A study of the sensitivity of long-range passive ranging techniques to atmospheric scintillation

Jason de Villiers[a,b], Fintan Wilson[a] and Fred Nicolls[b]

[a]Council for Scientific and Industrial Research, Pretoria, South Africa;
[b]University of Cape Town, Cape Town, South Africa

ABSTRACT

This work aims to quantify and improve the sensitivity of monocular passive ranging techniques in sequences that are degraded by the the quasi-periodic inter frame motion and localised blurring associated with atmospheric scintillation. Ranges from tens of meters up to 1km in uncontrolled outdoor conditions were investigated. This work has application in surveillance, border control, artillery, and the autonomous navigation of ground vehicles. A field trial with long focal length, large aperture cameras was held. The targets were placed at known distances and a scintillometer was used to measure the turbulence between the cameras and each target. A variety of depth from defocus and depth from focus algorithms were applied to a the video sequences to determine range to the targets. Increasing levels of atmospheric turbulence were observed ranging from no turbulence to mild, moderate and severe turbulence. Algorithms for the mitigation of atmospheric turbulence were applied to the recorded video sequences. The passive ranging techniques were then applied to the turbulence-mitigated video sequences and the results compared to that of the raw sequences. Finally, findings regarding the best performing combination of ranging and mitigation techniques to use under different scintillation conditions are presented.

Keywords: Depth from Defocus, Depth from Focus, Scintillation, Heat Shimmer, Passive Ranging

1. INTRODUCTION

Passive ranging has many applications for both military and civilian purposes. Some examples of these include ballistic measurements, enhanced surveillance, autonomous navigation, robotics, as well as measurement/bore-sighting for construction and manufacture. This work aims to investigate the usage of passive ranging techniques in uncontrolled outdoor environments for distances in the order of one kilometre.

1.1 Passive Ranging

Passive ranging techniques are those techniques that do not emit energy into the scene being observed to aid measurement, they thus exclude gated imaging, LIDAR, time of flight cameras and projecting fringe patterns. This work also further focusses on monocular ranging techniques and so eliminates stereo vision (or Ntple vision) with either a physical or temporal baseline.

This leaves the two related techniques of Depth from Defocus (DFD) and Depth from Focus (DFF). DFF is the process of determining at what which focal length each item in the scene being observed is maximally in focus. Once this focus is known the distance to the object can be calculated with geometric optics by using the equivalent thin lens parameters for the camera and the lens law:

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

where:

f = the focal length of the lens,

d = the distance to the object in world space, and

d_i = the distance of the imager from the equivalent thin lens optics.

Further author information: (Send correspondence to J.d.V.)
J.d.V.: E-mail: jdvilliers@csir.co.za

Visual Information Processing XXI, edited by Mark Allen Neifeld, Amit Ashok, Proc. of SPIE
Vol. 8399, 83990O · © 2012 SPIE · CCC code: 0277-786X/12/$18 · doi: 10.1117/12.920358

.

DFD, first described by Pentland[1] in 1987, attempts to quantify the amount of defocus at each point in the image and relate this to how far from the current focus distance the object is. The defocus is typically determined from only two images that were captured with varying camera parameters. This is in contrast with sweeping through some subset of the focal range for DFF.

1.2 Atmospheric Scintillation

Scintillation is the degradation experienced when viewing distant scenes horizontally through a large volume of air close to ground level. There are typically three types of degradation that are experienced:[2]

1. **Warping** is the random movement of pixels in a quasi-periodic fashion about their true position.
2. **Blurring** is the spatial and time variant local changing of focus across the image scene as observed by the camera.
3. **Washing** is the decrease in contrast, not only due to the blurring but also due to the varying attenuation of light along the optical path from the scene to the camera.

All of these effects are caused by air pockets of different temperature having slightly different refractive indices.

1.3 Contribution

This work investigates the feasibility of passive ranging at distances in the order of a kilometre as viewed horizontally through a turbulent atmosphere. The application of passive ranging techniques outdoors and to distances greater than several meters is a scarcely researched area. Conversely, the mitigation of atmospheric scintillation is an active area of research. This primary contribution of this work is the investigation into the feasibility of long range passive ranging and the application of scintillation mitigation techniques to determine passive ranging's sensitivity to scintillation.

2. DATA CAPTURE

2.1 Capture Equipment

Three sets of equipment were used for this work. The first was used to capture short range data with which to confirm the implementation of the ranging algorithms. This was done with an Allied Vision GT1920 1936×1456 machine vision camera with a Schneider Cinegon 4.8mm lens. This lens was replaced with a Fujinon HF75-SA 75mm lens for the midrange tests. Both of these lenses have external apertures of approximately 30mm diameter.

The long range image data for DFD and DFF ranging was captured with a custom 185mm focal length 4° field of view (FOV), 10cm aperture, F2.0 to F22.0 camera using an Allied Vision GX1910 1920×1080 sensor. A Scintec BLS900 Scintillometer was used to determine the refractive index structure function constant (C_n^2) which is the measurement of the atmospheric turbulence.

2.2 Short Range Data

The short range data consisted of a short range laboratory set up, using diffuse materials and controlled lighting. The purpose of this experiment was to verify the correct implementation of the DFF and DFD algorithms. Figures 1(a) and 1(b) show the near and far focussed images from the DFF scene. The data was captured by placing the GT1920 on a rigid tripod and manually changing the Schneider Cinegon's focus from in front of the closest object (Allied Vision GX1920 Camera) to behind the furthest object: the textured mouse pad. This was performed with the aperture fully open at F1.8.

The DFD data used the same apparatus, however the camera was focussed on the closest object in the scene (The Allied Vision GX Camera) with the aperture fully open and an image was captured. The aperture was then stopped down to F22, and the exposure period increased (to obtain a consistent brightness of objects between the two aperture settings) and the second image was captured.

2.3 Mid Range Data

This data set contained objects at distances from approximately 15m to 100m distant. The DFF data was captured with a 75mm lens. The lens was manually adjusted from focus in front of the closest object, the palisade fence (see Figure 3(a) to beyond the furthest objects (i.e. the supports for the shade netting storage area in the upper left of the images). Figure 3(b) contains a far focussed image. A video recording while the focus was slowly adjusted was captured and used.

The DFD data was captured with the custom 185mm focal length camera. Figures 4(a) and 4(b) contain the large and small aperture exposures respectively. The exposure period was again increased to attempt to equalise the brightness between the two exposures.

