Telecentric Optics for Constant-Magnification Imaging†

Masahiro Watanabe and Shree K. Nayar

CUCS-026-95

Department of Computer Science
Columbia University
New York, N.Y. 10027

September, 1995

†This research was conducted at the Center for Research in Intelligent Systems, Department of Computer Science, Columbia University. It was supported in part by the Production Engineering Research Laboratory, Hitachi, and in part by the David and Lucile Packard Fellwoship. Masahiro Watanabe is with the Production Engineering Research Laboratory, Hitachi.

Abstract

A novel approach to constant-magnification imaging is proposed. Magnification variations due to changes in focus setting pose problems for important vision techniques, such as, depth from defocus. It is shown that magnification of a conventional lens can be made invariant to defocus by simply adding an aperture at an analytically derived location. The resulting optical configuration is called "telecentric." It is shown that most commercially available lenses can be turned into telecentric ones. The procedure for calculating the position of the additional aperture and a detailed analysis of the photometric and geometric properties of telecentric lenses are discussed. Finally, the application of telecentric optics to the problem of depth from defocus is demonstrated. The proposed optics is shown to result in significantly improved depth maps than those obtained using a conventional lens.

1 Introduction

The problem of magnification variation due to change in focus setting has been one of the interests to machine vision researchers. The classical approach to solving this problem has been to view it as one of *camera calibration* [5, 17]. Willson and Shafer[17] conducted a careful analysis of the interaction between focus and magnification. They proposed a joint calibration approach that measures the relation between zooming and focusing for a given lens. Through this calibration model, it becomes possible to hold magnification constant while focusing; the calibration results are used to apply zoom adjustments so as to correct magnification changes caused by focusing. Though this approach provides a general scheme to tackle the problem, it has its drawbacks. One requires an expensive computer-controlled zoom lens and, of course, the extensive calibration procedure mentioned above, even if one only needs to vary the focus setting and not any other parameter. Further, the necessity to change two physical parameters (focus and zoom) simultaneously tends to increase errors caused by backlashes in the lens mechanism and variations in lens distortion.

An alternative approach to the magnification problem is a computational one, commonly referred to as image warping. Darrell and Wohn [3] proposed the use of warping to correct image shifts due to magnification changes caused by focusing. This method is simple and effective for some applications, but can prove computationally intensive for real-time ones. Furthermore, since warping is based on spatial interpolation and resampling techniques, it generally introduces undesirable effects such as smoothing and aliasing. These can be harmful for applications that rely on precise spatial-frequency analysis, such as depth from focus/defocus.

Depth from focus/defocus methods provide a powerful means of getting a range map of a scene from two or more images taken from the same viewpoint but with different optical settings. Depth from focus uses a sequence of images taken by incrementing the focus setting in small steps. For each pixel, the focus setting that maximizes image contrast is determined. This in turn can be used to compute the depth of the corresponding scene point [7, 10, 11, 3, 17, 9]. Magnification variations due to defocus, however, cause additional image variations in the form of translations and scalings. Estimation of image contrast in the presence of these effects will clearly result in erroneous depth estimates.

In contrast to depth from focus, depth from defocus uses only two images with different optical settings [14, 4, 2, 15, 18, 13]. As it attempts to compute depth from this minimal number of images, it requires all aspects of the image formation process to be precisely modeled. Since the two images taken correspond to different levels of defocus, they are expected to differ in magnification. Given that two images are all we have in the case of defocus analysis, the magnification variations prove particularly harmful during depth computation. As a result, most investigators of depth from defocus have been forced to vary aperture size [14, 4, 2, 15, 18] rather than the focus setting (for example, the distance between the sensor and the lens) to induce defocus variations between the two images. Aperture size changes have the advantage of not introducing magnification variations. However, this approach has two major drawbacks when compared with the idea of changing the focus setting: (a) One of the two apertures used must be significantly smaller than the other. This decreases the brightness of the corresponding image and consequently degrades its signal-to-noise ratio. (b) The sensitivity of aperture induced defocus to depth is lower and less uniform than in the case of focus setting changes. Therefore, depth from defocus based on focus change is clearly desirable if there exists a simple and inexpensive way of making magnification invariant to defocusing.

In this paper, a novel approach to constant-magnification imaging is described. Rather than resorting to the traditional approaches of detailed calibration or computational warping, we suggest a simple but effective optical solution to the problem. The magnification problem is eliminated in its entirety by the introduction of an optical configuration, referred to as *telecentric optics*. Though telecentricity has been known for long in optics [1, 8], it has not been exploited in the realm of computational vision. With this optics, magnification remains constant despite focus changes. The attractive feature of this solution is that commercially available lens (used extensively in machine vision) are easily transformed to telecentric ones by adding an extra aperture. The positions of aperture placement for a variety of off-the-shelf lenses are analytically derived. Further, extensive experimentation is conducted to verify the invariance of magnification to defocus in four telecentric lenses that were derived from commonly used commercial lenses.

We have successfully incorporated a telecentric lens into a real-time active range sensor that is based on depth from defocus [13]. Here, we demonstrate applications of telecentric optics to passive depth from defocus. In this application, the performance of the telecentric lens is shown to be far superior to that of a conventional one.

2 Telecentricity

2.1 Conventional Lens Model

To begin with, we discuss the lens model that is widely used in computer vision. Figure 1 shows the commonly used image formation model, where the main assumptions are that the lens is thin and the aperture position coincides with the lens. All light rays that are radiated by scene point P and pass the aperture A are refracted by the lens to converge at point Q on the image plane. The relationship between the object distance d, focal length of the lens f, and the image distance d_i is given by the Gaussian lens law:

$$\frac{1}{d} + \frac{1}{d_i} = \frac{1}{f}. \tag{1}$$

Each point P on the object plane is projected onto a single point Q on the image plane, causing a clear or *focused* image I_f to be formed. When the focus setting is changed by displacing the sensor plane, for example, to I_1 or I_2 from I_f, the energy received from P by the lens is distributed over a circular patch on the sensor plane[1]. Although this causes a blurred image, the effective image location of point P can be taken to be the center of the circle. This center lies on the ray R which is radiated from point P and passes through the center of the aperture O and refracted by the lens. This ray is called the *principal ray* [1, 8]

[1] Here, focus change is modeled as a translation of the sensor plane. This model is also valid for the case of lens translation which is used in most non-zoom lens systems, where the distance d between lens and object is typically much larger than the focal length f.

