Anisoplanatic Lucky Imaging for Surveillance

Sijiong Zhang, Frank F. Suess, and Craig D. Mackay

Institute of Astronomy, University of Cambridge, UK, CB3 0HA

Abstract

We have developed a system capable of delivering surveillance images with much greater detail than normally possible. It uses a high-speed camera to freeze the atmospheric turbulence and then selects small parts of each image when they are at their sharpest to synthesise a complete output image. We have demonstrated improvements in resolution of up to 4 times, depending on the percentage of frames used in the final image. The principal limitation is the poor sensitivity of current high-speed camera systems. The development of noiseless, high-speed electron multiplying CCD camera systems suggests that these limitations may be overcome in the future.

OCIS codes: 010.1290, 100.0100 and 100.3020.

1. Introduction

The importance of surveillance imaging is undoubtedly growing. Ultimately the image quality of any surveillance system is limited by atmospheric turbulence. Atmospheric turbulence is generally accepted to have a broadly power law dependence with most power being present in the largest turbulent scales. In general, turbulence on ground-ground paths is much worse than that of high-altitude turbulence experienced by astronomers [1]. With a high-quality surveillance system working at maximum magnification the appearance of the image is dominated by the apparent shaking of the image which is a consequence of the power on large scales manifesting itself as a tip-tilt error. When the turbulence is bad the tip-tilt errors appear less important as the

nature of turbulence becomes characterised by a boiling impression corresponding to a rubber sheet type distortion across the image. On a small scale the sharpness of the image is highly variable. Most people are familiar with this. On a hot day ground layer turbulence can make images particularly poor yet there are often moments the images become much sharper. A number of attempts have already been made to address the problem of achieving turbulent motion-blur restoration. Fraser et al [2] used a technique of estimating the point by point distortions in images and correcting each image accordingly. Vorontsov and Carhart [3] made a series of laboratory experiments using artificial turbulence layers to investigate the problems of anisoplanatic imaging through a turbulent atmosphere. Our programme emphasises the development of a real system that can be used under real conditions to deliver genuine performance improvement for surveillance system users.

This paper looks at techniques of high-speed framing together with image selection that deliver much higher resolution images than are possible with conventional imaging system. Similar Lucky Imaging techniques have been used for a number years by astronomers [4, 5] where it is used to deliver Hubble Space Telescope resolution from the ground. Astronomical Lucky Imaging provides images with a much larger isoplanatic patch angle than is found with adaptive optics techniques. Selecting the sharpest images corresponds to selecting the images with the best isoplanatic patch angle. Astronomers usually find that under reasonable conditions they are dealing with a single turbulent layer of finite thickness whereas with ground based imaging along a horizontal path the turbulence is distributed throughout the path in general. Further the linear scale size of the turbulence (r_0) is typically one order of magnitude smaller as is the coherence time (t_0). This makes the correction of turbulence much more difficult to achieve for surveillance purposes. These difficulties may be minimised by careful selection of

telescope diameter, by working as far into the red end of the spectrum as the camera system allows and by using a high-speed, high sensitivity camera for recording the images.

The basic principle we use follows others such as Vorontsov [6] which is to take images fast enough so as to freeze the turbulence. The images are divided into overlapping cells and the image quality within each cell is assessed from frame to frame. The best images within each cell are then registered with one another and with the best images from adjacent cells to allow a much higher quality output image to be synthesised. The choice of the cell size and a frame rate of the camera system must be chosen to match the turbulent conditions. Here we have applied these methods to conditions much closer to those encountered in real life where the criteria for establishing sharpness are less well defined.

2. Experimental setup

In order to characterise the turbulence we built a simple array of light emitting diodes that we imaged with 203 mm diameter Meade telescope to which a VDS-Vosskuhler HCC-1000 high-speed CMOS camera and band limiting filter was attached. The camera had 1024 x 1024 pixels each 10 μ in size and a 2 m focal length of the telescope gave a pixel scale of approximately 1 arc second per pixel. The quantum efficiency of the detector in varies between 20 and 40% over the wavelength ranges with which we experimented. Most of our work was carried out at 875 nm.