2.4 Long Range Data

The long range data was captured over two days on an grass runway which offered 1.4km line of sight and a slight upwardly concave elevation profile varying by no more than 2 metres over the viewable distance. The runway contained concrete slabs every 100 metres and the equipment was set up 50 metres from the first of these. A half-ton pick-up truck was used as the target, on it was placed mannequin torsos in camouflage uniform and large optical resolution charts containing horizontal bars at varying spatial frequencies. This vehicle was then driven to the marker 550 metres from the observing equipment. This is because 500m is the minimum working distance for the scintillometer. At every 100 metre marker from 550 metes to 1150 metres the vehicle was stopped and a 30 second recording of it at aperture stops F2.3, F5.0 and F11.0 was made. Thereafter a focal sweep was made from approximately 400 metres focus distance to infinity and a single image captured at each focus step. Concurrently with the image and video capturing, scintillometer readings were recorded for the current distance by using the scintillometer emitter which was placed adjacent to the vehicle.

3. DEPTH FROM FOCUS

The distinction between DFD and DFF, in this work, is the amount of input data available to the algorithm. This work uses a large number of frames captured with a constant maximally open aperture and a changing focus for DFF. The challenge thus becomes to determine at which frame, and therefore at which focal length, each pixel in the image is maximally in focus. The subsequent sections present the three metrics used to determine when an image patch is maximally in focus.

3.1 Michelson Contrast

Also known as the 'visibility' the Michelson contrast[3] was first defined by Nobel Laureate A.A. Michelson in 1927. It is the ratio of the difference in intensities to twice the average intensity and is expressed mathematically in Eq. (2). When an image point is out of focus, its blur radius bleeds onto adjacent pixels and it in turn is bled upon in the likely the case that its neighbouring pixels are also not perfectly focussed. This tends to average the pixel intensities and thus decrease the numerator of Eq. (2).

$$C(h,v,n) = \frac{I_{max} - I_{min}}{I_{max} + I_{min}} \tag{2}$$

where:

C = the Michelson contrast,
$I_{max} = max(I(i,j), \forall i \in [h-n, h+n], j \in [v-n, v+n])$,
$I_{min} = min(I(i,j), \forall i \in [h-n, h+n], j \in [v-n, v+n])$,
$I(i,j)$ = the intensity at coordinate (i,j),
(h,v) = the current pixel being considered, and
n = the window size.

3.2 Standard Deviation

Another metric used to determine the focus of an image is the standard deviation of an image patch. An in focus image patch will contain a greater number of intensities and possibly even sharp edge discontinuities, assuming these are present in the part of the scene currently being considered. This metric, like all focus and defocus metrics, relies on sufficient texture in the scene being viewed. Eq. (3) describes the metric used, it is a large sample number approximation to the true standard deviation.

$$S(h,v,n) \approx \sqrt{\frac{1}{(2n+1)^2}\sum_{i=v-n}^{v+n}\sum_{j=h-n}^{h+n}(I(i,j))^2 - \left(\frac{1}{(2n+1)^2}\sum_{i=v-n}^{v+n}\sum_{j=h-n}^{h+n}(I(i,j))\right)^2} \tag{3}$$

where:

S = the standard deviation,
$I(i,j)$ = the intensity at coordinate (i,j),
(h,v) = the current pixel being considered, and
n = the window size.

3.3 Shannon Entropy

Shannon described a measure to determine the information content, or entropy, of a signal.[4] If adjacent pixels are not independent (because they are out of focus and bleeding on to each other) then fewer bits are required to represent each of them on average as the greater correlation between pixels means that each new pixel contains less information. Eq. (4) provides Shannon's metric, it again was applied to a window region around each pixel. The Shannon entropy metric has also been shown to be a good indicator of image sharpness.[5]

$$Entropy = \int_0^{255} -p(i) \times log_2\left(p(i)\right)\, di \tag{4}$$

where:

$p(i)$ = the probability of a pixel having intensity i and is determined from the normalized 8-bit histogram.

4. DEPTH FROM DEFOCUS

DFD is the determination of distance based on only a few images and trying to quantify how out of focus an image point is and thus determine its distance from the focus plane. Techniques to quantify the amount of blur exist in both the frequency and spatial image domains and examples of each were evaluated in this work.

4.1 Pentland's Algorithm

Pentland proposed the original DFD algorithm[1] in 1987. It is an example of a spatial approximation to a frequency domain approach. In order to estimate the depth Pentland compares the 'normalised' high frequency components of the images, one of which is an all-focussed image created using a small aperture, to estimate the local blurring and then maps the blurring to the distance. This normalisation is two fold: firstly the output of the Laplacian filter (which is used to estimate the high frequency power) is averaged over a small window, and is then divided by the mean local brightness obtained with a Gaussian filter. The blur is estimated by dividing the blurry large-aperture image's normalised high frequency content by the that of the sharp small-aperture image. The mapping of the blur radiance to distance is given in Eq. (5).

$$D(h,v) = \frac{Fv_0}{v_0 - F - \sigma(h,v)f} \tag{5}$$

where:

D = the distance to the object at (h,v) in the image,
F = the focal length of the lens,
v_0 = the distance of the imager and the lens, and
f = the f number (ratio of focal length to aperture diameter) of the lens.

A direct frequency domain approach to Pentland's algorithm was also implemented. Using NVidia's Compute Unified Device Architecture[6] (CUDA) the Fast Fourier Transform (FFT) of a square window of length $2^n (n \in \mathbb{N})$ pixels was calculated for each pixel in the image and for both exposures. The output of each FFT was divided into a 4×4 grid and the sum of the magnitudes of the complex elements in the four blocks closest to 0Hz (i.e. the corners in our non-reordered 2D FFT) were taken to be the the low frequency power. The other 12 blocks' elements' magnitudes were tallied and the result taken as the power of the high frequency content. The high frequency power was divided by the low frequency power as originally suggested by Pentland to create a 'normalised high frequency content' with which to compare the two exposures. This implementation is referred to as the 'Pentland Frequency' algorithm in this work and the implementation of the 1987 paper[1] as the 'Pentland Spatial' algorithm.

4.2 Bove's Algorithm

Bove[7] considered the DFF problem in terms of the loss of entropy in a linear system, claiming that this provides superior results to regression (optimisation) based approaches. Bove provides a means to determine the entropy loss of the optical transfer function of the lens from the large and small aperture images, this is expressed in Eq. (6). Bove then provides a technique (Eq. (7)) to determine the entropy that a hypothesized blurring function would yield as a function of spatial frequencies and the distance of the object from the camera.

$$H(h,v) = -\frac{2}{W(h,v)} \sum_{i=0}^{N-1} (ln|S(w,h,v)| - ln|L(w,h,v)|) \tag{6}$$

$$H(h,v) = -\frac{1}{W(h,v)} \int_W ln|F(w,h,v)|^2 dw \tag{7}$$

where:

$H =$ the entropy,
$F(w,h,v) = \mathfrak{F}\big(f(z(h,v))\big)$
$S(w,h,v) = \mathfrak{F}\big(s(h,v)\big)$
$L(w,h,v) = \mathfrak{F}\big(l(h,v)\big)$
$s(h,v)$ = the small aperture image, and
$l(h,v)$ = the large aperture image.