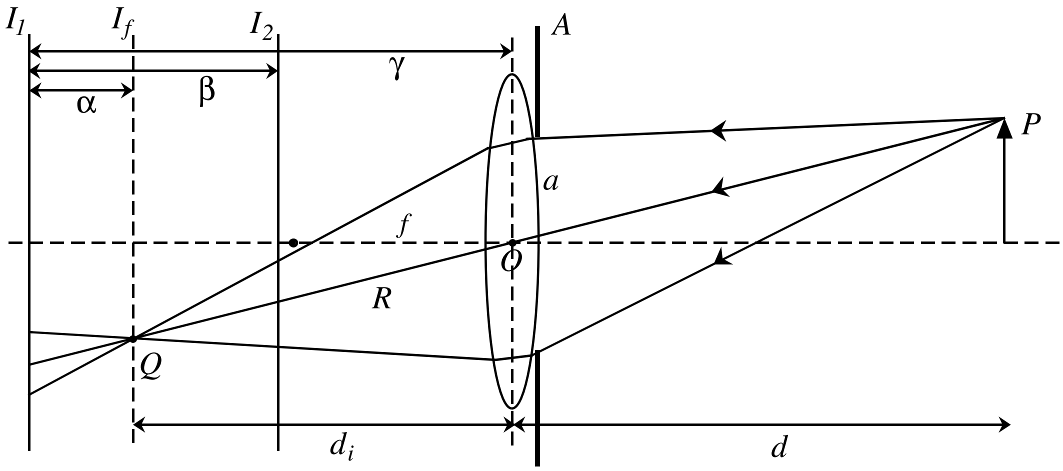

Figure 1: Image formation using a conventional thin lens model.

in optics. Since its intersection with the sensor plane varies with the position of the sensor plane, image magnification varies with defocus.

2.2 Telecentric Optics

Keeping in mind the image formation model shown in Figure 1, we proceed to discuss the constant magnification configuration, called *telecentric optics*. Figure 2 illustrates the principle underlying telecentric projection. The only modification made with respect to the conventional lens model of Figure 1 is the use of the external aperture A'. The aperture is placed at the *front-focal plane*, i.e. a focal length in front of the *principal point* O of the lens. This simple addition solves the problem of magnification variation with distance α of the sensor plane from the lens. Straightforward geometrical analysis reveals that the ray of light R' from any scene point that passes through the center O' of aperture A', i.e. the *principal ray*, emerges parallel to the optical axis on the image side of the lens [8]. Furthermore, this parallel ray is the axis of a cone that includes all light rays radiated by the scene point, passed through by A', and intercepted by the lens. As a result, despite defocus blurring, the effective image coordinates of point P on the sensor plane stays constant irrespective of the displacement of the sensor plane from I_f. In the case of depth from defocus, the magnification of any point in the scene, regardless of its position in the scene, remains the same in both images, I_1 and I_2. It is also easy to see from Figure 2 that this constant-magnification property is unaffected by the aperture radius a' as far as it is not large enough to cause severe vignetting.

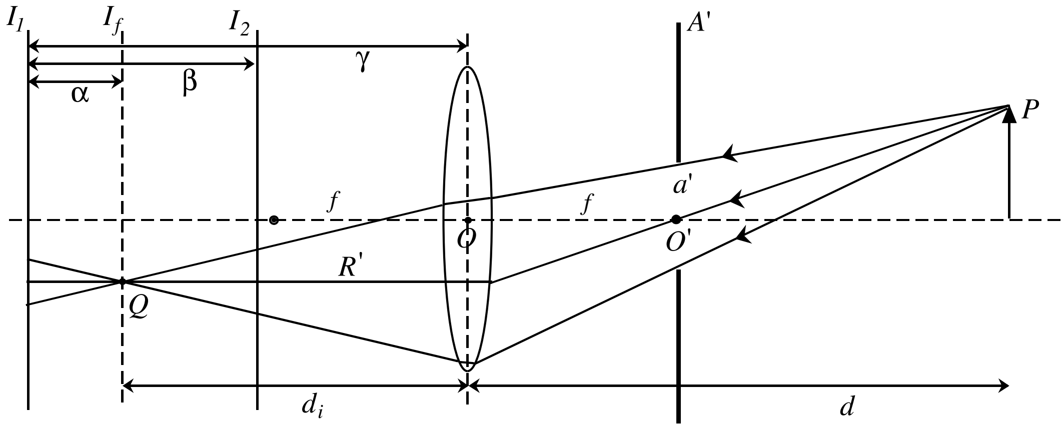

Figure 2: Telecentric optics achieved by adding an aperture to a conventional lens. This simple modification causes image magnification to be invariant to the position of the sensor plane, i.e. the focus setting.

2.3 Properties of Telecentric Optics

While the nominal and effective *F-numbers* for the conventional lens model in Figure 1 are $f/2a$ and $d_i/2a$, respectively, they are both equal to $f/2a'$ in the telecentric case. The reason for this invariance of the F-number for telecentric optics is the following. The effective F-number is defined as the ratio of the height to the diameter of the base of the light cone that emerges out of the lens. The nominal F-number is the effective F-number when the lens is focused at infinity. The above-mentioned light cone is bordered by the family of *marginal rays* which pass through the aperture A' touching its circumference. Consider these rays as they emerge from the lens on the image side. If the scene point P is displaced, the marginal rays on the image side only shift in a parallel fashion, keeping the angle subtended by the apex of the cone constant. This is due to the aperture being located at the front focal plane (see Figure 2): Light rays which pass through the same point in the front focal plane emerge parallel from the lens. As a result, the effective F-number does not change when the distance of the focused point from the lens changes. Further, when the lens is focused at infinity, the effective F-number is $f/2a'$ because the marginal rays pass through the aperture A' parallel to one another and strike the lens with the same diameter $2a'$ as the aperture, and converge on the image side of the lens at the focal length f.

