An algorithm-independent analysis of the quality of images produced using multi-frame blind deconvolution algorithms

Charles L. Matson and Alim Haji

Directed Energy Directorate, Air Force Research Laboratory, Kirtland AFB, NM 87117

ABSTRACT

Multi-frame blind deconvolution (MFBD) algorithms can be used to generate a deblurred image of an object from a sequence of short-exposure and atmospherically-blurred images of the object by jointly estimating the common object and all the blurring functions. In this paper we present fundamental limits (Cramér-Rao lower bounds) to the quality of restored images generated with MFBD algorithms. We show that image restorations are less noisy when using a Zernike-based point spread function (PSF) parameterization than when using a pixel-based PSF parameterization. We also show that the most noise reduction tends to occur near the edges of the true object support, even when a larger support region is used as a constraint. Finally, we show that Zernike-based PSF parameterization produces higher resolution in restored images than does pixel-based PSF parameterization.

1. INTRODUCTION

Atmospheric turbulence greatly limits the spatial resolution in images of space objects collected with ground-based telescopes. For example, when the atmospheric coherence parameter $r_o = 10$ cm, the resolution in a long-exposure image is r_o/λ regardless of the size of the telescope for telescopes whose diameters D are larger than ~10 cm, where λ is the wavelength at which the image is obtained. Adaptive optics can be used, under the right conditions, to remove much of the turbulence effects, producing images whose resolutions approach the diffraction limit of the imaging system. Often, though, images are collected with telescopes that do not have adaptive optics systems. Furthermore, even adaptive optics systems do not routinely produce diffraction-limited images. For these reasons, carrying out post-processing on these kinds of images produces increases in resolution. There are a variety of algorithms that can be used to post-process images. We choose to use a multi-frame blind deconvolution (MFBD) algorithm for at least three reasons. The first is that our algorithmic implementation of MFBD has been shown to achieve or closely approach the theoretical limits to image quality for the cases that we have investigated [1]. The second reason is that no separate star measurements are needed, unlike for speckle imaging methods [2]. This second fact is especially important for telescopes with adaptive optics systems because separate star measurements are needed for each set of system parameters, such as the adaptive optics system closed-loop bandwidth. The third reason is that MFBD algorithms can handle non-idealities such as image clipping and anisoplanatism more robustly than can speckle imaging algorithms.

We have been investigating how the quality of image restorations produced using MFBD algorithms depends on the observation and image processing parameters to determine how to optimize these parameters. In this paper we present results for support-constrained image restorations as a function of the model of the point spread function (PSF), the types of measurement noise, and the number of measurements. We quantify image quality in terms of the variances of estimates of image-domain pixel values and of spatial frequency values in MFBD image restorations. The former metric quantifies the accuracy of estimating image intensities while the latter metric quantifies the amount of resolution in the restored images.

To produce the results presented herein, we used Cramér-Rao lower bound (CRB) theory. CRBs are algorithm-independent lower bounds to the variances of any estimate of a vector of parameters, given a probability density function (PDF) of a set of measurements that is a function of the parameters that are being estimated. The CRBs are a function of the observation and image-processing parameters underlying the measurement data and restored images, however. This fact has motivated us to explore this functional dependence of CRBs to help us understand how to maximize image restoration quality. The outline of this paper is as follows: the imaging model that is used to produce the PDF of the measurements that is a function of the parameters to be estimated is given in Section 2, CRB theory is described in Section 3, the results of the CRB analysis are found in Section 4, and conclusions and future work are presented in Section 5.

2. IMAGING MODEL

We model a noise-free measured image as a space-invariant convolution of the object being imaged with the PSF of the telescope/atmosphere combination. This noise-free measurement is then corrupted by additive zero-mean noise that can be signal-independent noise such as CCD read noise, signal-dependent photon noise, or the sum of the two. We model the additive noise PDF as a Gaussian PDF where each pixel has a variance equal to the sum of the signal-independent noise variance and the intensity of the image at that pixel. We use the central limit theorem to justify modeling the photon noise as Gaussian. The joint PDF for the set of measurements is the multiplication of the individual additive noise PDFs for each measurement pixel, each with a mean equal to the noise-free measured image intensity in that pixel.

The system PSF is determined by three quantities: the optical system pupil function, the optical distortions in the system pupil, and the measurement wavelength. We assume that the measurement is sufficiently narrow band that a single-wavelength Fourier-optics model accurately describes the system PSF [2]. In terms of modeling the measured images, only an image-domain description of the system PSF is necessary; however, the model of how the system PSF is generated from the PSF parameters plays an important role in the quality of the restored images.

