Efficient Photometry In-Frame Calibration (EPIC) Gaussian Corrections for Automated Background Normalization of Rate-Tracked Satellite Imagery

Jacob D. Griesbach, Joseph D. Gerber
Applied Defense Solutions, Columbia, Maryland

Paul F. Sydney, Charles J. Wetterer
Integrity Applications Incorporated / Pacific Defense Solutions

ABSTRACT

Dark and flat calibration frames are routinely required for accurate photometric processing of electro-optical astronomical and resident space object non-resolved imagery. However, within a new photometric processing tool, called Efficient Photometry In-Frame Calibration (EPIC), an automated background normalization technique is proposed that eliminates the requirement to capture dark and flat calibration images, while still retaining accurate photometric measurements. The data is then autonomously corrected for constant false-alarm rate (CFAR) detection using typical Gaussian normalization techniques. This technique consists of bias and variance corrections that account for dark noise, shot noise, CCD quantum efficiency, and optical path vignetting effects. The proposed technique is explored herein, with comparisons using real sample data.

1. INTRODUCTION

Photometric processing of non-resolved Electro-Optical (EO) images has commonly required the use of dark and flat calibration frames that are obtained to correct for charge coupled device (CCD) dark (thermal) noise and CCD quantum efficiency/optical path vignetting effects respectively [1]. It is necessary to account/calibrate for these effects so that the brightness of objects of interest (e.g. stars or resident space objects (RSOs)) may be measured in a consistent manner across the CCD field of view. Detected objects typically require further calibration using aperture photometry [2, 3] to compensate for sky background (shot noise). For this, annuluses are measured around each detected object whose contained pixels are used to estimate an average background level that is subtracted from the detected pixel measurements.

In a new photometric calibration software tool developed for AFRL/RD, called Efficient Photometry In-Frame Calibration (EPIC) [4], an automated background normalization technique is proposed that eliminates the requirement to capture dark and flat calibration images. The proposed technique simultaneously corrects for dark noise, shot noise, and CCD quantum efficiency/optical path vignetting effects. With this, a constant detection threshold may be applied for constant false alarm rate (CFAR) object detection without the need for aperture photometry corrections. The detected pixels may be simply summed (without further correction) for an accurate instrumental magnitude estimate.

The noise distribution associated with each pixel is assumed to be sampled from a Poisson distribution that models photons independently arriving at a constant rate over a fixed time interval [5]. Since Poisson distributed data closely resembles Gaussian data for parametrized means greater than 10, the data may be corrected by applying Gaussian bias subtraction and standard-deviation division.

First, a mathematical basis for the proposed corrections is given in Section 2. Then the proposed bias and variance correction techniques are explored in Sections 3 and 4. Identification of hot-pixels and cosmic-rays are explored and mitigated in Section 5. Proposed corrections for time-varying effects such as atmospheric turbulence, clouds, stray light, etc. are explored in Section 6. Lastly, comparisons are made using real data obtained from an Andor Neo camera in Section 7.

2. MATHEMATICAL BASIS FOR CORRECTION

Herein, shot noise is used to refer synonymously to sky background noise. Background noise is used to refer to the combination of dark current noise and shot noise.

Accurate background and noise calibration for EPIC was mainly driven by the desire to incorporate matched filtering for streaked star detection. In order to apply a matched filter to the data, the noise must be uniformly and identically distributed spatially across the image to provide constant false-alarm rate (CFAR) detection characteristics. As a result, the stationary dark current and sky background noise must be accounted for and normalized. It should be noted that even meticulously collected dark and flat imagery does not provide a complete correction here, since traditional dark and flat calibration imagery are collected without the presence of sky background noise. As a result, even if darks and flats are collected, provided, and applied, background normalization is still necessary to account for sky background noise before matched filtering. Without complete normalization of the background noise across each image, the matched filter detection may erroneously detect even slight contours of the background noise as stars.

A primary assumption here is that shot noise uniformly illuminates the telescope aperture. However, the telescope optics vignette the shot noise so that it (as well as the signals of interest) vary in intensity across the CCD sensor. As a result, by normalizing the background noise across each image, we also properly restore the signal intensity as well.

Mathematically, we can model the signal and noise as

$$C_{i,j} = \frac{\left[\left(s_{i,j} + n\right)\left(VQ\right)_{i,j} + d_{i,j}\right]T + b_{i,j}}{g}. \tag{1}$$

Here,

$C_{i,j}$	$=$	measured counts for the i,j^{th} image pixel (counts)
$s_{i,j}$	$=$	signal of interest for the i,j^{th} image pixel (photons/sec)
n	$=$	shot noise (photons/sec)
$(VQ)_{i,j}$	$=$	combined optical vignetting factor (unitless) and quantum efficiency (electrons/photon) for the i,j^{th} image pixel (electrons/photon)
$d_{i,j}$	$=$	dark current for the i,j^{th} image pixel (electrons/sec)
T	$=$	exposure time (sec)
$b_{i,j}$	$=$	A/D converter bias (electrons)
g	$=$	gain (electrons/count)

We assume that all of the above values are deterministically constant for the period of the image deck with the exception of the shot noise, which we model as

$$n \in \mathcal{N}\left(\bar{n}, \sigma_n^2\right), \tag{2}$$

where $\mathcal{N}$ denotes the normal distribution and $\bar{n}$ and σ_n^2 are the mean and variance of the shot noise, respectively. While the shot noise might be more appropriately modeled as Poisson distributed, we model it here as a non-centralized Gaussian (Normal) distribution, as Poisson distributions are well matched to Gaussian distributions for values of $\bar{n} > 10$ [5]. The primary goal of the background and noise calibration is to remove the biasing effects of $d_{i,j}$ and $b_{i,j}$, while removing the spatially varying nature (the i,j dependency) associated with the vignetting factor and quantum efficiency, $(VQ)_{i,j}$, from (1). It turns out that the other factors involved, namely T and g, are simply absorbed into the zero point estimation.