This above fact results in another remarkable property of the telecentric lens: The image brightness stays constant in a telecentric lens, while, in a conventional lens, brightness decreases as the effective focal length d_i increases. This becomes clear by examining the image irradiance equation [6]:

$$E = L \frac{\pi}{4} \frac{\cos^4 \theta}{{F_e}^2}, \tag{2}$$

where, E is the irradiance of the sensor plane, L is the radiance of the surface in the direction of the lens and F_e is the effective F-number. In the case of a conventional lens, $F_e = d_i/2a$, while $F_e = f/2a'$ for a telecntric lens. θ is the angle between the optical axis and the principal ray which originates from object point P and passes through the center of the aperture.

In summary, telecentric optics provides us a way of taking multiple images of a scene at different focus settings while keeping magnification constant between the images. In practice, this can be accomplished by using a beam splitter (or for more than two images, a sequence of beam splitters) behind the telecentric lens [12]. In a real-time application, all images can be digitized simultaneously and processed in parallel as in [13].

3 Aperture Placement

Although the discussion on telecentricity in section 2 was based on the thin lens model, the results holds true for compound lenses. The thin lens is just a special case when the two *principal planes* [1, 6] coincides. Figure 1 is easily modified for the compound lens case, by replacing the thin lens with the two principal planes, U and U', of the compound lens, as shown in Figure 3 (a). The Gaussian lens law (1) remains valid when the distances are measured as follows; the scene point distance d from plane U and the image distance d_i from plane U'.

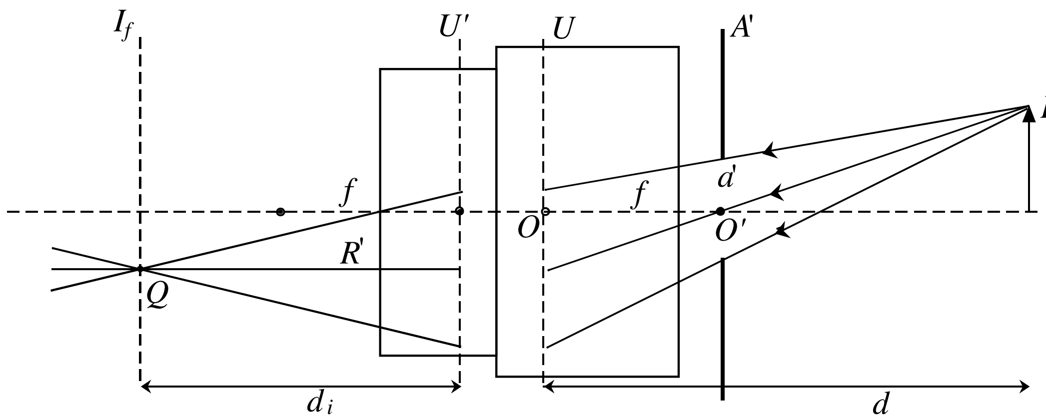

(a) Aperture placement outside the lens.

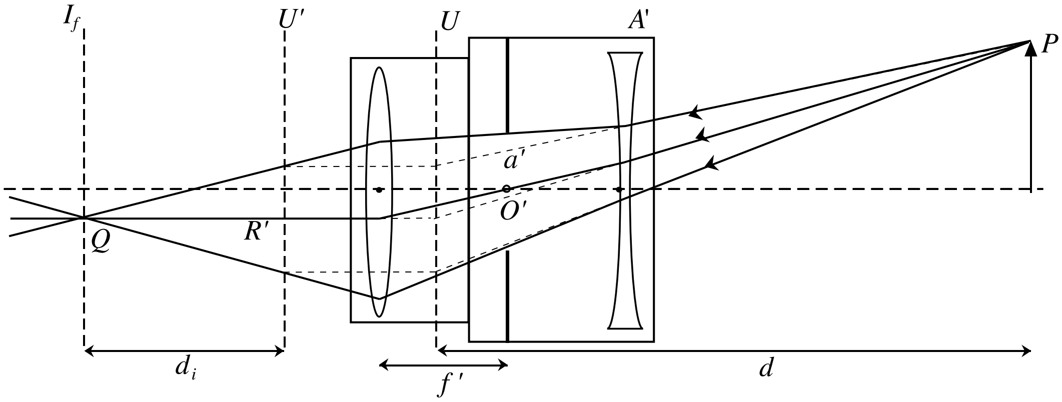

(b) Aperture placement inside the lens.

Figure 3: Telecentric aperture placement in the case of compound lenses.

Given an off-the-shelf lens, the additional telecentric aperture can be easily appended to the casing of the lens as far as the front focal plane is outside the lens. This is the case with most lenses including telephoto ones. However, for wide angle lenses with focal lengths shorter than the *back focal length* (distance from lens mount plane to sensor plane), the front focal plane is likely to be inside the lens. In such cases, one can still make the lens telecentric by placing an aperture inside the lens. The procedure is as follows. First, consider a set of parallel rays entering from the image side of the lens and find the point in the lens system where the parallel rays converge. In Figure 3(b), this would correspond to the point O'. If this point does not lie inside any of the physical lenses that comprise the compound lens, one can open the lens and place the telecentric aperture at the plane passing through the point and normal to the optical axis of the lens. Fujinon's CF12.5A F/1.4 is an example such a lens, which we converted to telecentric by placing an aperture inside and used to develop a real-time focus range sensor [13].

In practice, the exact location where the additional aperture needs to be placed can be determined in the following way. In some cases, the lens manufacturer provides information regarding the front focal position, which is customarily denoted as F in the schematic diagram. For instance, Nikon provides this data for their old line of SLR lenses, two of which, Micro-Nikkor 55mm F/2.8 and Nikkor 85mm F/2, are used in our experiments reported in section 4.1. If this information is not available from the manufacturer, it can be determined by the following procedure. Hold the lens between a screen (say, a white piece of paper) and a far and bright source (such as the sun or a distant lamp). In this setup the lens is held in the direction opposite to normal use; light enters from the back of the lens and the image is formed on the screen in front of the lens. The screen is shifted around to find the position that provides the clearest image. This position is the front focal plane where the telecentric aperture should be placed.

