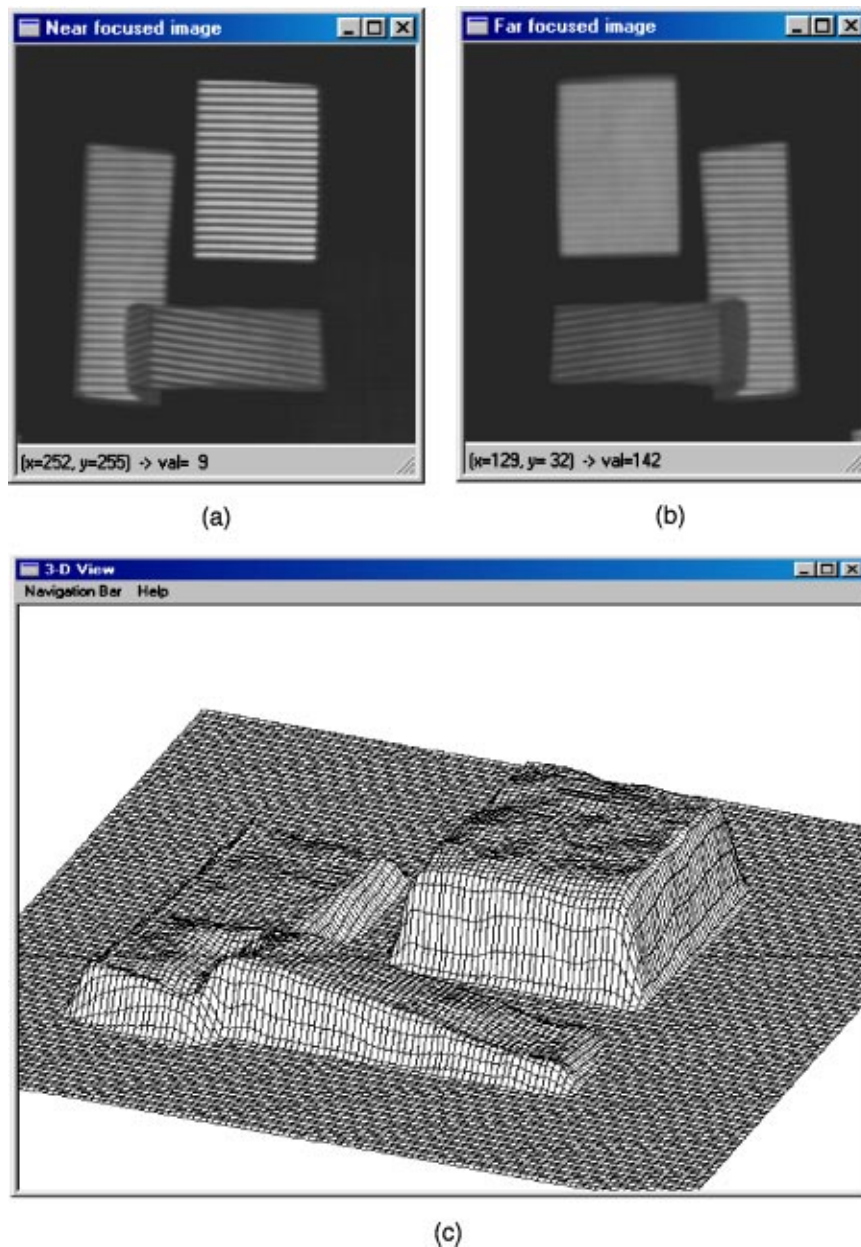
Fig. 7 Depth estimation for a scene defined by polyhedral objects. (a) Near focused image. (b) Far focused image. (c) Recovered depth.

6 Sensor Implementation

The developed sensor consists of two distinct parts, namely the sensing devices and the light projector. To capture the near and far focused images at the same time the sensor uses a beam splitter to separate the original image into two identical images. To capture the near and far focused images, one sensor is set in contact with the beam splitter while the second is positioned with a small gap (approximately 0.8 mm) from the beam splitter surface. The registration between the sensing elements is carried out by using a multiaxis translator which is attached to one of the sensing elements. Figure 3 illustrates the components of the developed range sensor.