In EPIC, background and noise calibration consists of the following four steps:

1. Background noise bias subtraction
2. Background noise variance normalization
3. Removal of hot pixels
4. Time-Varying noise bias subtraction

Each of the above steps is addressed in the next sections. All of the above processing steps typically takes 1-2 minutes for a deck of 100 images.

3. BACKGROUND BIAS CORRECTION: DARK NOISE + SKY NOISE

EPIC performs automated background normalization on rate-tracked satellite images using the following technique. A deck of approximately 50-100 images is combined by performing an independent median calculation along the deck dimension for each image pixel. Because the images are rate-tracked, moving objects (such as background stars) are quickly eliminated. Stationary RSO signatures are removed from the resultant median combined image using a local area median filter that smooths over the RSO responses. The result is a smoothed estimate of the background noise bias. The calculated bias estimate image is subtracted from each deck image, which effectively removes noise bias.

More formally, the above steps for bias correction can be enumerated in the following way:

1. All of the images within an image deck are median combined. This generally removes the effects of streaking stars, however, rate-tracked objects can still persist.
2. Hot pixels are detected by recording the pixel locations with values that are 10x greater than the mean of all 8 of its neighboring pixels using a fast convolutional filter technique.
3. A local 51x51 pixel median filter is applied to the median combined image to remove the effects of rate-tracked RSOs. It is important that the size of the median filter be 2-3x as large as the largest expected RSO. Smaller median filters will not effectively remove RSOs, and larger ones will degrade the ability of the estimation to detect vignetting fluctuations.
4. Detected hot pixels from step 2 are restored into the resultant image with their original values.
5. The resulting image is subtracted from each image in the original image deck to correct for background noise bias.

Note that the above algorithm strives to preserve hot pixel values as they are detected throughout the image deck. By subtracting their median value, this effectively reduces those pixels to a zero corrected count level. By doing so, such pixels will not cause matched filter detection false alarms during star detection.

Mathematically, the estimated median background bias image may be represented as

$$m_{i,j} = \frac{\left[\bar{n}(VQ)_{i,j} + d_{i,j}\right] T + b_{i,j}}{g} \tag{3}$$

Subtracting (3) from (1) results in the bias corrected image described as

$$C_{i,j} - m_{i,j} = \frac{T(VQ)_{i,j}}{g} s_{i,j} + \frac{T(VQ)_{i,j}}{g} \left(n - \bar{n}\right). \tag{4}$$

In (4), one can directly see that, while the undesirable d and $b_{i,j}$ terms have been canceled, the spatially varying vignetting factor and quantum efficiency, $(VQ)_{i,j}$, is still affecting the signal (left term) and the shot noise (right term). Because the vignetting factor is included in (3), we expect to visually see the vignetting captured in the estimated background bias image. Indeed, the vignetting is a prominent factor in the estimated background bias image, as seen in Figure 2 for the example raw image shown in Figure 1. Figure 3 shows the resulting bias corrected image after subtracting the estimated background bias image from the raw image.

If dark calibration images are provided [2], EPIC can apply them, instead of estimating the background bias, if the user configures EPIC to do so. By default, EPIC is configured not to utilize provided darks (and flats). If multiple dark images are provided, these images are median combined to form a super-dark image. If a single image is provided, the image is assumed to already be super-dark. This dark image will be subtracted from all provided imagery data to attempt to remove the effects of sensor dark current.

It is important to note, that while the estimated background bias is handled in the same fashion as if it were a provided super-dark image, the estimated background bias includes the effects of shot noise, while a dark image does not. As a