The lens is then mounted on an image sensor with the telecentric aperture attached to it, and the actual magnification variation due to defocus is used as feedback to refine the aperture position so as to drive magnification variation to zero. The method we used to measure magnification change is detailed in section 4.1. The above refinement process is recommended even when the front focal plane position is available from the manufacturer. This is because the precise position of the front focal plane may vary slightly between lenses of the same model. As we will see, after conversion to telecentric, magnification variations produced by a lens can be reduced to well below 0.1%.

While using the telecentric lens, the original aperture of the lens should be opened fully and the diameter of the telecentric aperture should be chosen so as to minimize vignetting. The degree of vignetting can be gauged by reducing the original aperture by one step from the fully open position and measuring image brightness in the corners of the image. If vignetting is significant, even this small reduction in the original aperture size will change image brightness. Table 1 summarizes all the information needed to convert five popular lenses into telecentric ones. We have in fact converted all five lenses and the telecentric versions of three of them are shown in Figure 4. In each case except Fujinon's CF12.5A, the telecentric aperture resides in the aluminum casing appended to the front of lens. In the case of Fujinon's CF12.5A, the telecentric aperture resides inside the lens body.

Table 1: Aperture placement for five off-the-shelf lenses. Outside aperture position is measured along the optical axis from the surface of outermost lens. Inside aperture position is measured from the stray light stop aperture toward the object direction.

Lens	Focal lenth	F-number	Aperture position	Max. aperture†
Fujinon CF12.5A	12.5mm	1.4	Inside : 4mm	F/8
Cosmicar B1214D-2	25mm	1.4	Outside: 4.7mm	F/8
Nikon AF Nikkor	35mm	2	Outside: 3.3mm	F/8.5
Nikon Micro-Nikkor	55mm	2.8	Outside: 46mm (38.5mm)*	F/13
Nikon Nikkor	85mm	2	Outside: 67mm (57.8mm)*	F/6.8

* Number in () is the maker supplied front focal position.

† Max. aperture is maximam aperture which does not cause vignetting (converted into F-number).

There is yet another way to convert a conventional lens into a telecentric one. This is by placing a convex lens between the lens and the image sensor to make the principal ray parallel to the optical axis. Though this is better in the sense that one can use the full aperture range without worrying about vignetting, it changes the focal length and the position of the image plane. In addition, the convex lens must reside deep inside the the camera.

It is worth mentioning that there are some commercially available lenses that are called telecentric. But these lenses are telecentric on the object side, i.e. principal rays come into the lens parallel to the optical axis, which is opposite to what has been discussed here. These lens are used in profile projectors where magnification changes caused by the variation of object

(a) The Nikon Nikkor f=85mm SLR lens (left)

and the Nikon AF Nikkor f=35mm SLR lens (right).

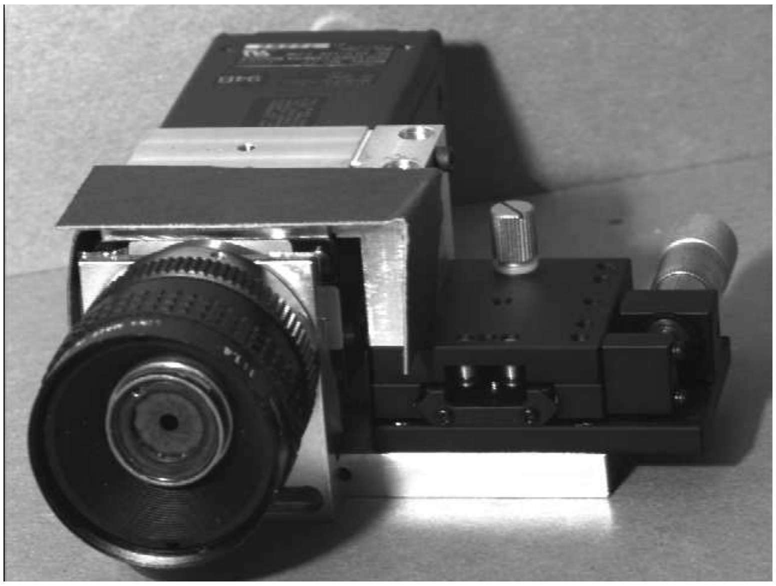

(b) Fujinon CF12.5A f=25mm lens on a micrometer stage.

Figure 4: Popular lenses converted into telecentric ones. In each case, the telecentric aperture resides in the aluminum casing appended to the front of the lens. The micrometer stage in (b) allows constant-magnification focus variation by displacing the image sensor from the lens.

distance from the lens is a serious problem. In effect, these lenses realize precise orthographic projection and not constant-magnification focusing.

It turns out that zoom lenses for 3-CCD color cameras are made telecentric on the image side to avoid color shading caused by RGB color separation. One can tell this by looking at the exit pupil position in the specification sheet provided by the manufacturer. If the exit pupil position is at ∞, the lens is image-side telecentric [1, 8]. An example is Fujinon's H12$\times$10.5A. But this does not mean that this zoom lens has magnification that is invariant to focus change, because zoom lenses usually change focus by complex movements of some of the lens components. To achieve constant magnification in such cases, a special mechanism must be added to shift the relative position between the zoom lens and CCD sensor.

4 Experiment

4.1 Magnification

To verify the constant-magnification capability of telecentric lenses, we have taken a series of images by changing the focus level and measured the magnification. To detect magnification change between images, the following method was used.

1. **FFT-phase-based local shift detection:** We compute the Fourier Transform of corresponding small areas in the two images and find the ratio of the two resulting spectra. Then a plane is fitted to the phases of the divided spectra. The gradient of the estimated plane is nothing but the image shift vector. As the two image areas should contain the same scene areas to get sub-pixel accuracy, the area used for FFT computation is refined iteratively by using the computed image shift. The image area used to compute local shifts is 64$\times$64 pixels.