The measurements were performed at Lord's Bridge, about 8 kilometres southwest of Cambridge in England at the University's Mullard Radio Astronomy Observatory. The first group of measurements used a baseline of 1.2 km with the telescope about 3 m above the ground. The target was approximately 2 m above the ground. This site is flat and surrounded by open fields. The light emitting diodes were 4 mm in diameter corresponding to about 0.7 arc seconds

at 1.2 km range and were separated by 40 mm. The camera could be run at rates between 60 and 500 frames per second, but only for the duration of 1000 frames. Data were taken on more than 30 separate occasions under different weather conditions generally representative of conditions in Cambridge. Initial tests were carried out at three LED wavelengths of 525 nm, 644 nm and 875 nm. The expected wavelength dependence with the turbulence strength decreasing in proportion to wavelength $\lambda^{1.2}$ was observed so all further tests were restricted to the longest wavelength of 875 nm.

The important parameters we wished to measure included the Fried parameter, r_0 given by:

$$r_0 = 0.98\frac{\lambda}{FWHM} \tag{1}$$

where λ is the wavelength and FWHM (in radians) is the full width at half maximum of a long exposure image simply obtained by summing the images from a single run. We also measured the temporal correlation function C(t) by measuring the vector displacements $\mathbf{p}(\tau)$ of the images of one LED according to:

$$C(t) = \frac{<\mathbf{p}(\tau)\cdot\mathbf{p}(\tau+t)>}{<|\mathbf{p}(\tau)|^2>} \tag{2},$$

where t denotes the time interval between two correlated speckle images and $<\ldots>$ indicates the ensemble average or time average. Coherence time t_0 was defined as the time corresponding to a *1/e* decay of the C(t) values.

Another important parameter is the isoplanatic angle, the angular distance over which an instantaneous set of turbulent parameters hold. In this case we measure the degree of isoplanicity of two adjacent images by calculating the normalised cross-correlation function as:

$$C(x) = \frac{< \boldsymbol{p}(x',t') \cdot \boldsymbol{p}(x'+x,t') >}{< \left| \boldsymbol{p}(t') \right|^2 >} \qquad (3),$$

where x denotes the angular separation of between two point sources.

In each case the measurements were made on a frame to frame basis and the distribution of values obtained used to guide the image selection criteria. The isoplanatic angle θ was defined as the angle corresponding to a *1/e* decay of the C(x) values.

In all cases it is important to realise that although we are using parameters which we call t_0 and r_0 what we measure are averages over the thick turbulent layer between the target and the imaging telescope.

3. Measurement of the turbulence parameters r_0 and t_0

The measurements presented in this paper were collected in the middle of the day between 4 March to 23 May 2006. The maximum temperatures in March, April and May in Cambridge were 9.0, 13.5 and 17.4 degree Celsius, respectively. The values of r_0 vary between 75 and 11 mm diameter, with an average of 24mm and t_0 varied between 33 and 3 ms with an average value of about 8ms. There was a clear correlation between r_0 and t_0 in our data in the sense that both become smaller when the weather is warmer. These results are summarised in figure 1.

4. Anisoplanatic Lucky Imaging strategy

It is clear from the data presented in the previous section that atmospheric conditions are highly variable. Not only does the nature of the turbulence characterised by t_0 and r_0 vary from day to day, it also varies on very short timescales. We know from our experience of astronomical turbulence that conditions change on essentially all timescales. In addition, most ground-ground imaging depends on natural illumination so that frequently, particularly in the

British climate, the signal level detected can change on very short timescales indeed. This means that we need to be able to accommodate not simply changes in the turbulent characteristics of the atmosphere but also in illumination levels and therefore in the signal-to-noise ratio of any features we might choose to guide our image selection.

Our basic strategy is to capture a time sequence of images at high resolution and high-speed so that the motion that normally blurs the recorded images is eliminated as far as possible. The parts of the images that are unusually sharp are combined to form an output image significantly sharper than would be achieved with a conventional imaging system.