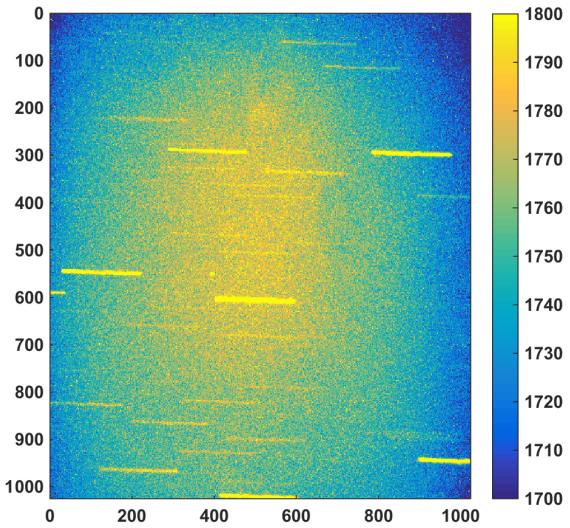

Fig. 1: Example Raw Image

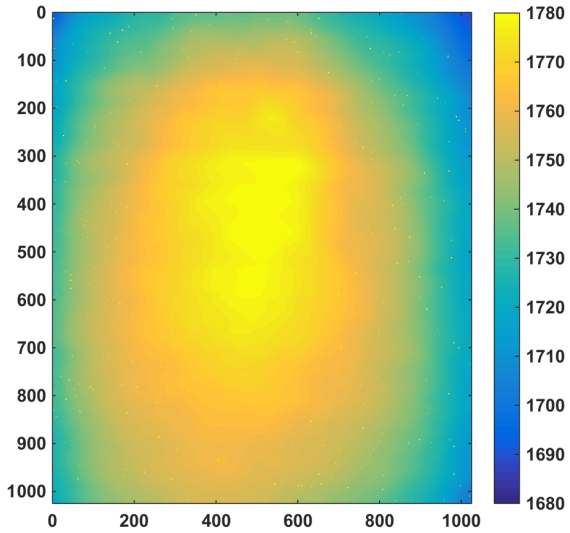

Fig. 2: Estimated Background Bias

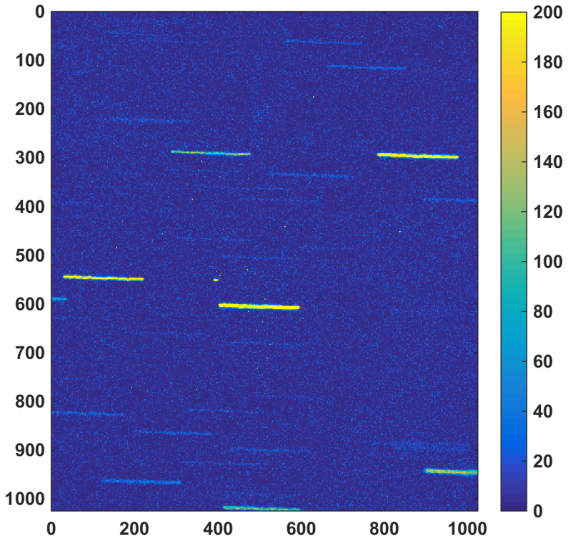

Fig. 3: Bias Corrected Image

result, the background bias image does not attempt to estimate the dark image. Because of this, even if a user supplies both dark and flat calibration images, a correction for the remaining shot noise must still be applied in lieu of aperture photometry before matched filtering can occur. The time-varying correction (step 4 as enumerated in Section 2) can fulfill this requirement.

4. BACKGROUND VARIANCE CORRECTION: FLAT ESTIMATION

A variance correction is applied to address spatially varying noise by dividing each pixel by the square-root of its measured variance. The proposed technique uses the sky background noise (assumed to be uniform over the telescope aperture) to normalize the CCD quantum efficiency/optical path vignetting effects.

After background bias subtraction is completed, EPIC estimates the background noise variance normalization using the following steps, similar to the bias estimation process. This is referred to as the "deck-based" variance estimation technique.

1. For each image pixel, the variance is computed using the data for the selected image pixel from all of the images

in the image deck. An iterative approach is used to eliminate 3-sigma outliers until the measured variance converges.

2. A local 51x51 pixel median filter is applied to the image containing the measured variances for each pixel. As for the background bias estimation process, this is done to remove the effects of rate-tracked RSOs from the results.

3. A square root is applied to convert all of the variance values to standard deviation estimates in the smoothed image.

4. The resulting smoothed image values are normalized by dividing by the largest smoothed variance measurement. Afterwards, the largest variance measurement has a value of one (typically in the image center), while the other pixels have values less than one.

5. The normalized smoothed image is divided from each image in the original image deck (after dark correction) to normalize the noise power across the image.

Alternatively, it may be noted that if indeed the pixel noise is Poisson distributed, that for such data, the variance is equal to the mean. With this approach, the following "Poisson-based" algorithm may be used:

1. Start with the bias estimated image from Step 3 of the algorithm in Section 3. This is the smoothed image variance estimate.

2. A square root is applied to convert all of the variance values to standard deviation estimates in the smoothed image.

3. The resulting smoothed image values are normalized by dividing by the largest smoothed variance measurement. Afterwards, the largest variance measurement has a value of one (typically in the image center), while the other pixels have values less than one.

4. The normalized smoothed image is divided from each image in the original image deck (after dark correction) to normalize the noise power across the image.

For a comparison of the algorithms, see Section 7.

It should be noted that the variance normalization purposely does not attempt to detect or correct for hot pixels. Doing so would potentially cause such pixels to be multiplied by a very large value in attempt to widen their near zero variances to the other (healthy) pixels. Instead, such pixels are purposely left with small resultant variances.