3. CRAMÉR-RAO LOWER BOUND THEORY

Given a vector of parameters to be estimated, CRB theory produces lower bounds to the variances of any estimates of that vector's elements. These CRBs are not necessarily achievable; however, they often can be achieved or be approached closely, as we have discovered with our MFBD algorithm [1]. The CRBs for any unbiased estimate of a vector of parameters are the diagonal elements of the inverse of the Fisher information matrix (FIM) associated with the measurements [3]. The FIM is a function of the derivatives of the measurement PDF and can be calculated easily for the types of noises we are considering in this paper. When it is desired to generate CRBs for functions of a vector of parameters, the inverse of the FIM can be transformed in a straightforward way such that the diagonal elements of this transformed matrix are the desired CRBs.

For the results presented in this paper, the parameters to be estimated are the object parameters and the PSF parameters. The object parameters are always the pixelized intensities of the object because we do not make any assumptions about the structure of the object. Unbiased estimates of these parameters are impossible, in general, because imaging systems are low-pass filters. For this reason, the object CRBs presented in this paper are of either regularized estimates of the true object intensities or unbiased estimates of the true object energy spectrum values inside the bandpass of the optical system.

The PSF parameters to be estimated depend upon the forward model assumed for the PSF. If the only information about the optical system is the image-domain description of the PSF, the PSF parameters are the image-domain pixelized PSF intensities. If, however, it is known that a Fourier-optics model well describes the generation of the system PSF, then the PSF can be parameterized in terms of the field in the optical system pupil. Various parameterizations have been described in the literature. We have chosen to use a Zernike-polynomial-based parameterization because it produces excellent results with our MFBD algorithm. With this parameterization, the PSF parameters are the coefficients of the Zernike polynomials.

4. RESULTS

In this section we present object-estimation CRBs for a variety of scenarios. Key results are: (1) Zernike-based PSF parameterization produces image-domain object CRBs that are a factors of six or more lower than the image-domain object CRBs for pixel-based PSF parameterization, (2) the image-domain object CRBs are lowest near the edges of the true object support region, even when larger support regions are used for constraints, and (3) higher spatial resolution is possible with Zernike-based PSF parameterization than with pixel-based PSF parameterization. All of the CRBs are calculated using a support constraint on the object estimate. For pixel-based PSF parameterization, a centered circular support constraint is used for the PSF since the FIM is non-singular only for invertible PSFs. We minimize the impact of applying support constraints on infinite-support PSFs (always the case for unaliased imaging systems) by choosing the PSF support size large enough to contain greater than 99.9% of its energy. For Zernike-based PSF parameterization, no PSF support constraint is required.

Fig. 1. The (a) OCNR satellite model and (b) a PSF realization used as inputs to the CRB calculations.

The object used for these results is the OCNR satellite shown in Fig. 1(a). The PSFs were generated from random realizations of phase-only atmospheric turbulence distortions in the pupil of the telescope for $D/r_o = 8$. One of the PSFs is shown in Fig. 1(b). The read noise variance is 100 and the photon level is 10^8 when photon noise is included. All of the CRBs include the effects of an ideal low-pass filter as a regularizing filter whose spatial frequency cutoff is two pixels smaller than the optical system spatial cutoff frequency.
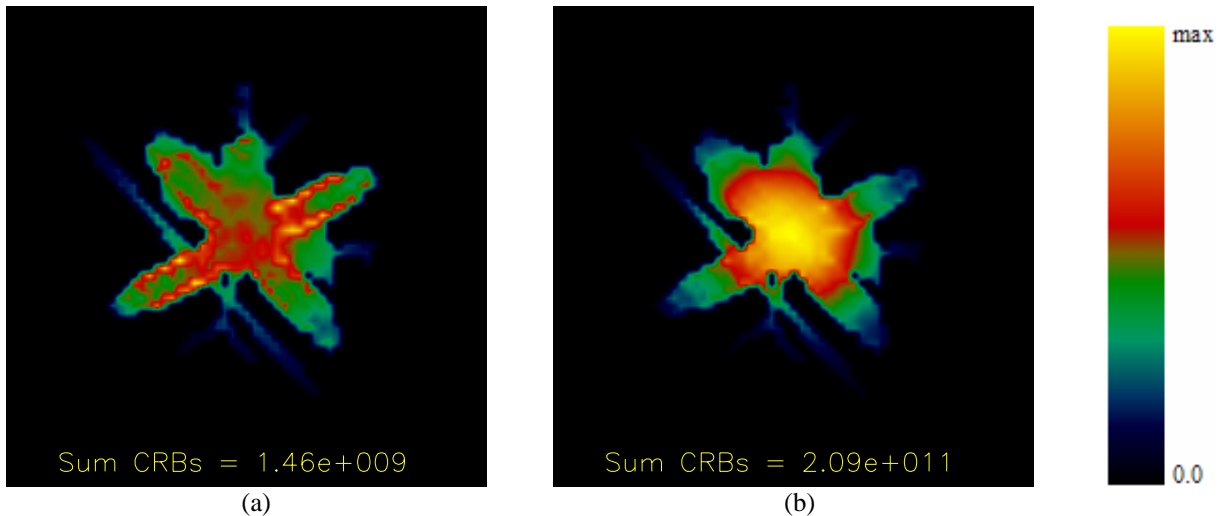

Fig. 2. Image-domain object CRBs for (a) read noise and (b) photon noise when the true object support is used as a constraint and the PSFs are estimated pixel by pixel.