The structured light is projected onto the scene using a MP-1000 projector fitted with a MGP-10 Moire grating (stripes with density of 10 lines per millimeter). The system uses two AF MICRO Nikkor 60 mm lenses, where one is used to image the scene while the other is attached to the light projector. The aperture of the lens attached to the light projector should be very small in order to obtain a lens with a large depth of field. Therefore, the illumination pattern projected onto the scene will be nonblurred and defocus will be introduced only by the focal settings of the sensing elements. On the other hand, a pinhole aperture will contribute to a severe reduction in illumination level arriving at the sensing elements. To compensate for this issue we need to employ a very powerful source of light, a solution difficult to apply in practice due to safety considerations. Since our light projector is fitted with a 50 W incandescent bulb, this approach is not feasible. Thus, we set the aperture of
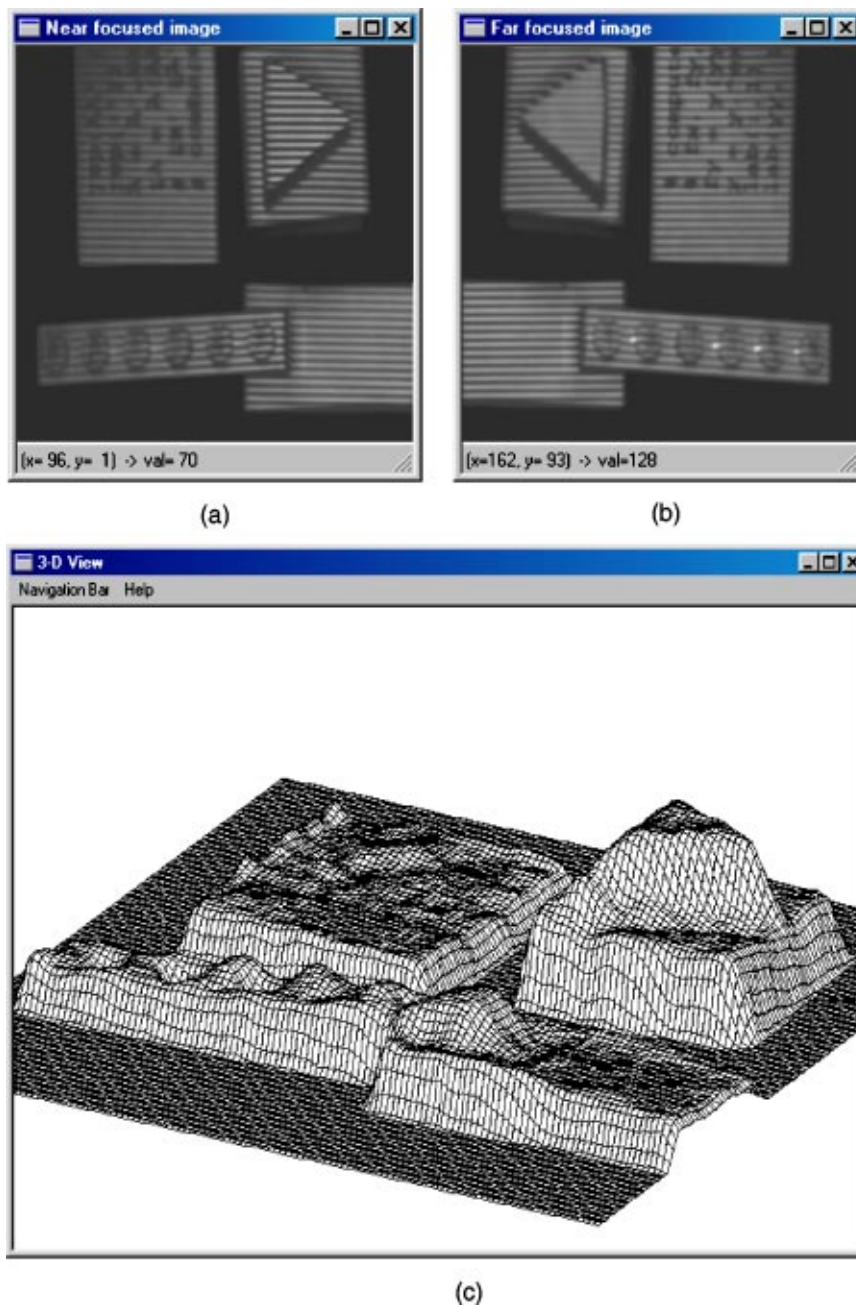
Fig. 8 Depth estimation for a scene defined by textureless, textured, and mildly specular objects. (a) Near focused image. (b) Far focused image. (c) Recovered depth.