Mathematically, proceeding from (4) and noting that the process of computing the variance with smoothing and outlier rejection over the image deck strives to remove the signal response from the variance estimation process, we examine the variance associated with the noise term. Using (2), we realize that

$$n - \bar{n} \in \mathcal{N}\left(0, \sigma_n^2\right), \tag{5}$$

and

$$\frac{T(VQ)_{i,j}}{g}\left(n - \bar{n}\right) \in \mathcal{N}\left(0, \left(\frac{T(VQ)_{i,j}}{g}\right)^2 \sigma_n^2\right). \tag{6}$$

As a result, variance estimation for the i, j^{th} pixel should estimate the variance to be

$$\hat{v} = \left(\frac{T(VQ)_{i,j}}{g}\right)^2 \sigma_n^2. \tag{7}$$

Steps 4 and 3 the respective processes above divides all of the pixel variance estimates by the largest, resulting in a correction factor as

$$r_{i,j} = \sqrt{\frac{\hat{v}}{\hat{v}_{\max}}} = \frac{\frac{T(VQ)_{i,j}}{g}\sigma_n}{\frac{T(VQ)_{\max}}{g}\sigma_n} = \frac{(VQ)_{i,j}}{(VQ)_{\max}}. \tag{8}$$

It could be noted that one could apply $r_{i,j} = \sqrt{\hat{v}}$ instead, but this would normalize all of the noise variances in the corrected image to a unit variance of 1, which may severely decrease the overall amplitude of the image. To avoid this, we use (8) since it has the same effect as flat correction, which is to generally preserve the center of the image while amplifying the borders to reverse the effects of vignetting.

Applying (8) to (4) by dividing by $r_{i,j}$ results in a corrected image value of

$$\frac{C_{i,j} - m_{i,j}}{r_{i,j}} = \frac{T(VQ)_{\max}}{g}\left(s_{i,j} + (n - \bar{n})\right). \tag{9}$$

With this, we have achieved the goal of removing all bias terms and i, j specific spatially varying terms from the right hand side. This results in zero mean images with noise variances that are consistent across the image. The remaining deterministic constants, T, $(VQ)_{\max}$, and g, may be estimated individually or accounted for by sensor properties. Regardless, they will be absorbed into the zero point calculation and pose no inaccuracies if left unaccounted for.

Figure 4 illustrates a plot of the variance correction factor, $r_{i,j}$, for the example image shown in Figure 1. Figure 5 shows the resulting bias and variance corrected image. It is hard to notice a difference between Figures 5 and 3, but one may notice that the bright star in the lower right corner has been corrected (amplified) to account for the image vignetting reducing its brightness.

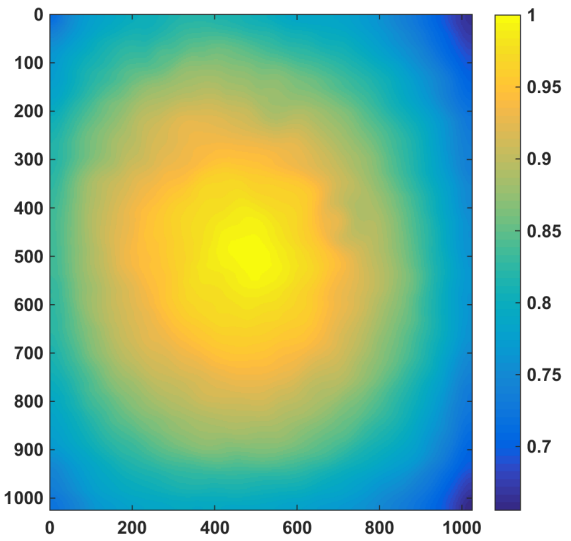

Fig. 4: Estimated Background Variance (Flat)

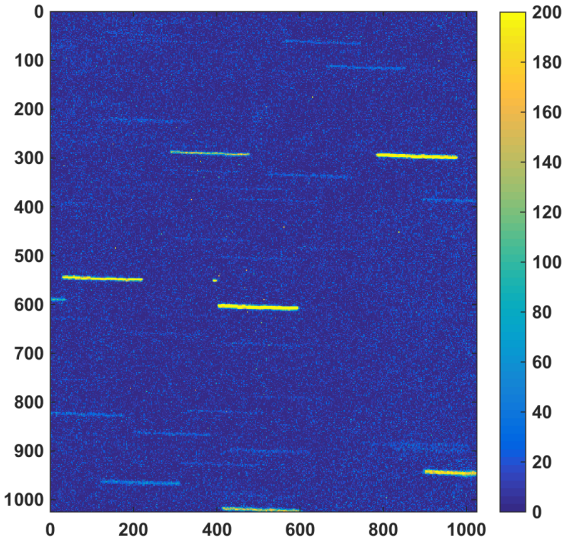

Fig. 5: Bias and Variance Corrected Image

If flat images are provided, EPIC can use them, instead of estimating the background variance, if the user configures EPIC to do so. By default, EPIC is configured not to utilize provided flats (and darks). If multiple flat images are provided [2], these images are median combined to form a super-flat image. If a single image is provided, this image is assumed to already be super-flat. This flat image will be divided from all provided imagery data to attempt to remove the effects of vignetting.

It is important to note, that while the estimated background bias is considered to be different from a dark image (because the estimated background bias includes shot noise), the estimated variance background correction image (as shown in Figure 4) directly correlates to a flat. Therefore, it is fair to state that the estimated background variance

image is an estimated flat. While a true flat is measured using a highly controlled uniform illumination, the estimated background variance strives to compute the same result using shot noise.