For anisoplanatic imaging there is no space-invariant point spread function over more than a small part of the field of view since light rays originating from different directions propagate through different parts of atmosphere and encounter different perturbations. Therefore, we work within an isoplanatic window so that we can deal with each small part of the field of view at a time. The size of the isoplanatic window must be chosen to match the actual turbulent conditions we encounter.

Each windowed section within the image sequence is searched for the sharpest images which are then shifted and added into the output image. Within each frame, the sharpness of each section will be essentially uncorrelated with the sharpness of others.

5. Image processing

After capturing a sequence of images $I_i(x, y)$ from the high-speed camera, we carry out the following procedure to construct a sharper image $R(x, y)$. The intensities of each frame of the image sequence are normalised to match the sequence average $\bar{I}(x, y)$. After setting the size of the isoplanatic window W appropriately for the turbulent conditions current at the time the

image sequence was captured, we calculate the Fisher information $F(I_i(x_w y_w))$ for each of the sub-images in the window given by:

$$F(A) = \sum_{x,y} \frac{\left|\nabla \psi(x,y)\right|^2}{\psi(x,y)} = 4\sum_{x,y} \left|\nabla \sqrt{\psi(x,y)}\right|^2, \quad \psi(x,y) = \frac{I_i(x,y)}{\sum_{x,y} I_i(x,y)}$$

The Fisher score allows us to select the sharpest given percentage of sub-images. Each of these sub-images is shifted into alignment with the sequence average $\bar{I}(x,y) = \frac{1}{N}\sum_{i=1}^{N} I_i(x,y)$ and added, allowing the output image $R(x,y)$ to be synthesised after renormalizing the mean intensities. At this stage overlapping parts of windows if the exist are appropriately rescaled.

For image sequences with a sufficiently large number of frames , the sequence average is equal to the long-exposure image. Every pixel in the output image must be present in at least one isoplanatic window and by definition, the window size should be comparable to the size of the lucky isoplanatic patch. Other than that, we may choose any shape and position for the windows that we think is appropriate. Rectangular windows are the simplest and most efficient shape to implement. All of the results presented in this paper were calculated using square, non-overlapping windows.

6. An experimental example of anisoplanatic imaging

We constructed an extended target composed entirely of English text with 33mm inter-line spacing. We used text because it has a uniform sharpness across the entire scene and human users are very familiar with the recognition of text. The imaging setup is the same as outlined earlier. In this case the distance between the target and telescope was about 400m. A long-pass near-infrared filter with cut-off wavelength 860nm was introduced between the telescope and

camera. The combined effect of the filter and the CMOS sensor's quantum efficiency as a function of wavelength results in an effective measurement wavelength around 875nm. The angle subtended by the target at the telescope is about 800 arcseconds. This field of view is much bigger than the lucky isoplanatic angle for even moderate turbulence. Images were captured at about 150 frames/second with an exposure time of 6.5 milliseconds. In order to provide a more quantitative evaluation of the imaging performance we added a single LED to the target to act as a point source. Sequences of images were captured as described above. In order to understand the trade-off between output image resolution and selection percentage, a number of output images were generated with selection percentages ranging from 0.1% to 100%. Note that with 1000 frames in an image sequence, 0.1% means selecting only the sharpest image. Also note that the 100% image will be sharper than the long-exposure image as each frame is shifted and aligned before being added into the 100% image. This gives tip-tilt correction of the imaging system, removing the largest turbulent scales, and camera shake if present.

Figure 2 shows a series of images to illustrate how the various percentage selections were able to deliver improvements in image resolution. Figure 2a is the sequence average, which is roughly what you would see in a long-exposure image. Figures 2b, 2c and 2d are lucky images produced with 10%, 1% and 0.1% selection percentages respectively. Figure 2e shows the frame in the image sequence with the highest measure of sharpness, It illustrates how variable image resolution can be, even within the sharpest frame. Selecting just the best short-exposure image is unlikely to achieve a reliable or consistent improvement in image resolution.