2. **Object pattern:** A crisp and dense pattern is designed to ensure the phase estimates are accurate. The pattern we have used is shown in Figure 5 (b). The period of the pattern must be larger than the FFT computation area to avoid phase ambiguities.

3. **Needle diagram of local shifts:** A sparse needle map (5×4 needles) of the local shift vectors are chosen for display over the 640×480 pixel image.

4. **Magnification change and translation:** A shift vector $(\Delta x, \Delta y)$ at image coordinate (x, y), where x and y are measured from the center point of the image, is modeled as a similarity transformation:

$$\Delta x \;=\; a_x + m\,x\,, \quad \Delta y \;=\; a_y + m\,y\,. \tag{3}$$

   Least-mean-square fitting is used to estimate the parameters m and (a_x, a_y). This way we separate out the global magnification change m from the global translation (a_x, a_y). The translation factor (a_x, a_y) is introduced in the above model for two reasons; (a) the optical center of the image is assumed to be unknown, and (b) the optical center itself can shift between images due to the shift that result from any possible misalignment or wobble in the focus adjustment mechanism. The residual error of the above fit reveals the local shift detection accuracy and the validity of the above transformation model.

Figures 5 and 6 show the magnification change for the f=25mm lens without and with the telecentric aperture, respectively. In each of these two figures, (a) is the image focused at infinity, (b) at 787mm from the lens, and (c) at 330mm from the lens. Figures (d) and (e) are needle diagrams that show the local shifts of image (b) relative to image (a) and that of image (c) relative to image (a), respectively. The needles are magnified by a factor of 5 for clarity of display.

In the same manner, images were taken and processed for the 35mm, 55mm and 85mm lenses and the results are summarized in Table 2. As is clear from the table, magnification variations in the telecentric versions of the lenses are remarkably small. Pixel shifts due to large focus setting changes are as small as 0.5 pixels. With careful tuning of the aperture position, this number can be further reduced to 0.1 pixels. Small translations are observed in the cases of the 35mm and 55mm lenses. This is because the focus is changed using the focus ring of the lens. For the smaller lens, the translation caused by the wobbling of the lens is larger, as expected. But such translations can be easily compensated for by shifting the

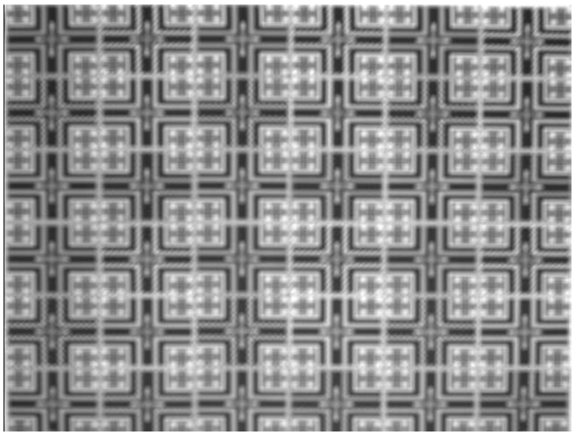

(a) d =infinity

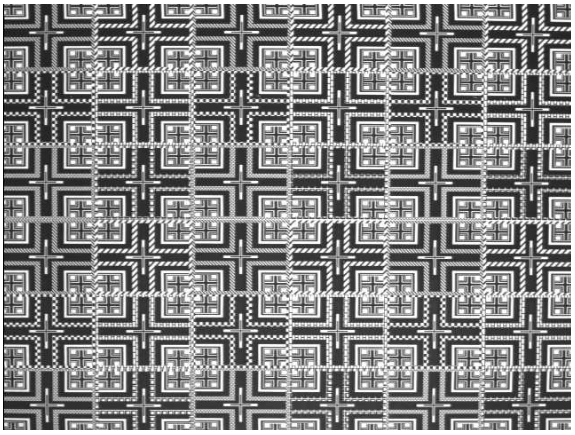

(b) d =787mm

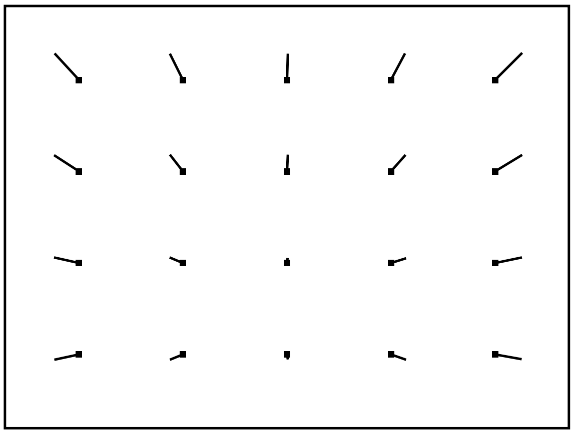

(d) shift:(b)-(a) $\times 5$

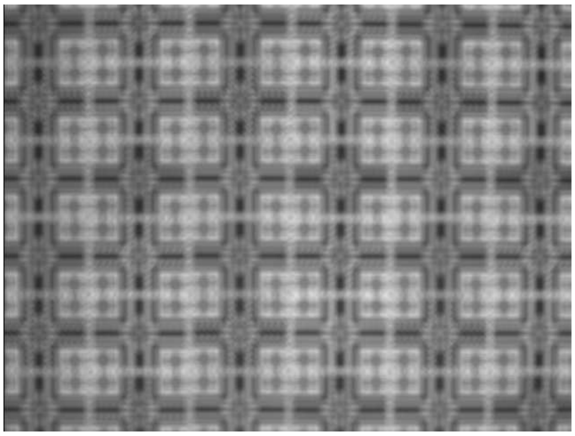

(c) d =330mm

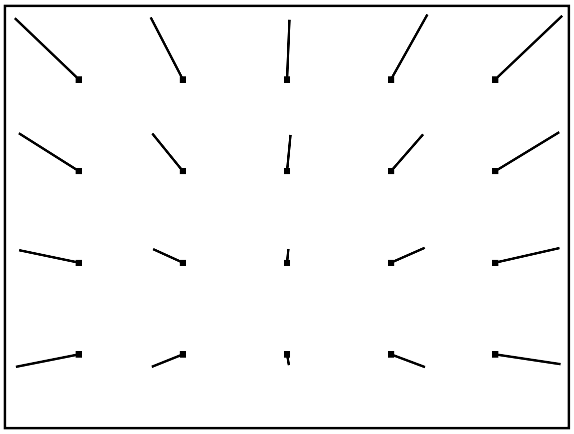

(e) shift:(c)-(a) $\times 5$

Figure 5: Magnification variations in non-telecentric optics due to focus change.