5. AUTOMATIC HOT-PIXEL AND COSMIC-RAY MITIGATION

There are commonly a significant number of hot pixels and/or cosmic rays present that produce very strong single pixel measurements. While the bias estimation process (or provided darks) should account for (stuck) hot pixels, it will not correct for cosmic ray events. A simple filtering algorithm was written to detect and remove such strong single pixels, whose values lie 10x beyond those of its neighboring 8 pixels. The algorithm uses an 8-pixel convolutional filter kernel that allows efficient Fast Fourier Transform (FFT) -based processing.

The algorithm is as follows:

1. Filter a given image with a 2-D filter consisting of a 3x3 matrix where all the values are equal to 1/8 except for the middle value, which is set to 0. The convolution of this kernel with an image replaces each image pixel with the average value of all the surrounding pixels.

2. One then compares the original image to the filtered image: if the original image has any pixels whose value is larger than 3-sigma and larger than the filtered image pixel values by 10x, then the original image pixel is replaced by the filtered image pixel.

It should be noted that this approach only works for single pixels. If a cosmic ray affects two neighboring pixels, this approach will leave those pixels unchanged. Figure 6 shows the remaining strong single pixels associated with bias and variance corrected data from Figure 5 (plotted with a different color scale). Figure 7 shows the resulting image after cosmic ray reduction has been applied. One can easily see that the result is dramatically cleaner afterwards.

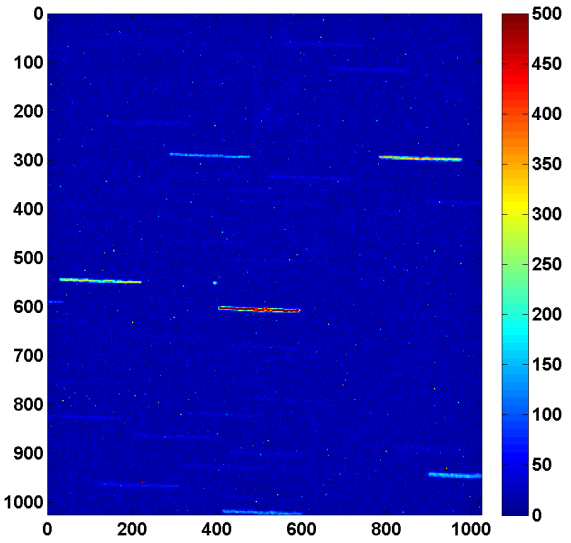

Fig. 6: Before Mitigation

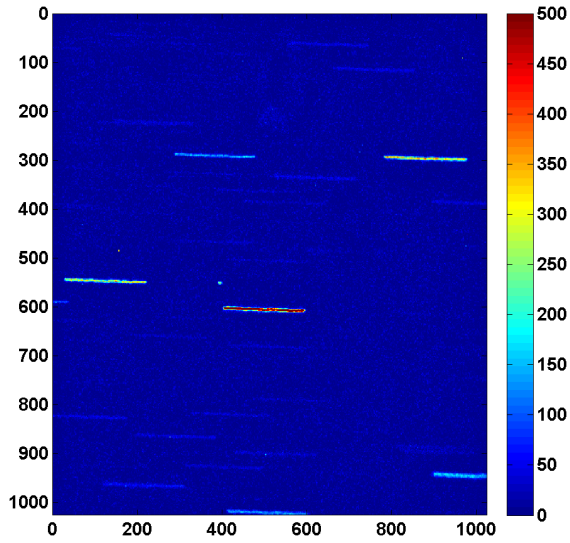

Fig. 7: After Mitigation

6. TIME-VARYING BACKGROUND CORRECTION

The aforementioned bias and variance estimation techniques provide time-invariant image correction assuming the background noise is stationary for all images in a deck. I.e., a constant image correction that is applied equally to all images in an image deck. Even after performing this correction, the median associated with each individual image in the image deck can still be observed to vary around zero. In this final step, a time-varying image background estimation is applied to account for changing atmospheric conditions (clouds) and/or time-varying image noise. This effectively corrects the median to zero in an image-by-image basis. Hence, it is a time-varying correction.

This step is imperative if provided dark and flat images are used. As described earlier, provided dark images do not account for shot noise, as so, the shot noise will remain in the images after provided dark/flat correction is applied. If left uncorrected, this will cripple the matched filter detection. This step provides the needed correction.

Once performed, this step removes the need to perform aperture photometry (using an annulus to estimate background noise around detections) [1]. Since this step strives to normalize the background bias such that noise is zero mean, if one did perform aperture photometry after applying this correction, the background measurement should effectively be zero.

If estimated bias and variance correction is used (instead of provided dark and flat images), then this step usually results in only minor corrections.

1. For a given image pixel, a 128x128 local pixel neighborhood is established, where the neighborhood is used to estimate the median (noise bias) of the local area. Since such a large local median computation can be expensive, a median-of-medians approach is used instead, where medians are first computed for each column in the 128x128 local area, and then the final median is computed as the median of the column median results.

2. Step 1 is repeated for every pixel in a given image, where each pixel location is replaced with its 128x128 local median result.

3. The median image estimate is subtracted from the image to produce the time-varying corrected image.

4. The above steps are repeated for each individual image in the image deck.

7. REAL-DATA COMPARISONS

At first, data obtained from a local telescope hosting an Andor 897 camera was going to be used for this comparison, but it had a very flat field-of-view with less than 3% vignetting. With this data, all of the estimation techniques yielded very close performance to the provided dark and flat, essentially showing only negligible differences. As a result, a different telescope/camera assembly was found that illustrated much more dramatic vignetting, which started to illustrate differences between the proposed techniques. The data that follows was obtained from an 11" telescope hosting an Andor Neo camera.