the lens at the minimum value (2.8 setting) that assures a sufficient level of light to image the scene objects irrespective of their color. Nevertheless, in this situation the illumination pattern was supplementary defocused. To alleviate this problem we set the surface of best focus of the projected illumination pattern at the same position with the surface of best focus of the near focused sensing element. Using this approach, the level of blurring in the near focused image is almost linear with depth. On the other hand, the level of blurring in the far focused image will be disturbed due to the attenuation of the illumination pattern. This problem can be observed in Fig. 4 where the intensity output of the near and far focused images after the applica-

tion of the focus operator is plotted against depth.

This generates errors when dealing with far situated objects with respect to the sensor's position. To compensate for this problem the blurring profile of the far focused sensor is linearized in agreement with the blurring profile of the near focused sensor. The linearization procedure is implemented using a look-up table where the depth is estimated directly from the intensity outputs of the near and far focused image after the application of the focus operator.

7 Experiments and Results

In this paper our aim is to evaluate how the focus operator affects the overall performance of the range sensor. To
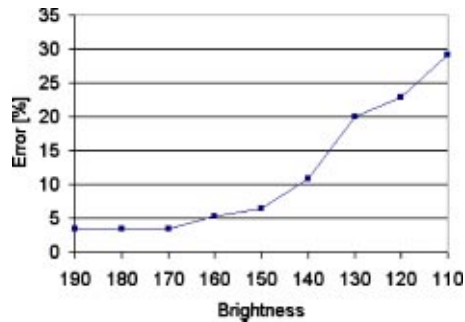
Fig. 9 Relation between the depth error and the brightness of the object surface.

achieve this goal, the range sensor was tested initially on textureless scenes then on scenes defined by arbitrary textured objects.

The relative accuracy was estimated for successive measurements and was defined by the maximum error between the real and estimated depth values contained in the depth map. During operation the sensor is placed at a distance of 86 cm above the base line of the workspace. Figure 5 illustrates the effect of the focus operator when the sensor was tested on a simple scene defined by a planar textureless object which is placed at different elevations from the base line of the workspace. As expected since there is no additional texture to disturb the illumination pattern, the depth is estimated almost similarly irrespective of the choice of the focus operator. However, when the sensor was tested on textured scenes, the experimental results indicated that the Laplacian operator cannot reject the influence of the object texture while the four peak focus operator and the simple cell are more robust to arbitrary texture (see Fig. 6).

Our results are similar with those reported by Nayar *et al.*[19] when the four peak Laplacian was employed as focus operator. Also it can be noticed that the depth estimation is less precise for objects situated at distances close to the calibration point where the depth values are over determined. Figures 7 and 8 depict additional results when the sensor was applied to various scenes.

In line with other active techniques, this approach returns unreliable depth estimation when it is applied to highly specular scenes or scenes defined by objects with very dark surfaces. Figure 9 illustrates how the accuracy is affected when the sensor was applied to scenes defined by objects with different surface colors. Figure 10 indicates the
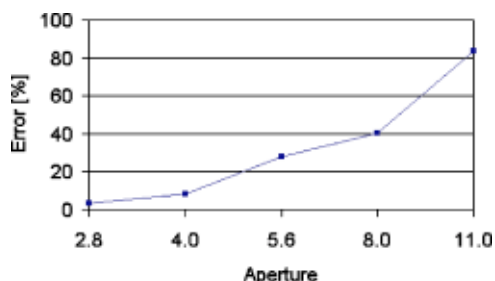

Fig. 10 Relation between the depth estimation and the level of illumination.

performance of the sensor when the illumination level of the light projector is reduced by changing the lens aperture.