One of the most tricky parts of this work was establishing a quantitative measurement of the improvement in image sharpness. The application of these systems to surveillance means that they will be examined and assessed by human observers. The human vision system is very

good at recognizing sharp features in the presence of an extended diffuse background. Further, we know from our experience in astronomy that lucky imaging has the effect on the point spread function of pulling flux from the extended profile in to a sharp central core, as well as generally reducing the width of the point spread function. Visual assessment of astronomical lucky images tends to exaggerate the importance of these cores giving an impression of sharpness that is less evident when using full width at half maximum, for example, as a sharpness criterion. Astronomical turbulence is predominantly high-altitude where an essentially fixed turbulent screen moves across the aperture of the telescope driven by wind. Ground-ground surveillance through ground layer turbulence may be rather different because of the clear importance of turbulent motions in the atmosphere along the imaging path. Nevertheless we would expect ground-ground lucky image profiles also to show a significant core which would be a significant contribution to the impression of sharpness.

A useful way of describing the extent of improvement in image sharpness is to equate this improvement to what would be obtained by a specific reduction in the distance between camera and target. We have tried to assess this in several different ways. The most useful appears to be to convolve the lucky image with an appropriate function to degrade the Fisher information to that of the long exposure image. For simplicity we used a Gaussian smoothing function. The results are shown in figure 3 where the Fisher information is plotted against the half width of the Gaussian function used to convolve each of the different lucky images. Figure 3 shows graphs of the Fisher value plotted against the FWHM of the convolving Gaussian function for each image. The lowest curve is the graph for the sequence average. The point at which the Fisher information is reduced to a level that equals the long exposure value allows an estimate of the equivalent range reduction we might associate with each lucky image.

The interpretation of figure 3, however, is further complicated because the Fisher information from the long exposure (bottom curve) contains a contribution from detector nonuniformities that is greatly attenuated in all the other images that are assembled after a shift and add procedure. Further, the imaging system is somewhat undersampled (we use approximately 1 arcsecond pixels with a diffraction limit at 870 nm of about one arcsecond). Nevertheless it appears that the improvement in image sharpness corresponds to a factor slightly under 2 when using 100% of the frames (simply by removing the largest turbulent scale) and an additional factor of up to 2 when using less than 1% of the images to construct the output image.

An independent assessment of image resolution may be had from the calibration LEDs present at the top of each image. The LEDs subtend an angle of approximately 2 arcseconds at the range of 400 m used in these sequences. We can measure the full width at half maximum of these images. These results are shown in figure 4 which plots the LED cross-sections for 0.1%, 1%, 10% and 100% selection and for the sequence average. The image for 100% selection shows the effect of tip-tilt correction alone. The FWHM ranges from about 5 arcseconds for 0.1% selection to 9 arcseconds for 100% selection. These figures should be compared with a long-exposure value of 13 arcseconds obtained from the sequence average. Again we find an improvement of a factor slightly under 2 when using 100% of the frames but a smaller additional factor of about 1.3 when using less than 1% of the images. This apparently poorer improvement in resolution when using a FWHM measure can be understood since this gives no particular weight to the sharp central core in the image profile that we expect from lucky imaging strategies.

It is clear from the images presented in figure 2 that the legibility of the target text is dramatically improved by lucky image selection. Our overall conclusion therefore is that these

lucky imaging strategies can deliver an improvement in image resolution of perhaps a factor of four overall.

The results we obtained are consistent with our experience of astronomical turbulence and the use of lucky imaging techniques to minimise its effects [4,5]. A good guide to ground-based turbulence is that it is typically an order of magnitude worse in terms of scale sizes and correlation times than is found when looking up through a high altitude turbulent layer. In astronomy we find that an ideal telescope diameter is about 2.5 m, approximately 10 times bigger than the telescope we use here. As in astronomy there are different regimes where these techniques work in rather different ways. If the turbulence is very low then considerable improvements may be achieved. If the turbulence is on very small scales then the tip tilt compensation produces very little improvement since it is principally smaller cells within the image that are tipping and tilting and the net tip tilt of the whole image is generally small on average. Under these conditions it is important to reduce the diameter of the telescope to ensure that the number of turbulent cells across its diameter is kept to a manageable level. Since turbulence levels change on relatively short timescales, it is important in any practical system to be able to reconfigure the telescope diameter, camera frame rate and isoplanatic window size on relatively short timescales.