This vignetted data poses a challenge for the EPIC processing, primarily it seems, for the flat estimation. With the very flat Andor 897 telescope assembly, EPIC routinely achieves relative accuracies of 2.5-5%, on a good, photometric night. However, while this particular collection did have some light atmospheric haze, EPIC only achieves relative accuracies slightly better than 10%, whether dark/flat calibration imagery is used or not.

71 Dark Comparison

It's difficult to compare the difference between using a provided dark and the proposed estimated background bias technique, because they measure different things. A provided dark only seeks to capture the sensor dark current, while the estimated background bias strives to estimate the dark current + sky noise. As seen in Figures 8 and 9, perhaps the only conclusion that can be reached is that the provided dark might indeed serve as a base for the estimated dark + sky image. It should be noted that both techniques calculate the same zero point to within 2 hundredths of a Vm.

If the dark subtraction is incorrect, this can introduce a logarithmic curvature into the zero point data. While difficult to assess, there is no indication of distinctive curvature seen when examining Figures 10 and 11.

72 Flat Comparison

The effects of uncorrected vignetting is easier to assess, since stars with (near) constant brightness traverse the images in the deck horizontally. With a frame rate of 3 seconds rate-tracking a GEO object, the same star can be imaged and measured many times. By comparing the relative deviations in brightness measurements one can begin to discern if residual vignetting exists.

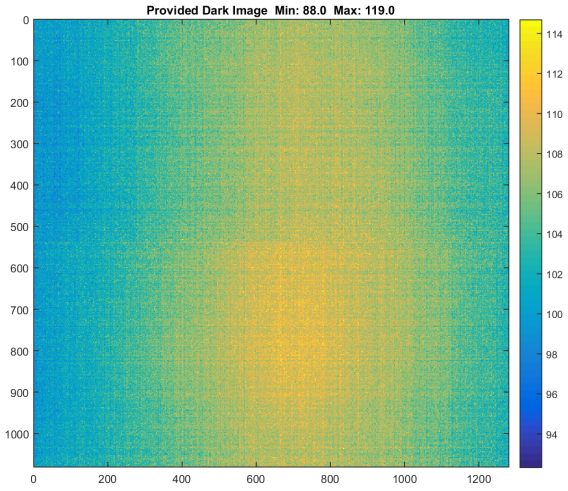

Fig. 8: Provided Dark

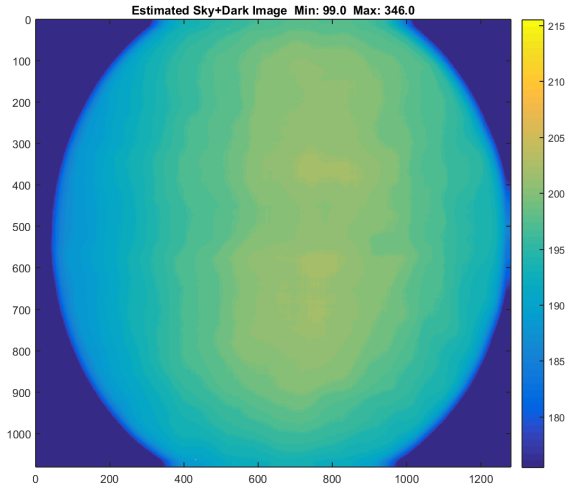

Fig. 9: Estimated Background Bias (Dark + Sky)