Eq. (6) is calculated directly from the 2D FFT of a window centred around each pixel, and for a chosen blur kernel ($f(z(h,v))$ can then be equated to Eq. (7) to calculate the blur radius and thus the distance to the target. This is also a frequency domain approach to DFD and is comparable to the DFF entropy metric for determining best focus described in §3.3.

4.3 Watanabe and Nayar's Algorithm

Watanabe and Nayar[8] used a set of predetermined rational filters to try improve upon the inaccuracies of relative defocus measurement that resulted from using Laplacian and Gaussian filters. They claim that those kernels are overly sensitive to the frequency spectra of the scene being viewed. Simple operations with a few filter kernels are performed on linear combinations of the near and far focussed images in order to determine the depth in a manner invariant to the frequency of the scene. The four filter kernels are:

- g_F: a smoothing filter applied to subtracted and added images,
- g_{M1}: a filter applied to the smoothed subtraction image,
- g_{P1}: a filter applied to the smoothed added image, and
- g_{P2}: a filter applied independently to the smoothed added image to produce a second filtered added image.

Using these filters and the near (I_C) and far (I_F) focussed images, Eq. (8) shows how an initial depth estimate (B_0) is refined using a single Newton-Rhapson iteration to provide the final depth estimate, B. Watanabe and Nayar provided a means to determine the filter kernels for a chosen blur kernel, they also provided a set of optimised filter kernels for a Gaussian blur kernel. This work makes use of the 7×7 filter kernels Watanabe and Nayar provided. This method is therefore a pure spatial domain method.

$$B(h,v) = B_0(h,v) - \frac{\frac{C_{P2}(h,v)}{C_{M1}(h,v)} B_0^3(h,v)}{\frac{C_{P1}(h,v)}{C_{M1}(h,v)} - 3\frac{C_{P2}(h,v)}{C_{M1}(h,v)} B_0^2(h,v)} \tag{8}$$

where:

$$\begin{aligned}
B_0 &= \frac{(I_C(h,v) - I_F(h,v))C_{M1}(h,v)}{(I_C(h,v) + I_F(h,v))C_{P1}(h,v)}, \\
C_{M1} &= ((I_C - I_F) \otimes g_F) \otimes g_{M1}, \\
C_{P1} &= ((I_C + I_F) \otimes g_F) \otimes g_{P1}, \\
C_{P2} &= ((I_C + I_F) \otimes g_F) \otimes g_{P2}, \\
I_C &= \text{the close focussed image, and} \\
I_F &= \text{the far focussed image.}
\end{aligned}$$

5. SCINTILLATION MITIGATION

Section 1.2 describes the primary image degradations caused by atmospheric scintillation. The mitigation of these effects is an active area of research. Robinson[2] discussed a means to mitigate the blurring and washing induced by the atmosphere. Robinson's algorithm dramatically improves the apparent detail in the image, however the global Wiener filtering applied to sharpen the image may remove some of the blur caused by lens defocus effects, thus adversely affecting both DFF and DFD.

Delport[9] worked on mitigating the warping, or random movement, induced in the image. This work is complementary to Robinson's and may be combined to yield imagery with decreased warping and improved sharpness and contrast. Delport's work however does not affect the blur of the imagery when used alone and thus is suitable for this work. A high-level outline of the algorithm is given below, with more information available in the paper.[9]

1. Iterate through the current frame, and find all images patches with sufficient contrast. This increases efficiency by ensuring that only parts of the image where there is sufficient detail to make strong matches are processed.
2. Match each image patch to the best fit nearest neighbour from the last frame.
3. Calculate the displacement vector for each image patch.

4. Add the vectors to a history of the vectors for each patch.
5. If the direction of an image patch is constant over time, assume it is true movement in the scene and not scintillation and remove it from the list.
6. Interpolate between identified matches to create a complete de-warping mesh for the image.
7. Use de-warping mesh to create stabilised image.

6. RESULTS

The resultant depth images in this paper are small in order to conserve space, full resolution images at either 1920×1080 or 1936×1456 are available online*.

6.1 Short Range Results

Figure 1 and 2 provide the results of the laboratory experiments for DFF and DFD respectively.

Each of the 3 algorithms explained in $S3$ was run on a 341 frame video sequence captured while the focus was being adjusted from in front of the Prosilica to behind the wood back drop. Window sizes of 9, 17, 33, 65 and 129 pixels per side were evaluated. The best results for each algorithm are given in Figure 1(a). Darker shades indicate closer objects and lighter shades further objects. From the results it can be seen that the standard deviation based algorithm yields the best results. Both the contrast and entropy functions yield poor results in the vicinity of sharp edges, this can also be seen on the blurriness of the sharp edges of the cartoon character and checkerboard in Figures 1(c) and 1(d). These edge anomalies may also be attributable to minor magnification changes induced when changing the focus as the lens was not telecentric. The contrast algorithm in particular produced noisy results and appeared sensitive to vignetting, while the entropy algorithm was susceptible to producing ringing of the edges. The entropy algorithm did however yield a smooth gradient of depth for the surface on which the scene's objects were placed. It can be deduced that the three algorithm were correctly implemented and were able to discriminate depth.

Figures 2(c), 2(d), 2(e) and 2(f) show the depth maps obtained from the spatial domain implementation of Pentland's algorithm, the algorithm of Watanabe and Nayar, the frequency domain implementation of Pentland's algorithm, and Bove's algorithm respectively for the laboratory experiment. Each of the algorithms were run with window sizes ranging in integer powers of 2 from 4 to 64 pixels.

The two frequency domain algorithms provided the best results, both obtained their best results with a window size of 16 pixels and seemed to provide more depth information in the less uniform regions around the edges. Higher intensities in both algorithms indicate increased distance from the camera.