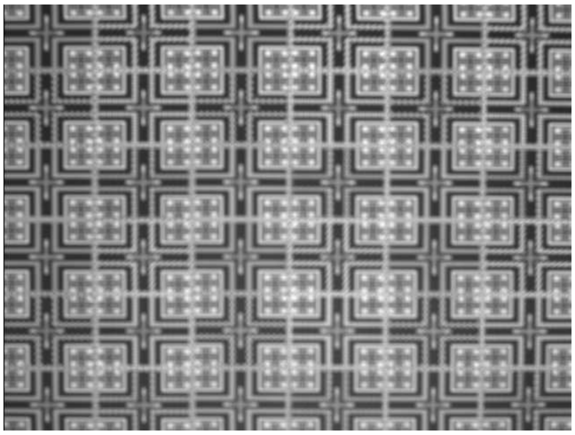

(a) d =infinity

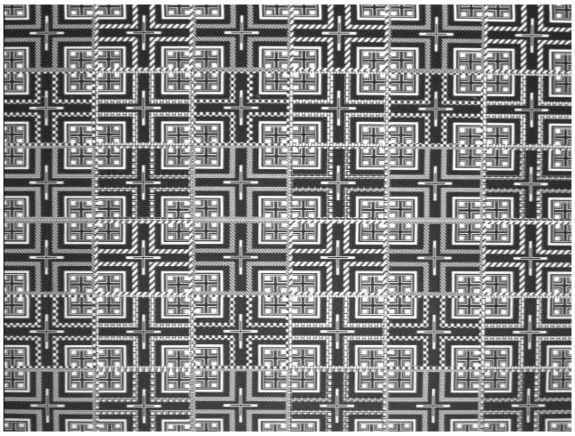

(b) d =787mm

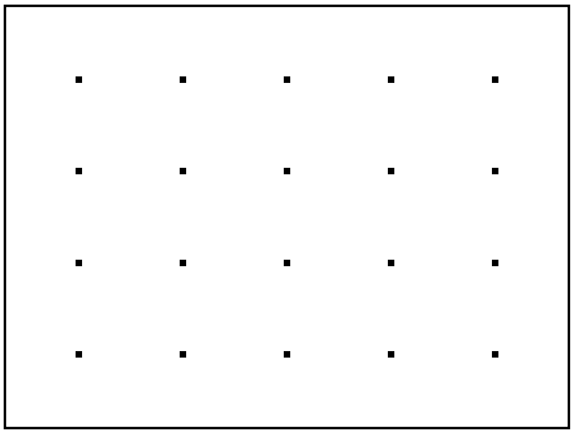

(d) shift:(b)-(a) $\times 5$

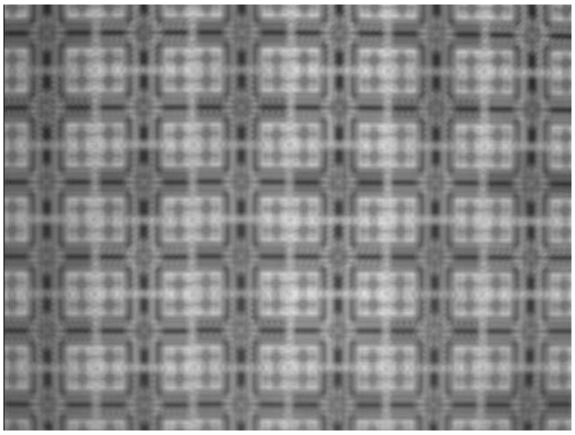

(c) d =330mm

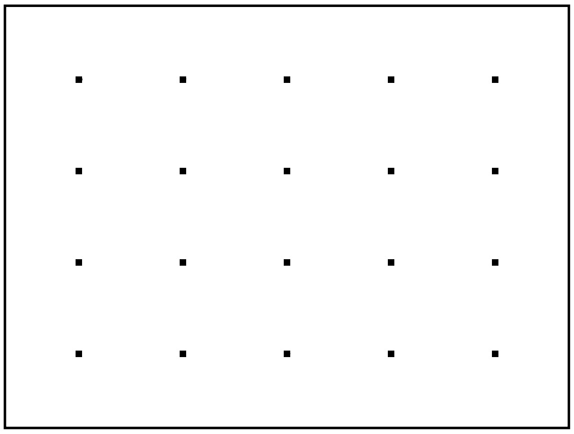

(e) shift:(c)-(a) $\times 5$

Figure 6: Magnification variations in telecentric optics due to focus change. The needle maps can be compared with those in Figure 5 to see that magnification remains unaffected by focus change.

image without any harmful side effects. This is because, telecentricity holds the magnification constant and hence the shifts that need to be applied are pure translations. In the case of the 25mm lens, which is mounted on a precise micrometer stage, (see Figure 4 (b)), translation between images is only of the order of 0.5 pixels.

Table 2: Magnification variations for four widely used lenses and their telecentric versions.