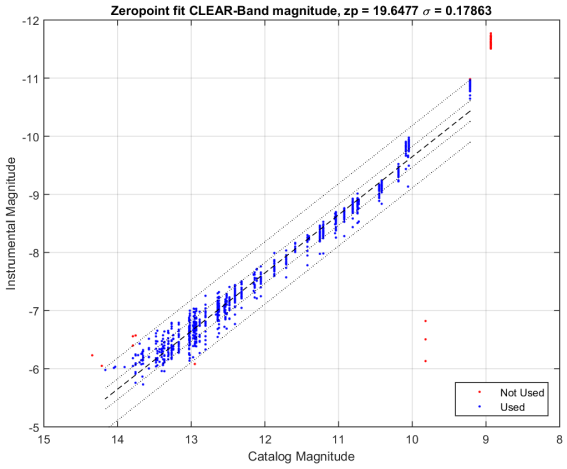

Fig. 10: Provided Dark/Flat Zero Point

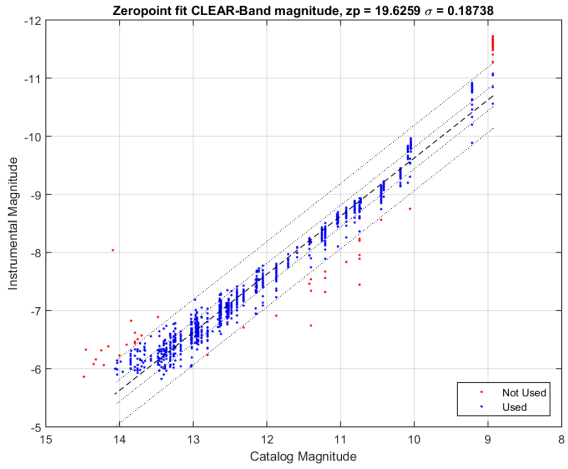

Fig. 11: Estimated Bias/Flat Zero Point

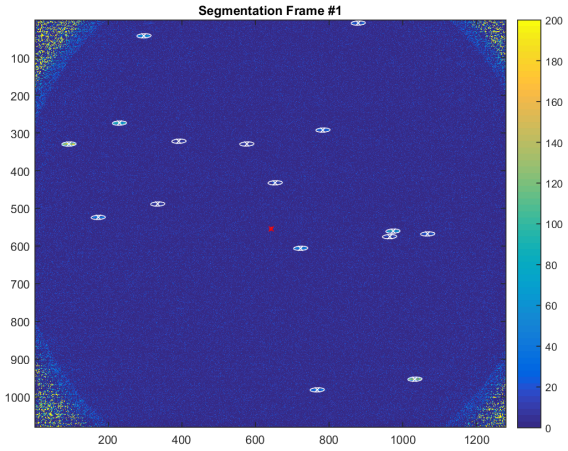

Fig. 12: Provided Dark/Flat Example Segmented Image

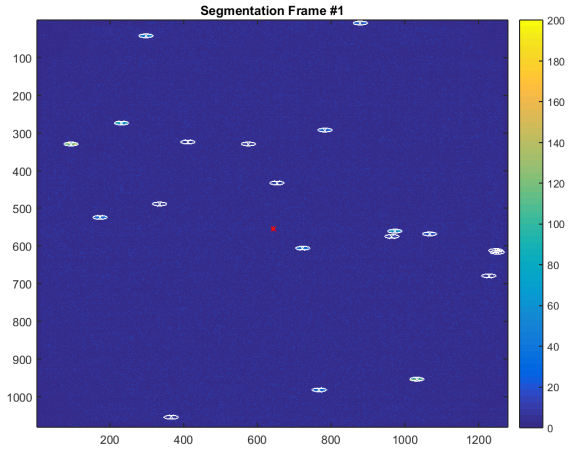

Fig. 13: Estimated Bias/Flat Example Segmented Image

Sometimes, one can even visually see if vignetting effects remain in the corrected imagery. Indeed, the provided flat seems to over-correct for the vignetting, as the extreme regions of the resultant imagery have elevated noise levels. This can be readily seen in Figure 12 using the provided flat and is not apparent in Figure 13, which uses an estimated flat. This strongly suggests that the provided flat may be corrupt.

Figure 14 shows results obtained from relative star flux measurements. As individual measurements are obtained horizontally across the image a curve may be fit to the data to discern a trend. This is done by fitting polynomials of order 2, 3, and 4 to the data, which are shown in each of the plots. General agreement of the various polynomial orders give confidence to the curve fitting.

The visual magnitude deviation input to the curve fitting is calculated in the following way:

$$\text{residual Vm}_s = -2.5 \log_{10} \left(\frac{f_{s,i}}{\text{median}\,(f_{s,i})} \right) \tag{10}$$

where $f_{s,i}$ denotes the measured flux for a given star s at horizontal position i in the imagery. The median is performed over all horizontal positions, i for a given star. In this way, a residual Vm measurement is produced for each individual star flux observation.

In Figure 14a, the average relative star deviation curve fits are shown for 5 different configurations, all using exactly the same input data imagery. Figure 14a illustrates the measured residual vignetting using the provided dark and flat calibration imagery. The remaining subfigures all use the estimated (dark+sky) bias correction to provide a constant comparison baseline. Figure 14b illustrates the combination of the provided flat with the estimated (dark+sky) bias correction, while the remaining three illustrate different techniques for flat estimation. The Poisson Flat and Deck Flat techniques are described in Section 4. Figure 14e illustrates results obtained from a "Spatial-Flat" technique where the variance is estimated using local spatial estimation in each individual image before median combining results across the deck. This technique is useful for small decks (10-25 images) and can perform more reliably than the "Deck-based" technique for such decks.

According to the curve fitting plots, all of the techniques over-correct for the vignetting somewhat. However, the Poisson-based flat estimation appears to perform the best of the 5 techniques, even achieving better performance than the provided dark and flat.

In a final accuracy comparison, the mean relative accuracies are shown numerically in Table 1 along with the zero point fit standard deviations for each technique. Here, lower relative accuracies denote better performance, where 0% would denote the algorithms ability to measure individual stars with exactly the same flux repeatedly. From a mean relative accuracy statistic, the "Deck-based" technique appears to perform the best, while all of the estimated techniques achieve the same or better accuracies than using the provided dark and flat data, at least for this set of imagery data.

The zero point fit standard deviations are also provided in Table 1 for documentation purposes, but these call into question the accuracy of the correlated star catalog, which can provide an additional source of error. These numbers (as well as the zero points illustrated in Figures 10 and 11) were generated using the SST-RC5 photometric star catalog.