8 Conclusions

In order to achieve accurate and dense depth estimation using active DFD, we have to address a large number of problems including mechanical, optical and computational. While the physical implementation of this sensor has been previously detailed,[30] in this paper we place the emphasis on the computational components. To robustly extract the relative blurring between two images captured with different focal settings, we have to confront problems such as sensitivity of the focus operator to the object texture and the variation in image magnification. In order to achieve insensitivity to object texture we developed a focus operator that responds strongly to the frequency derived from the illumination pattern. The problems associated with the variation in image magnification were addressed by employing image interpolation. All these components were included in the implementation of a real-time active DFD range sensor which was successfully applied in the development of a vision sensor for robotic bin-picking.[30]

References

1. A. Pentland, "A new sense for depth of field," *IEEE Trans. Pattern Anal. Mach. Intell.* **9**, 523–531 (1987).
2. P. Grossman, "Depth from focus," *Pattern Recogn. Lett.* **5**, 63–69 (1987).
3. E. Krotkov, "Focusing," *Int. J. Comput. Vis.* **1**, 223–237 (1987).
4. I. Nourbakhsh, D. Andre, C. Tomasi, and M. Genesereth, "Obstacle avoidance via depth from focus," *Proc. of Image Understanding Workshop (IUW 96)*, pp. 1339–1344 (1996).
5. Y. Xiong and S. A. Shafer, "Depth from focusing and defocusing," *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition*, pp. 68–73 (1993).
6. N. Asada, H. Fujiwara, and T. Matsuyama, "Seeing behind the scene: analysis photometric properties of occluding edges by the reversed projection blurring model," *IEEE Trans. Pattern Anal. Mach. Intell.* **20**, 157–166 (1998).
7. W. N. Klarquist and W. S. Geisler, "Maximum likelihood depth from defocus for active vision," *Proc. of the IEEE Conf. on Intel. Robots and Systems* **3**, 374–379 (1995).
8. A. Pentland, T. Darrell, M. Turk, and W. Huang, "A simple, real-time range camera," *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition*, pp. 256–261 (1989).
9. M. Subbarao, "Parallel depth recovery by changing camera parameters," *Proc. of the IEEE Conf. on Computer Vision*, pp. 149–155 (1988).
10. M. Subbarao and G. Surya, "Depth from defocus: A spatial domain approach," *Int. J. Comput. Vis.* **13**, 271–294 (1994).
11. Y. Y. Schechner and N. Kiryati, "Depth from defocus vs. stereo: How different really are they?," *Int. J. Comput. Vis.* **39**, 141–162 (2000).
12. M. Watanabe and S. K. Nayar, "Rational filters for passive depth from defocus," Technical Report CUCS-035-95, Dept. of Computer Science, Columbia University, New York (1995).
13. D. Ziou, "Passive depth from defocus using a spatial domain approach," *Proc. of the Intl. Conf. of Computer Vision (ICCV)*, pp. 799–804 (1998).
14. A. N. Rajagopalan and S. Chaudhury, "Optimal recovery of depth from defocused images using an MRF model," *Proc of the Intl. Conf. on Computer Vision (ICCV), Bombay, India*, pp. 1047–1052 (1998).
15. F. Deschenes, D. Ziou, and P. Fuchs, "Improved Estimation of defocus blur and spatial shifts in spatial domain: A homotopy-based approach," *Pattern Recogn.* **36**(9), 2105–2125 (2003).
16. S. Bhasin and S. Chaudhuri, "Depth from defocus in presence of partial self occlusion," *Proc. of the Intl. Conf. on Computer Vision (ICCV)*, Vol. 1, pp. 488–493, Vancouver, Canada (2001).
17. P. Favaro, A. Mennucci, and S. Soatto, "Observing shape from defocused images," *Int. J. Comput. Vis.* **52**(1), 25–43 (2003).
18. A. Pentland, S. Scherock, T. Darrell, and B. Girod, "Simple range cameras based on focal error," *J. Opt. Soc. Am. A* **11**, 2925–2935 (1994).