Lens	Shift of CCD*	Change in focused distance†	Non-telecentric		Telecentric	
			Mag.(max. shift‡)	Translation	Mag.(max. shift‡)	Translation
Cosmicar 25mm	2.05mm	∞ – 330mm	5.9% (18.9pix)	(0.4, 5.8)	-0.03% (0.1pix)	(0.5, 0.2)
AF Nikkor 35mm	2.08mm	4022– 550mm	4.5% (14.3pix)	(1.4, 4.7)	-0.15% (0.5pix)	(-0.7, 3.3)
Micro-Nikkor 55mm	2.94mm	∞ –1085mm	5.8% (18.6pix)	(-1.9, 0.0)	-0.07% (0.2pix)	(-0.2, 0.1)
Nikkor 85mm	2.85mm	1514–1000mm	4.7% (14.9pix)	(-0.9,-1.8)	0.01% (0.03pix)	(0.4,-1.0)

* Measured in image side.
† Measured in object side.
‡ Max. shift is the maximum shift caused by the magnification change.

The above experiments relate to the geometric properties of both conventional and telecentric lenses. Figure 7 shows the brightness variation due to focus change for Cosmicar 25mm lens and Nikkor 85mm lens. As explained in section 2.3, image brightness varies with focus change in the case of the conventional lens while it remains constant in the telecentric case due to the invariance of the effective F-number (see equation 2). As seen from the plot, brightness variations in the conventional lens is significant. This requires brightness values for different focus settings to be normalized. This process is avoided in the telecentric case. Table 3 summarizes our experiments on the photometric properties of the two configurations.

4.2 Example Application: Depth from Defocus

In order to demonstrate the application of telecentric optics to computer vision, we conducted experiments on the depth from defocus method. A description of the exact algorithm used

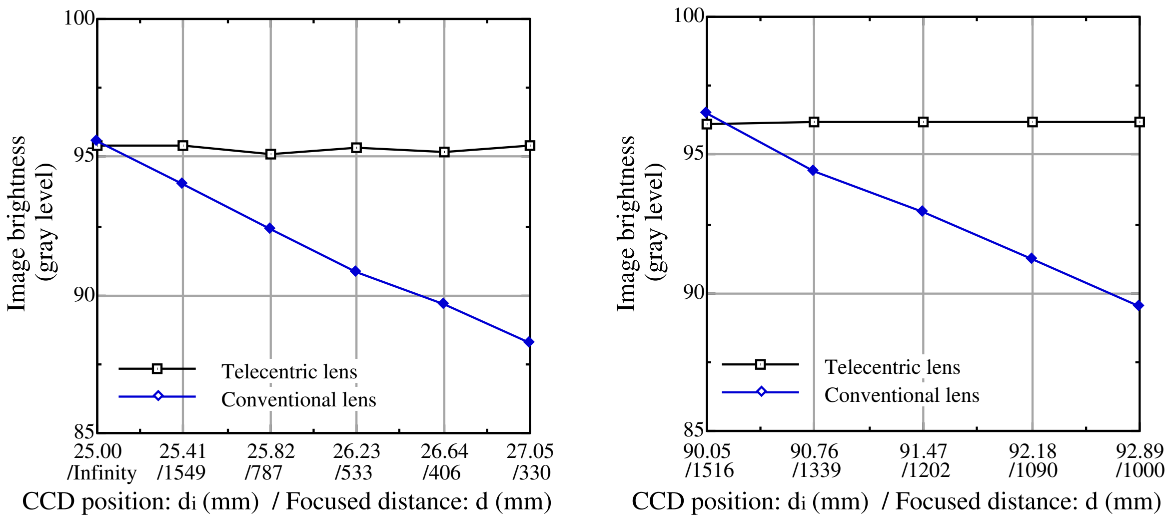

(a) Cosmicar 25mm

(b) Nikkor 85mm

Figure 7: Brightness variation due to focus change: brightness does not change with telecentric lenses.

for computing depth from relative defocus in two images is beyond the scope of this paper. Briefly, the algorithm we have used here applies a set of filters, that are sensitive to defocus but invariant to object texture, to two defocused images of a scene. (see [16] for details). The filter outputs are used to efficiently compute depth at each pixel, producing a full resolution depth map of the scene.

The 25mm lens was used in this experiment. The lens aperture was set to F/8.3. The first image was taken with the lens focused at 869mm from the camera, and the other image with the lens focused at at 529mm. These two distances were chosen so that all scene points lie between them. Figure 8 summarizes the results obtained using the lens in its conventional mode and the telecentric mode. Figures 8 (a) and (b) are the two defocused images taken without the telecentric aperture. The images obtained with the telecentric aperture appear the same to a human observer though they do not include the magnification change of 1.4% that exists between the images in (a) and (b). Figures 8 (c) and (d) are the gray-coded map and its wireframe plot, respectively, of the depth map computed from the images without the telecentric aperture. Large errors are found in this depth map that result primarily due to

Table 3: Brightness variation due to focus change for four widely used lenses and their telecentric versions.

Lens	Shift of CCD*	Change in focused distance†	Brightness variation	
			Non-telecentric	Telecentric
Cosmicar 25mm	2.05mm	∞ – 330mm	-7.6%	0.0%
AF Nikkor 35mm	2.08mm	4022– 550mm	-5.1%	0.7%
Micro-Nikkor 55mm	2.94mm	∞ –1085mm	-9.5%	0.6%
Nikkor 85mm	2.85mm	1514–1000mm	-7.2%	0.1%

* Measured in image side.

† Measured in object side.

the magnification change between the two images. Figures 8 (c) and (d) are the gray-coded depth map and its wireframe plot, computed after converting the lens into a telecentric one. This simple experiment shows how telecentric optics can be used to significantly enhance the performance of computer vision techniques.

5 Conclusion

In this paper, we have shown that constant-magnification imaging can be achieved using what is called *telecentric optics*. Most commercial lenses can be turned into telecentric ones by simply attaching an additional aperture. The procedure for aperture placement and the photometric and geometric properties of telecentric lenses were discussed in detail. We have demonstrated that magnification changes can be reduced to as low as 0.03%, i.e. maximum magnification induced shift of 0.1 pixel, by the proposed aperture placing method. The only drawback of this optics is that the available F-number that does not cause vignetting is larger (aperture is smaller) than in the lens prior to aperture placement. This needs to be compensated by using either a brighter illumination or a more sensitive image sensor.