	Relative Accuracy (%)	Vm Standard Deviation (σ)
Provided Dark & Provided Flat	9.40	0.179
Estimated Sky+Dark & Provided Flat	9.25	0.185
Estimated Sky+Dark & Poisson Flat	9.25	0.187
Estimated Sky+Dark & Deck Flat	9.03	0.185
Estimated Sky+Dark & Spatial Flat	9.29	0.181

Table 1: Accuracy Comparison

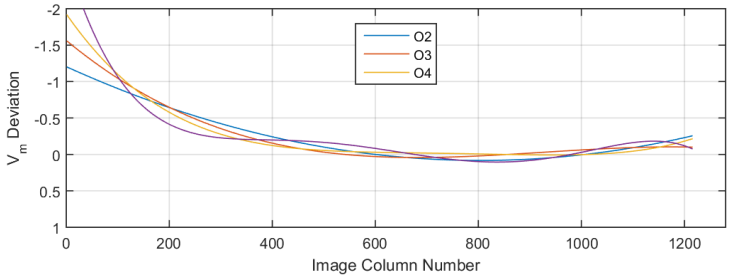

(a) Provided Dark & Provided Flat

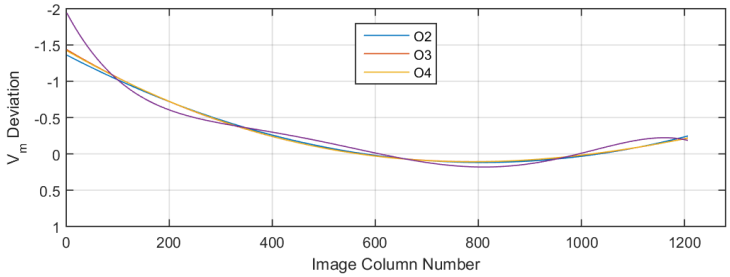

(b) Estimated Sky+Dark & Provided Flat

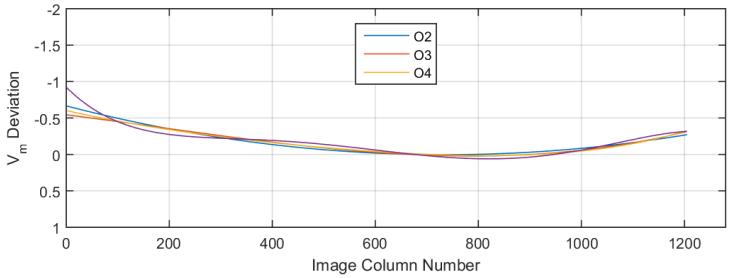

(c) Estimated Sky+Dark & Poisson Flat

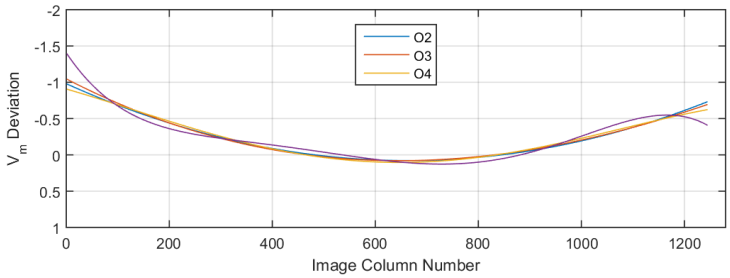

(d) Estimated Sky+Dark & Deck Flat

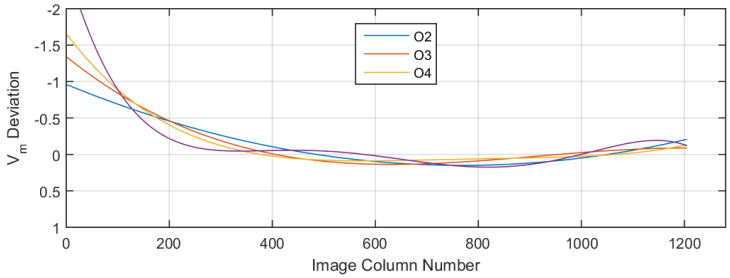

(e) Estimated Sky+Dark & Spatial Flat

Fig. 14: Average Horizontal Brightness Deviations for Tracked Stars

8. CONCLUSIONS

The Efficient Photometry In-Frame Calibration (EPIC) software can apply automated corrections, if a user does not wish to provide or trust the dark/flat calibration imagery. The software's primary technique, the "Poisson-based" technique, is also the fastest to execute and is shown to perform well in the comparisons. In many cases, the supplied calibration images were found to perform worse than the proposed estimation techniques, most likely because the calibration images were somehow corrupted or mismatched.

It is the authors' view that good, meticulously collected calibration imagery can outperform the estimated techniques. However, if such calibration imagery is not available or is too difficult to collect, the proposed techniques should provide a valuable alternative.

References

[1] R. Berry and J. Burnell, *The Handbook of Astronomical Image Processing*. Willmann-Bell, 2011.

[2] S. B. Howell, *Handbook of CCD Astronomy*. Cambridge University Press, 2006.

[3] B. D. Warner, *A Practical Guide To Lightcurve Photometry and Analysis*. Springer, 2006.

[4] J. D. Griesbach and J. D. Gerber, "Efficient Photometry In-Frame Calibration," *Small Telescope Workshop*, Chantilly, VA, 2015.

[5] Wikipedia, "Poisson Distribution," 2015. https://en.wikipedia.org/wiki/Poisson_distribution [Online; accessed 23-July-2015].