19. S. K. Nayar, M. Watanabe, and M. Noguchi, "Real-time focus range sensor," *Proc. of the Intl. Conf. on Computer Vision (ICCV)*, pp. 995–1001 (1995).
20. B. Girod and S. Scherock, "Depth from defocus and structured light, Optics, Illumination and Image Sensing for Machine Vision IV," *Proc. Soc. Photo-Opt. Instrum. Eng.* **1194**, 209–215 (1989).
21. A. M. Darwish, "3D from focus and light stripes," *Proc. of SPIE Sensors and Control for Advanced Automation II*, Vol. 2247, pp. 194–201, Frankfurt, Germany (1994).
22. M. Noguchi and S. K. Nayar, "Real-time focus range sensor," *Proc. of the Intl. Conf. on Pattern Recognition* (1994).
23. A. R. FitzGerrell, E. R. Dowski, and W. T. Cathey, "Defocus transfer function for circularly symmetric pupils," *Appl. Opt.* **36**(23), 5796–5804 (1997).
24. H. C. Lee, "Review of image-blur models in a photographic system using principles of optics," *Opt. Eng.* **29**(5), 405–421 (1990).
25. A. Mennucci and S. Soatto, "On observing shape from defocused images," *Proc. of the Intl. Conf. on Image Analysis and Processing*, pp. 550–555 (1999).
26. J. P. Jones and L. A. Palmer, "An evaluation of the two-dimensional Gabor model of simple receptive fields in cat striate cortex," *J. Neurophysiol.* **58**, 1233–1258 (1987).
27. N. Petkov and P. Kruizinga, "Computational models of visual neurons specialised in the detection of periodic and aperiodic oriented visual stimuli: bar and grating cells," *Biol. Cybern.* **76**, 83–96 (1987).
28. T. Darrell and K. Wohn, "Depth from defocus using a pyramid architecture," *Pattern Recogn. Lett.* **11**, 787–796 (1990).
29. M. Watanabe and S. K. Nayar, "Telecentric optics for computational vision," *Proc. of Image Understanding Workshop (IUW 96), Palm Springs* (1996).
30. O. Ghita and P. F. Whelan, "A bin picking system based on depth from defocus," *Mach. Vision Appl.* **13**(4), 234–244 (2003).
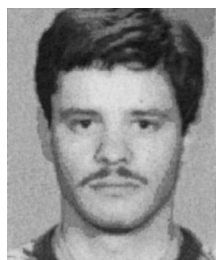
Ovidiu Ghita received his BE and ME degrees in electrical engineering from Transilvania University, Brasov, Romania. From 1994–1996 he was an assistant lecturer in the Department of Electrical Engineering at Transilvania University. Since then, he has been a member of the Vision Systems Group at Dublin City University (DCU) during which time he received his PhD for work in the area of robotic vision. Currently, he holds a position of postdoctoral research assistant in the Vision Systems Laboratory at DCU. His current research interests are in the area of range acquisition, shape representation, machine vision, and medical imaging.

Paul F. Whelan received his BEng(Hons) degree from the National Institute for Higher Education Dublin, his MEng degree from the University of Limerick, and his PhD from the University of Wales, Cardiff. During the period 1985–1990 he was employed by Industrial and Scientific Imaging Ltd. and later Westinghouse (WESL), where he was involved in the research and development of industrial vision systems. He was appointed to the School of Electronic Engineering, Dublin City University (DCU) in 1990 and currently holds the position of associate professor and director of the vision systems laboratory. As well as a wide range of scientific publications, Professor Whelan coedited *Selected Papers on Industrial Machine Vision Systems* (1994), and was the coauthor of *Intelligent Vision Systems for Industry* (1997) and *Machine Vision Algorithms in Java* (2000). His research interests include applied morphology, texture analysis, machine vision, and medical imaging. He is a Senior Member of the IEEE, a Chartered Engineer, and a member of the IEE, SPIE, and IAPR. He is also a member of a number of machine vision related conference program committees. He currently serves on the IEE Irish center committee, as member of the governing board of the *International Association for Pattern Recognition* (IAPR) and as the president of the *Irish Pattern Recognition and Classification Society*.

John Mallon received his BEng (H1) degree from Dublin City University, Ireland. Currently he is a member of the Vision Systems Group and is working towards a PhD degree in the area of autonomous navigation. His research interests include range data acquisition and multisensor fusion.