7. Technological limitations

Undoubtedly the biggest problem we encountered was the amount of light that we had to work with. As we have seen, ground-based turbulence works on time scales that are typically an order of magnitude shorter than high altitude turbulence. This calls for correspondingly higher frame rates. Increasing the frame rate means shorter exposure times and less signal per frame. Indeed our experience was that we only got satisfactory signal-to-noise images under bright

summer conditions and preferably in full sunlight. There are many circumstances where it is entirely practical to illuminate the target with a floodlight, particularly when working in the near infrared beyond the limit of human vision. Nevertheless it is important to look at alternate technologies to deliver the speed and sensitivity that this procedure calls for.

Recently there have been considerable improvements in the performance of CCD cameras working at extremely low light level. Both E2V Technologies (Chelmsford, UK) and Texas Instruments have developed electron multiplying CCDs which are able to work at high speed without any significant readout noise being added to the image. At present these devices are only available in limited formats where the frame rates required for surveillance work can only be achieved if a small part of the image area is read out. Clearly faster devices could be manufactured with multiple output ports. Our estimates of the performance level of the current CMOS camera used in these tests suggest that a properly designed electron multiplying CCD camera system could give a 30 fold improvement in sensitivity.

8. Conclusions

We have developed an imaging system that achieves dramatic improvements in image resolution in circumstances where the resolution would otherwise be limited by the blurring and distortion caused by atmospheric turbulence. In the examples given we have been able to reduce the effective range of a target by factors of up to 4 by processing images to select only those parts of each image that are of the highest quality and recombining them into an output image.

The process is highly computer intensive and to be effective will require significant investment in the development of real-time image processing hardware and the associated software. Nevertheless for surveillance applications the methodology is already applicable.

Acknowledgement

This research was funded by EPSRC under its Crime Prevention and Detection Programme.

References

[1] Clare E. Dillon, Heather I. Campbell, Sijiong Zhang and Alan H. Greenaway, “Anisoplanatic atmospheric measurements for terrestrial imaging,” Proc. of the SPIE, Vol. 6018, 271-278 (2005).

[2] Donald Fraser, Glen Thorpe and Andrew Lambert, “Atmospheric turbulence visualization with wide-area motion-blur restoration,” J. Opt. Soc. Am. A 16, 1751-1758(1999).

[3] Mkhail A. Vorontsov and Gary W. Carhart, “Anisoplanatic imaging through turbulent media: image recovery by local information fusion from a set of short-exposure images,” J. Opt. Soc. Am. A 18, 1312-1324(2001).

[4] Lucky Imaging website at: http://www.ast.cam.ac.uk/~optics/Lucky_Web_Site/index.htm

[5] C. D. Mackay, J. Baldwin, N. Law, P. Warner, “High resolution imaging in the visible from the ground without adaptive optics: new techniques and results,” Proc. of the SPIE Vol. 5492, 128-135(2004).

[6] Mkhail A. Vorontsov ,J. , "Parallel Image Processing Based on an Evolution Equation with Anisotropic Gain: Integrated to Electronic Architectures" J. Opt. Soc. Am. A 16, 1623-1637(1999).

[7] J. Green and B. R. Hunt, “Improved Restoration of Space Object Imagery,” J. Opt. Soc. Am. A 16, 2859-2865(1999).

[8] Vladan Rankov, Rosalind J. Locke, Richard J. Edens, Paul R. Barber and Borivoj Vojnovic, “An algorithm for image stitching and blending,” Proc. of the SPIE Vol. 5701, 190-199(March 2005).