The two spatial domain algorithms yielded poorer results, the Pentland spatial results required filtering due to the division of the Laplacian-estimated high frequency content causing instability in low detail regions (this is not a problem in the frequency domain regions because the true spectral power for the high frequency band is determined rather than being estimated). The remaining data provides good results in the edges with higher intensities indicating decreased distance from the camera. The Watanabe-Nayar algorithm provides poor results and seems especially susceptible to the change in vignetting induced by changing the aperture. The provided rational filters seem to be specific to their apparatus in terms of resolution and blur size as measured in pixels. Development of custom rational filters for long range passive ranging maybe possible and beneficial but will require further investigation.

In general the DFD algorithms appear to produce poorer results than DFF and are more susceptible to vignetting, which is to be expected since only two images are used to determine the depth map. The laboratory results instill confidence in the algorithms' implementations.

*http://prism.csir.co.za/

(a) Close Focussed Image

(b) Far Focussed Image

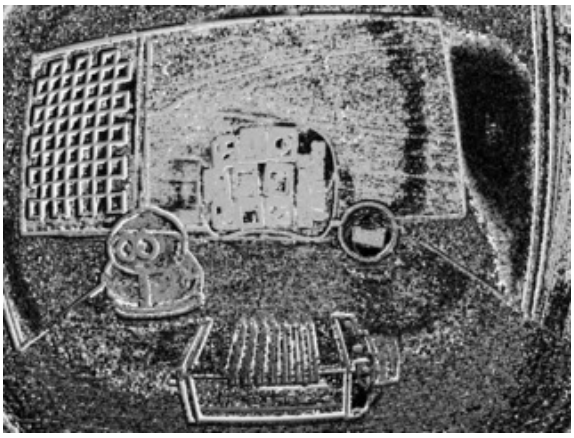

(c) Michelson Contrast Depth Map

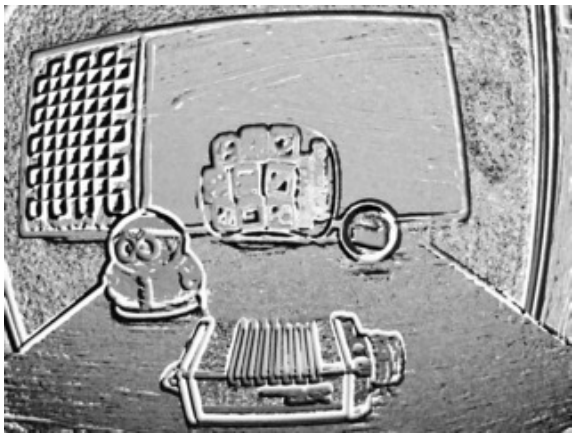

(d) Shannon Entropy Depth Map

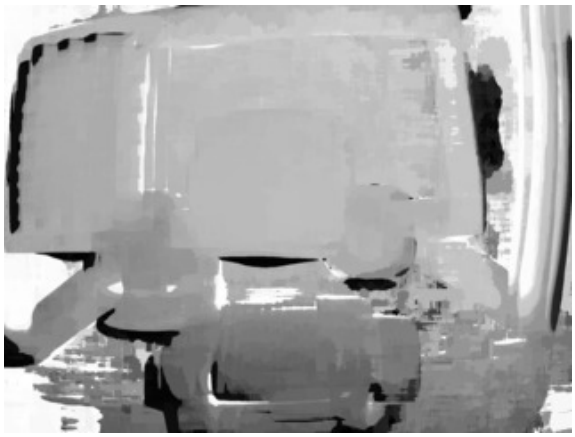

(e) Standard Deviation Depth Map

Figure 1. Short Range Depth From Focus Results

(a) Large Aperture Image

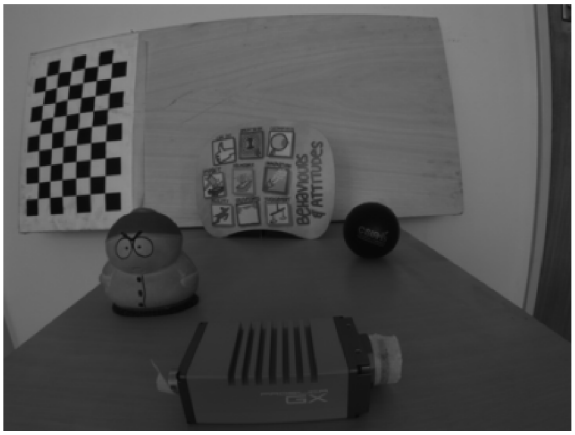

(b) Small Aperture Image Image

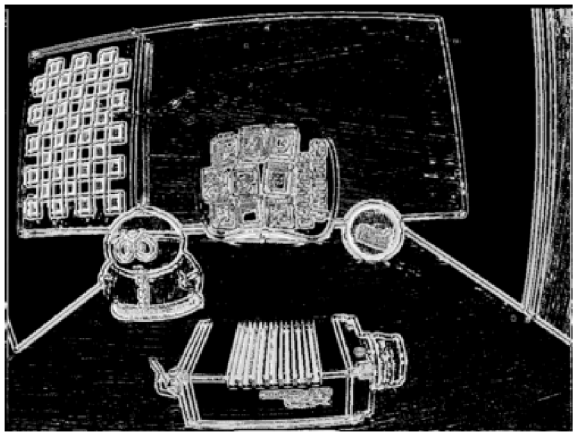

(c) Spatial Pentland Depth Map

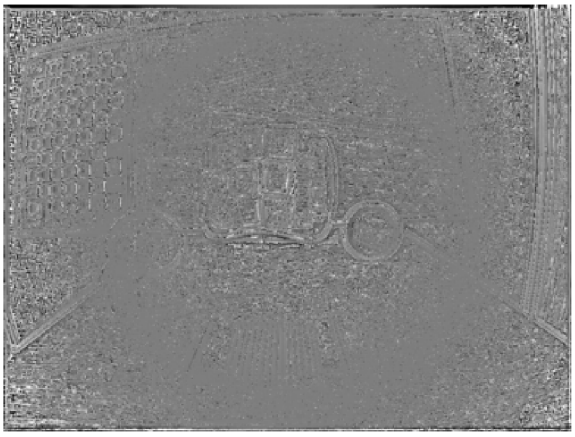

(d) Watanabe-Nayar Depth Map

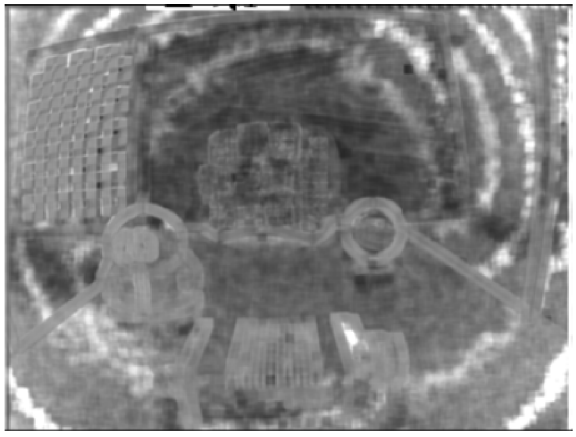

(e) Pentland Frequency Depth Map

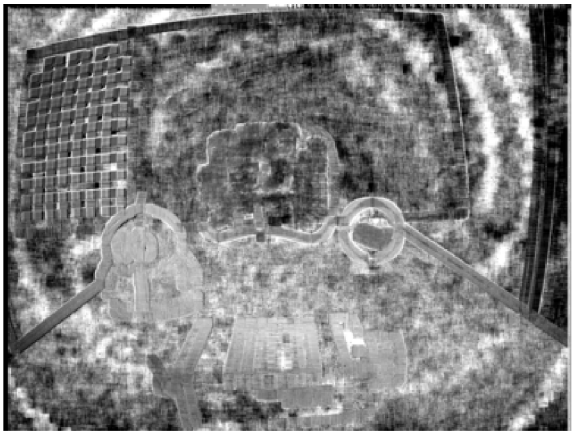

(f) Bove Depth Map

Figure 2. Short Range Depth From Defocus Results

6.2 Mid Range Results

Once the correctness of the algorithms' implementations was verified with the short range laboratory tests, mid range tests were conducted to determine the feasibility of using DFF and DFD for longer ranges. The results are given in Figure 3 for DFF and Figure 4 for DFD.

Figures 3(c), 3(d) and 3(e) show the results of running the DFF algorithms on a medium range data set. Each algorithm was again run for window sizes varying from 2 to 64 pixels in integer powers of 2 and were experimentally determined to give best results with window sizes of 8 pixels. These are the results provided.

A sequence of dots is visible on the road just below the far side in the entropy and standard deviation images. This is due to a small bird which flew across the scene near the end of the data capture. This is a fundamental flaw in DFF: the span of time that it takes to capture the data increases the probability of movement in the scene.