(a) Far-focused image

(b) Near-focused image

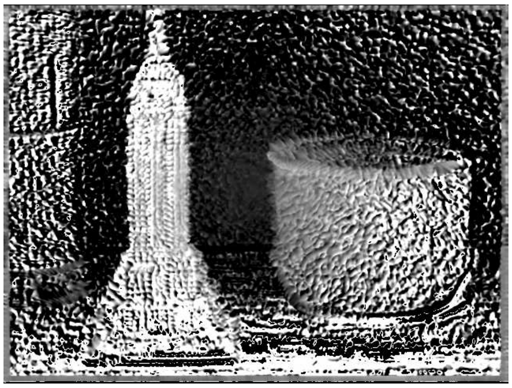

(c) Gray-coded depth map (non-telecentric lens)

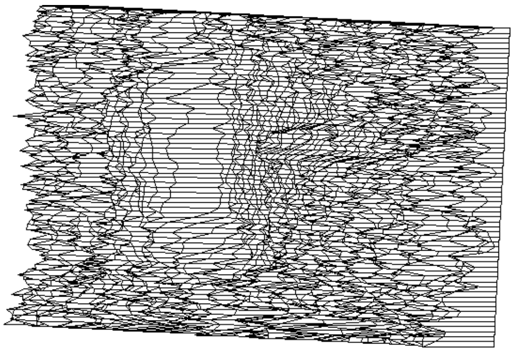

(e) Wire-frame (non-telecentric lens)

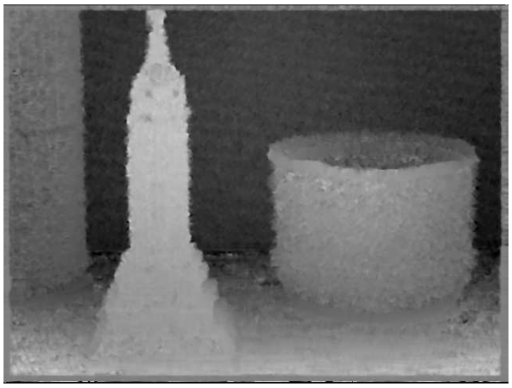

(d) Gray-coded depth map (telecentric lens)

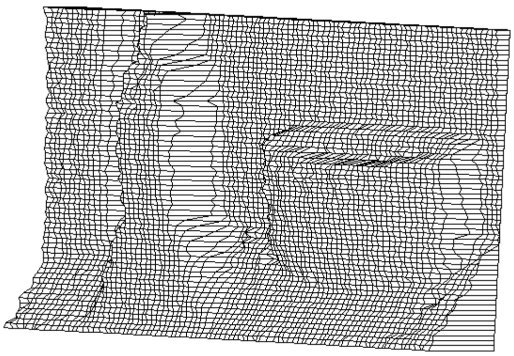

(f) Wire-frame (telecentric lens)

Figure 8: Depth from defocus without ((c) and (e)) and with ((d) and (f)) telecentric optics.

However, it should be noted that a lens designed from scratch can be made telecentric without substantial reduction in brightness. As an example of the application, we applied telecentric optics to the depth from defocus method. The computed depth maps were found to be far superior to ones computed using a conventional lens.

References

[1] M. Born and E. Wolf. *Principles of Optics.* London:Permagon, 1965.

[2] V. M. Bove, Jr. Entropy-based depth from focus. *Journal of Optical Society of America A*, 10:561–566, April 1993.

[3] T. Darrell and K. Wohn. Pyramid based depth from focus. *Proc. of IEEE Conf. on Computer Vision and Pattern Recognition*, pages 504–509, June 1988.

[4] J. Ens and P. Lawrence. A matrix based method for determining depth from focus. *Proc. of IEEE Conf. on Computer Vision and Pattern Recognition*, pages 600–609, June 1991.

[5] B. K. P. Horn. Focusing. Technical Report Memo 160, AI Lab., Massachusetts Institute of Technology, Cambridge, MA, USA, 1968.

[6] B. K. P. Horn. *Robot Vision.* The MIT Press, 1986.

[7] R. A. Jarvis. A perspective on range finding techniques for computer vision. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 5(2):122–139, March 1983.

[8] R. Kingslake. *Optical System Design.* Academic Press, 1983.

[9] A. Krishnan and N. Ahuja. Range estimation from focus using a non-frontal imaging camera. *Proc. of AAAI Conf.*, pages 830–835, July 1993.

[10] E. Krotkov. Focusing. *International Journal of Computer Vision*, 1:223–237, 1987.

[11] S. K. Nayar and Y. Nakagawa. Shape from focus: An effective approach for rough surfaces. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 16(8):824–831, August 1994.

[12] S. K. Nayar and M. Watanabe. Passive bifocal vision sensor. Technical Report Technical Report, (in preparation), Dept. of Computer Science, Columbia University, New York, NY, USA, Oct 1995.

[13] S. K. Nayar, M. Watanabe, and M. Noguchi. Real-time focus range sensor. *Proc. of Intl. Conf. on Computer Vision*, pages 995–1001, June 1995.

[14] A. Pentland. A new sense for depth of field. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 9(4):523–531, July 1987.

[15] G. Surya and M. Subbarao. Depth from defocus by changing camera aperture: A spatial domain approach. *Proc. of IEEE Conf. on Computer Vision and Pattern Recognition*, pages 61–67, June 1993.

[16] M. Watanabe and S. K. Nayar. Depth from defocus using a texture invariant operator set. Technical Report Technical Report, (in preparation), Dept. of Computer Science, Columbia University, New York, NY, USA, September 1995.

[17] R. G. Willson and S. A. Shafer. Modeling and calibration of automated zoom lenses. Technical Report CMU-RI-TR-94-03, The Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, USA, January 1994.

[18] Y. Xiong and S. A. Shafer. Moment and hypergeometric filters for high precision computation of focus, stereo and optical flow. Technical Report CMU-RI-TR-94-28, The Robotics Institute, Carnegie Mellon University, Pittsburg, PA, USA, September 1994.