Figures and Captions

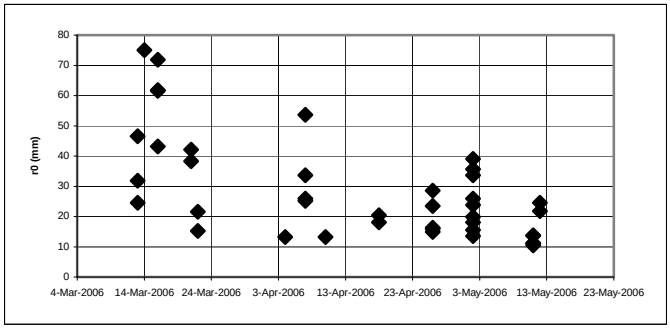

Fig.1 (a) r_0 over the observing period

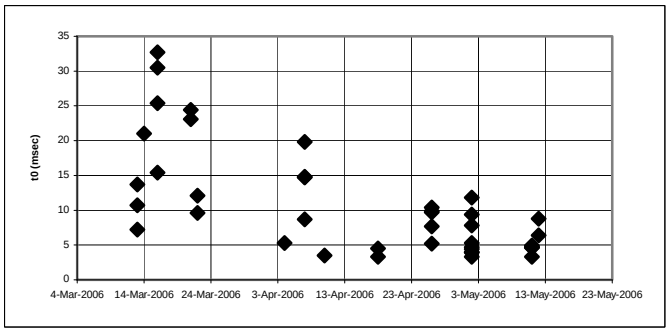

Fig.1 (b) t_0 over the observing period

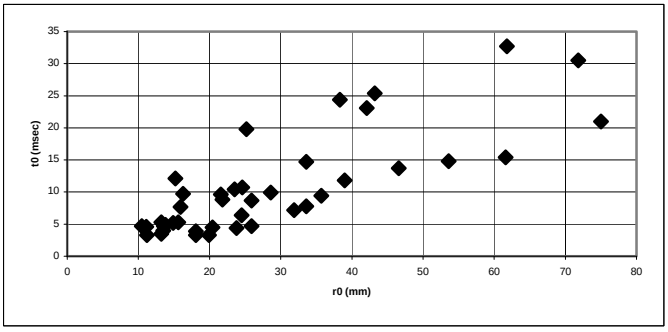

Fig.1 (c) comparison of t_0 and r_0

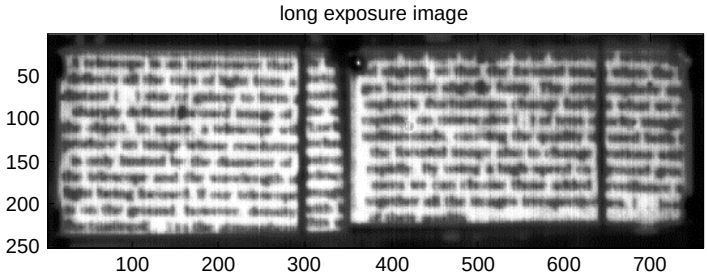

(Fig 2a, Sequence average)

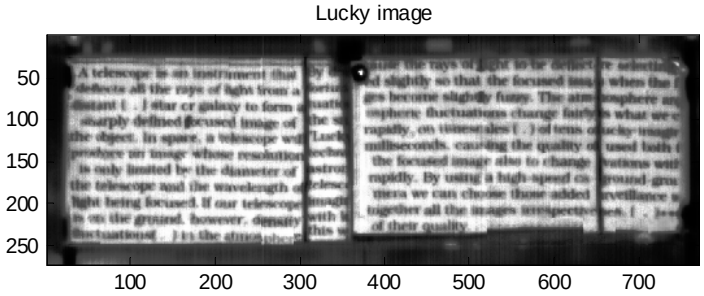

(Fig 2b, 10% selection)