The three algorithms produce similar results. The main image features are very well segmented by depth. The palisade fencing in the extreme foreground is darker than the carports and the aerial of the vehicle parked beneath them. The rear of the signboard is then next closest followed by the corner of the container and finally the poles of the distant storage area. The information is best where there is either strong texture or sharp edges. The standard deviation and entropy algorithms provide the best results.

Figures 4(c), 4(d), 4(e) and 4(f) show the depth maps obtained from the spatial domain implementation of Pentland's algorithm, the algorithm of Watanabe and Nayar, the frequency domain implementation of Pentland's algorithm, and Bove's algorithm respectively for the mid range experiment. Each of the algorithms were run with window sizes ranging in integer powers of 2 from 4 to 64 pixels, with 16 pixels providing the best results.

The frequency domain algorithms yielded superior results. The segmenting and depth separation of these algorithms is indeed very good and can be attributed to the high quality large aperture optics used in the custom 185mm focal length camera. The over and under exposed areas yield anomalous results, suggesting that high dynamic range imaging would be beneficial. In the spatial domain Pentland's algorithm noticeably out performs the Watanabe-Nayar algorithm. Pentland's algorithm provides a good indication of depth at the edges of objects, whereas Watanabe-Nayar produced little discernible depth discrimination.

6.3 Long Range Results

This section presents the early morning long range results captured shortly after sunrise before the scintillation started to become observable. This can thus be seen as a feasibility study for long range passive ranging after the successful mid range tests.

Figure 5 provides the results for DFF at 650m and at 1050m. The analysis of the results is complex. First it is worth noting that there is a general trend along the mowed runway of light at the bottom and dark at the top in keeping with the true distance. As there is no texture in the sky the distances are arbitrary according to sensor noise and contain no information. In the long grass to the right of the runway, the time elapsed during the data capture (approximately 30 seconds), causes some movement induced noise in the determined results. However if one locates the scintillometer emitter in the depth map, it is a clear discontinuity and the cause of the blooming creating the vertical bars evident in all the images. The range to the scintillometer is of course incorrect - being affected by the saturation of the camera pixels. Adjacent to the scintillometer is the bakkie, it is discernible from the background by a few shades of grey. This is particularly apparent in the standard deviation and entropy images. This is the desired result, as it is not vastly closer than the background which it obscures which is few dozen meters more distant.