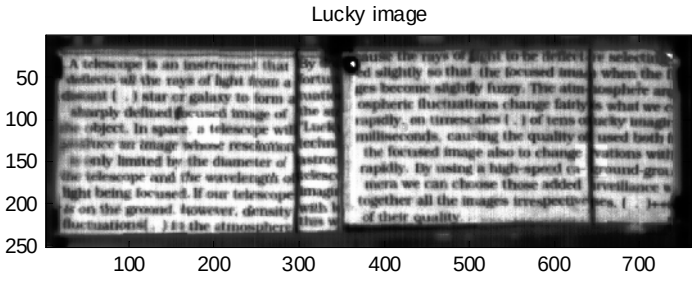

(Fig 2c, 1% selection)

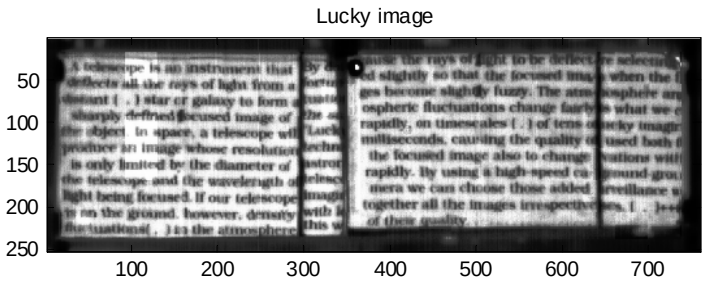

(Fig 2d, 0.1% selection)

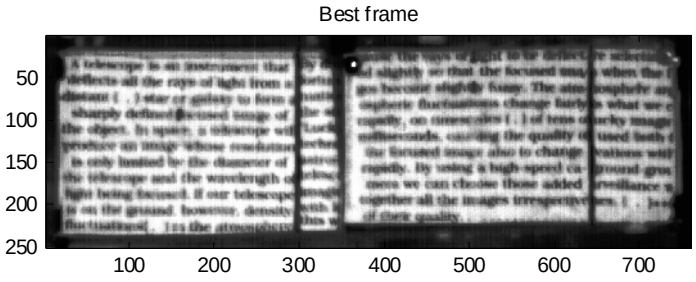

(Fig 2e, Sharpest frame)

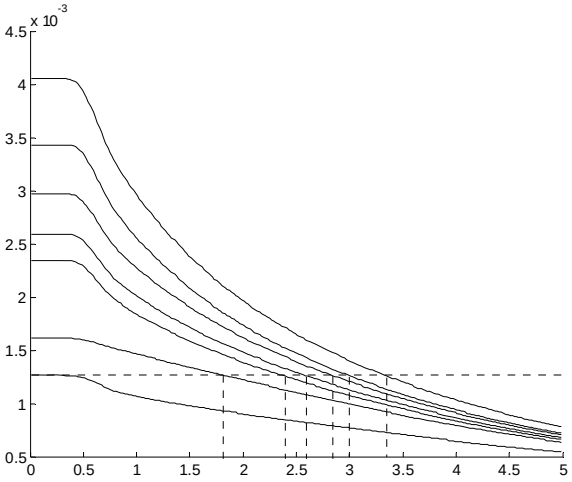

Fig. 3 Fisher information (vertical axis) vs. FWHM of the Gaussian function use to convolve each lucky image..

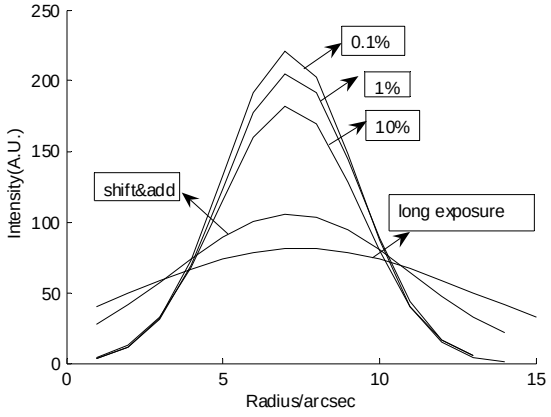

Fig. 4 LED cross-sections from each of the lucky images.