(a) Close Focussed Image

(b) Far Focussed Image

(c) Michelson Contrast Depth Map

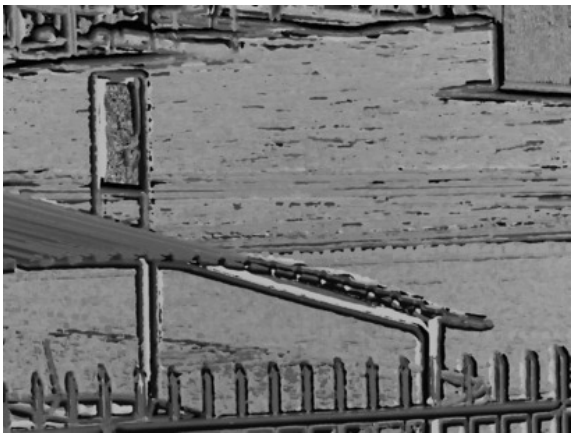

(d) Shannon Entropy Depth Map

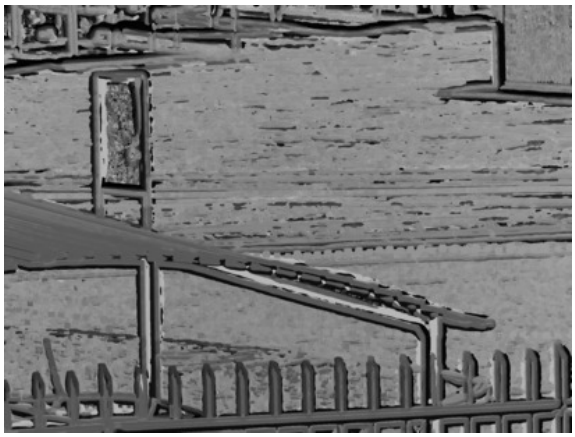

(e) Standard Deviation Depth Map

Figure 3. Mid Range Depth From Focus Results

(a) Large Aperture Image

(b) Small Aperture Image Image

(c) Spatial Pentland Depth Map

(d) Watanabe-Nayar Depth Map

(e) Pentland Frequency Depth Map

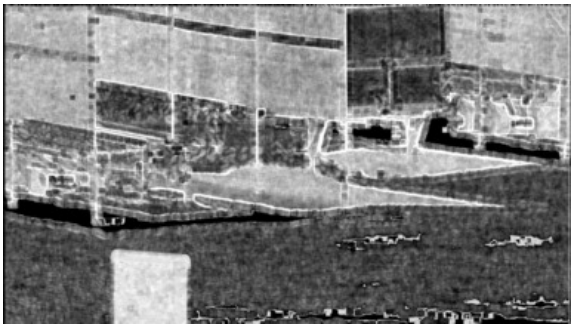

(f) Bove Depth Map

Figure 4. Mid Range Depth From Defocus Results

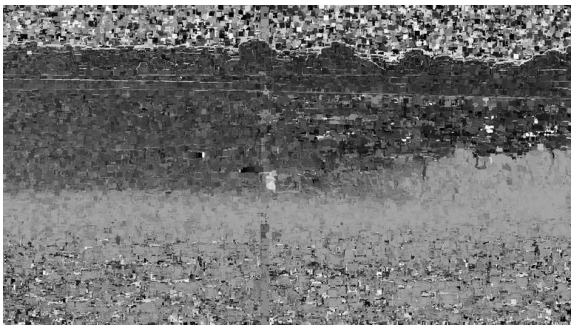

(a) 650m Michelson Contrast Depth Map

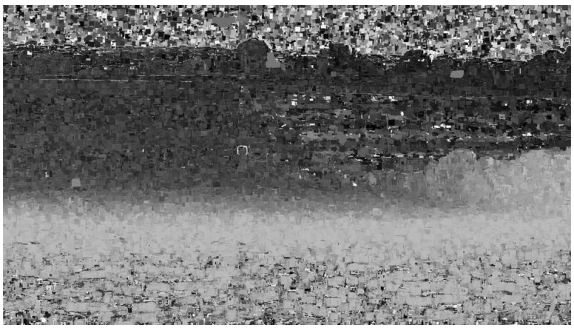

(b) 1050m Michelson Contrast Depth Map

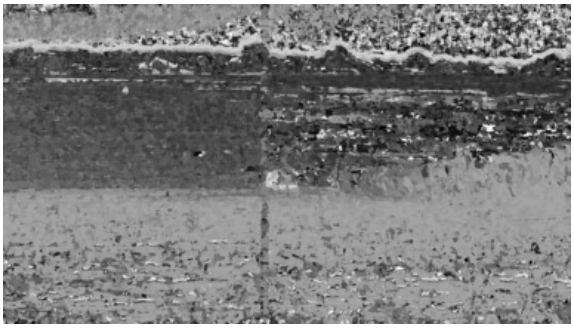

(c) 650m Shannon Entropy Depth Map

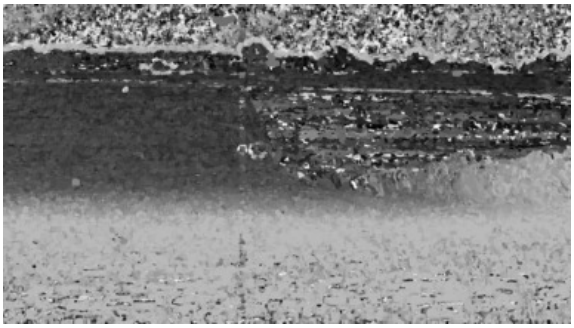

(d) 1050m Shannon Entropy Depth Map

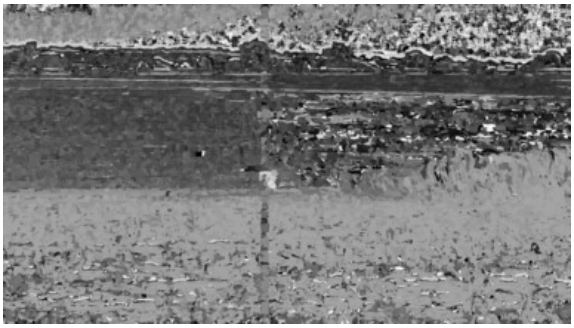

(e) 650m Standard Deviation Depth Map

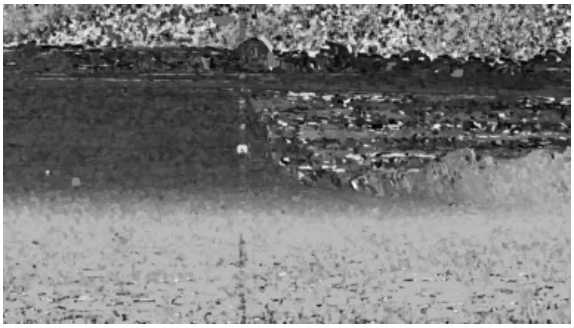

(f) 1050m Standard Deviation Depth Map

Figure 5. Long Range, Depth From Focus Results

Figures 6 and 7 provide the DFD results for a distance 550m and 1150m respectively. Due to the generic nature of the field trial, which aimed to capture data applicable to several projects, the images were captured many seconds apart. This is apparent from Figures 6(a) and 6(a) where the foreground bird has changed position and the vehicle driver has exited and is standing alongside the scintillometer in Figure 6(a). For both long range input sets the four DFD algorithms were run with window sizes ranging from 4 to 64 in integer powers of 2.

In the analysis of the 550m range data, it was found that for all of the algorithms the best results were obtained with a window size of 16 pixels. Fewer pixels created a noisy depth map, and more pixels created a smoother image at the expense of loss of shape and structure of the scene and its contents. The two frequency domain algorithms far outperformed the spatial domain algorithms. Both frequency algorithms identified the two new objects in the scenes as being extremely close to the camera, leading to potential usefulness as change detection algorithms for automated surveillance with the Pentland algorithm's results (Figure 6(e)) slightly outperforming Bove's algorithm's results (Figure 6(f)). They also both created a smooth depth map correctly showing the contour of the grass runway with both the scintillometer and vehicle segmented from the background and only a few shades of grey different from the background. The spatial algorithms performed poorly, the spatial

(a) Large Aperture Image

(b) Small Aperture Image Image

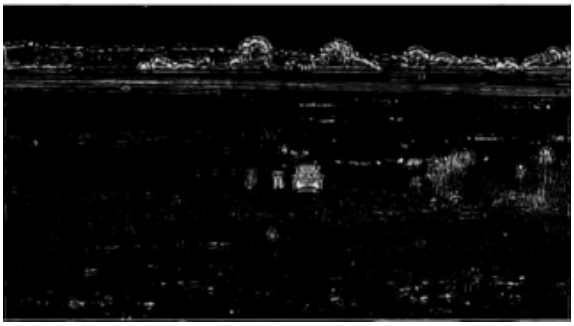

(c) Spatial Pentland Depth Map

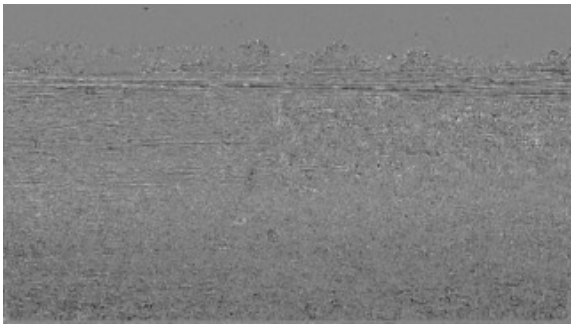

(d) Watanabe-Nayar Depth Map

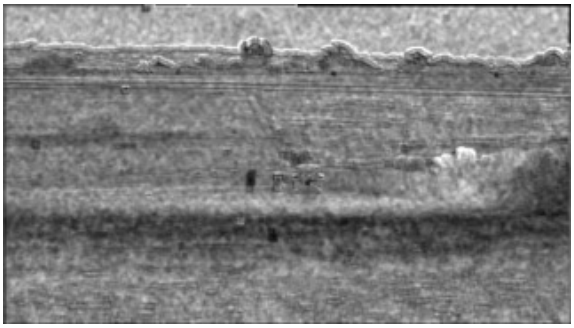

(e) Pentland Frequency Depth Map

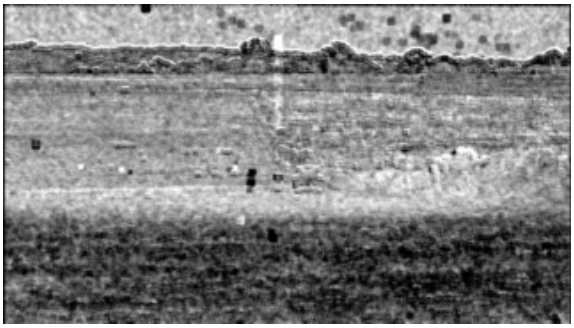

(f) Bove Depth Map

Figure 6. Long Range 550m, Range Depth From Defocus Results

(a) Large Aperture Image

(b) Small Aperture Image Image

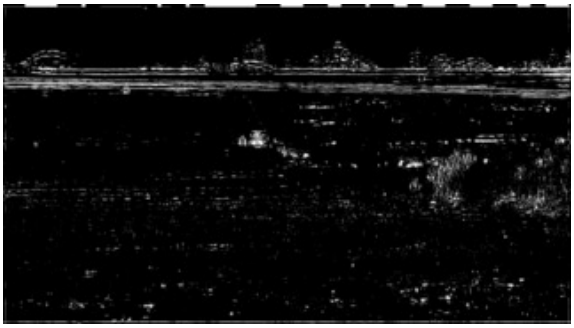

(c) Spatial Pentland Depth Map

(d) Watanabe-Nayar Depth Map

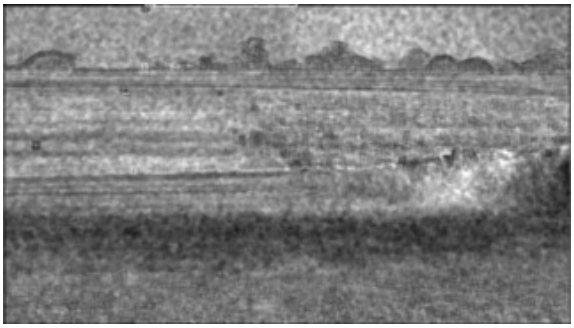

(e) Pentland Frequency Depth Map

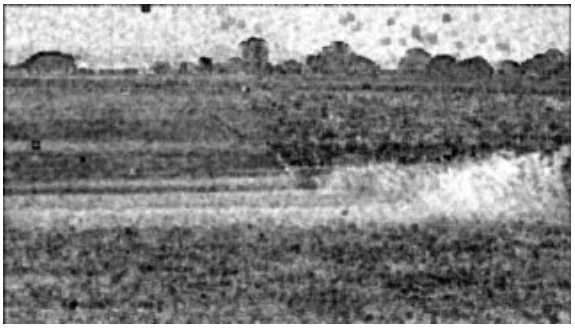

(f) Bove Depth Map

Figure 7. Long Range 1150m, Range Depth From Defocus Results

Table 1. Scintillation Metrics

Day	Run	Scintillation Description	Distance (m)	Scintillometer Measurement C_n^2 ($m^{2/3} \times 10^{-13}$)
1	2	Severe	550	7.6
			750	11.3
			950	18.0
			1150	20.1
1	3	Mild	550	1.81
			750	0.50
			950	3.81
			1150	0.22
2	1	Moderate	550	0.22
			750	2.07
			950	5.24
			1150	5.86

Pentland algorithm could be used to detect potential objects of interest but no depth information is apparent in the intensity of the edges, see Figure 6(c). The Watanabe-Nayar algorithm produced only noise (Figure 6(d)), thereby underlining the need for application-specific filter kernel development if it is to be used.

The 1150m results in Figure 7 yield results similar to those of the 550m experiments. However there is a noticeable variation in the illumination in the scene between the large (Figure 7(a)) and small (Figure 7(b)) aperture images due to the time delay and moving clouds. The illumination in the vicinity of the vehicle remained fairly constant. Once again the frequency algorithms outperformed the spatial algorithms. This time however the frequency Pentland algorithm's results (Figure 7(e)) was noticeably superior to the results of Bove's algorithm (Figure 7(f)) which seems sensitive to illumination variances. The spatial domain implementation of Pentland's algorithm again identified salient scene features without yielding depth information (Figure 7(c)). Watanabe and Nayar's algorithm did not yield useful information as seen in Figure 7(d).

6.4 Scintillation Results

The long range field trial did not capture sufficient data to evaluate the sensitivity of DFF with regards to scintillation. This would have required capturing a short video sequence at each focus setting so that a dewarped image could be produced. However video sequences of the (stationary) target vehicle on the grass runway were captured with apertures of F2.3, F5.0 and F11.0. The first and last of these can be used to evaluate the sensitivity of DFD with regard to atmospheric scintillation. Preliminary results are presented.

Each video was dewarped using Delport's algorithm[9] as described in §5. The 100th frame of the dewarped sequence was then used as the input to the DFD algorithms. Sections 6.1, 6.2 and 6.3 describe how the two frequency domain algorithms were found to be superior to the spatial algorithms in each of the short, mid and long ranges respectively. Therefore, only Bove's algorithm and the frequency domain version of Pentland's algorithm were evaluated for their scintillation sensitivity. Table 1 provides the average scintillometer measurements recorded while the video was being captured. The values are not always monotonic as even between consecutive measurements (provided at 30 second intervals) the values can vary by 50%. This is at least partly due to the scintillation being affected both by crosswinds and by turbulence along the optical path. Typically changes of C_n of an order of a magnitude are required to have a discernible affect.

Figure 8 shows the results of the algorithms on ranging the vehicle at 1.2km by comparing the dewarped F2.3 and F11.0 aperture videos. Figure 8(a) corresponds to the Day 1, Run 3, Position 12 in Table 1. Figure 8(c) corresponds to Day 2, Run 1, Position 12 and Figure 8(e) corresponds to Day 1, Run 2, Position 12. In no case is the vehicle discernible from the background of the scene.

(a) Bove Results, Low Scintillation

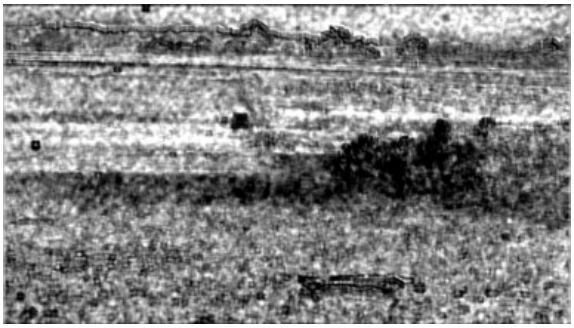

(b) Frequency Pentland Results, Low Scintillation

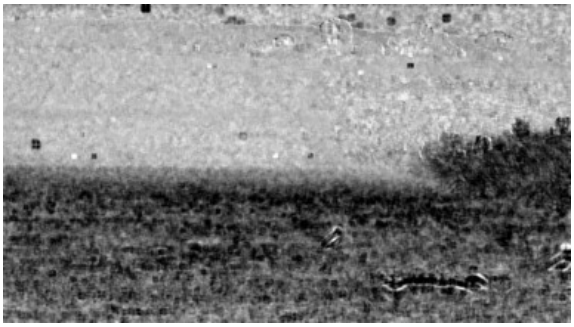

(c) Bove Results, Moderate Scintillation

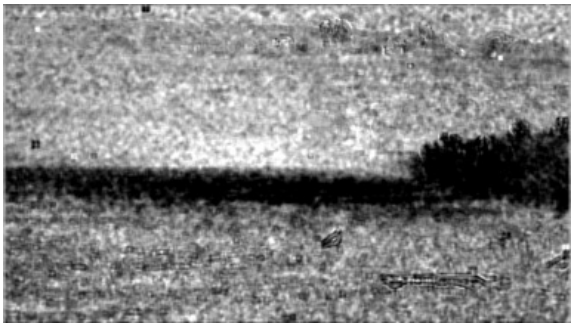

(d) Frequency Pentland Results, Moderate Scintillation

(e) Bove Results, High Scintillation

(f) Frequency Pentland Results, High Scintillation

Figure 8. DFD Scintillation sensitivity tests at 1150m

Figure 8 also shows the results of Pentland's algorithm on the same data in the same order. The vehicle is discernible in both the medium and high scintillation, and less so in the low scintillation data, which was where Bove's algorithm provides its strongest result. This suggests a scintillation metric based weighted sum of the two algorithms that favours Bove's algorithm for low scintillation conditions and favours Pentland's for high scintillation conditions would provide superior results in all conditions.

7. CONCLUSIONS

This work showed that it is possible to use focus based passive range techniques at ranges far beyond those found in the laboratory experiments which are typically reported. Both ranges of up to 100m with standard lenses and over 1km with large aperture optics were tested.

Depth from focus provides results superior to depth from defocus due to the increased input data but suffers from requiring additional data acquisition time. This limits DFF's usefulness in dynamic scenes. The standard deviation based focus metric provided the best results.

Of the depth from defocus algorithms evaluated, the frequency based techniques performed far better. The frequency domain implementation of Pentland's algorithm provided slightly better results than Bove's algorithm, particularly when subjected to higher atmospheric scintillation. Both frequency algorithms detected change in the scene well.

The standard Pentland algorithm provided good distance information at edges in the scene for short and mid range results and identified salient features if not depth at long ranges. The standard filters suggested by Watanabe and Nayar performed poorly, suggesting that new kernels should be developed to match the characteristics of the large aperture lens and the higher resolution cameras used.

Preliminary results for performing DFD on high scintillation data that had been processed to reduce the warping of the image were presented. The results were positive and indicate that further research in this area could yield a practical system.

Further work could include capturing a short video clip at every focus position to evaluate the performance of DFF in scintillation conditions. Improvements in accuracy of both DFF and DFD could be made, by increasing the speed at which data is captures, testing more focus settings, increasing the focal length and using even larger aperture optics.

REFERENCES

[1] Pentland, A. P., "A new sense for depth of field," *IEEE Trans. Pattern Anal. Mach. Intell.* **9**, 523–531 (July 1987).

[2] Robinson, P., Walters, B., and Clarke, W., "Sharpening and contrast enhancement of atmospheric turbulence degraded video sequences," in [*Proceedings of the 21st Symposium of the Pattern Recognition Society of South Africa*], *PRASA 2010* **1**, 247–252 (2010).

[3] Michelson, A., [*Studies in optics*], Dover Publications (1995).

[4] Shannon, C. E. and Weaver, W., [*A Mathematical Theory of Communication*], University of Illinois Press, Champaign, IL, USA (1963).

[5] de Villiers, J. P., "A comparison of image sharpness metrics and real-time sharpening methods with GPU implementations," in [*AFRIGRAPH '10*], 53–62, ACM, New York, NY, USA (2010).

[6] NVIDIA-Corporation, "NVIDIA CUDA programming guide version 4.0."

[7] Bove, V. M., "Entropy-based depth from focus," *Optical Society of America* **10**(4), 561–566 (1993).

[8] Watanabe, M. and Nayar, S. K., "Rational filters for passive depth from defocus," *International Journal of Computer Vision* **27**, 203–225 (1998).

[9] Delport, J. P., "Scintillation mitigation for long-range surveillance video," in [*Proceedings of the 3rd CSIR Biennial Conference: Science Real and Relevant*], (2010).