The first set of CRBs, shown in Fig. 2, are image-domain object CRBs for a single measurement frame with read noise and with photon noise. The true object support and pixel-based PSF parameterization were modeled for these results. The displays in Fig. 2 are not restored images; rather, they are the CRBs within the support constraint region. We call these displays CRB images. The color bar for each of the CRB images in this paper is normalized to the maximum value in the CRB image; thus, only CRB morphologies can be compared from figure to figure. The sum of the CRBs for each spatially-resolved CRB image are shown at the bottom of each figure in these and the

following figures. In both images, it can be seen that, on average, the CRBs are lowest at the edges of the support region and highest in the interior. We believe that this property results from the fact that knowledge that non-zero values outside the object support are purely noise has the most impact near the edges of the support region. The photon-noise CRBs are more uniformly higher in the interior of the object support than are the read-noise CRBs. This occurs because the photon noise levels are higher in the center of the object support since the object is brightest in the center of the object support, while the read noise level is constant throughout the object support. At this time, we do not understand why the read-noise CRBs have the spatial structure that they show in Fig. 2(a).
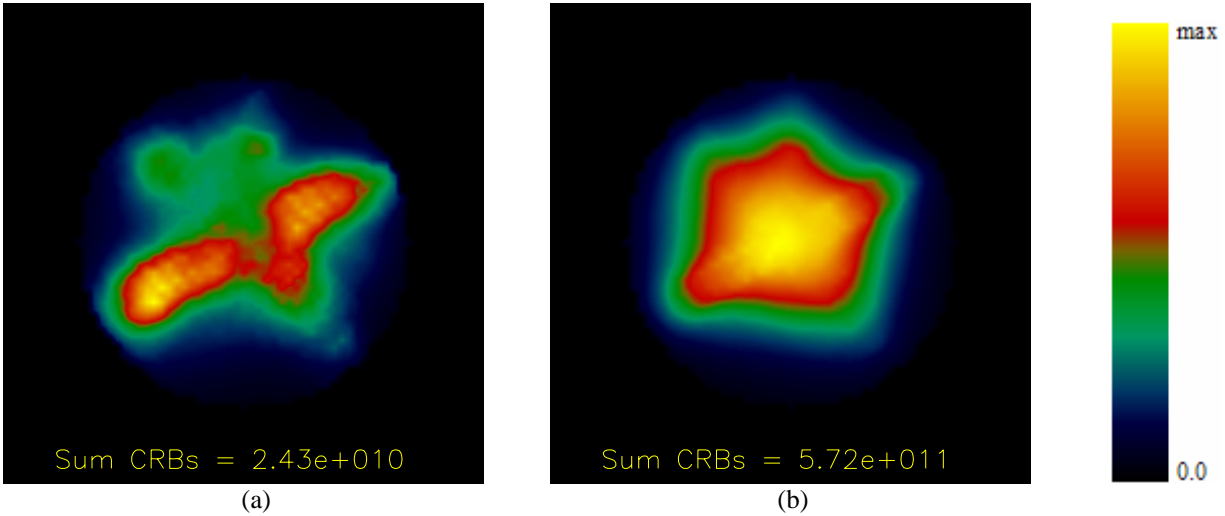

Fig. 3. Image-domain object CRBs for (a) read noise and (b) photon noise when a circular region containing the true object is used as a constraint and the PSFs are estimated pixel by pixel.

The second set of CRBs, shown in Fig. 3, are for the same set of values as for Fig. 2 except that a circular object support constraint was modeled. Notice that the sum of the CRBs has gone up in value because the support constraint size is increased. This increase occurs because we are providing less information about the underlying object. The spatial structure of the CRBs is qualitatively similar, though, to the spatial structure in Fig. 2. An interesting property of these (blind) CRBs is that they are largest in the region of the true object support, even though a circular object support constraint was modeled. The corresponding non-blind CRBs are purely a function of the size of the object support constraint, not the true object support [4].

The third set of CRBs, shown in Fig. 4, are for the same set of values as for Fig. 2 except that Zernike-based PSF parameterization using 100 terms was modeled. It can be seen that the sum of the CRBs in each CRB image are lower than the corresponding CRB sums in Fig. 2. This is due to the fact that there are fewer free parameters (100) to estimate for the PSF when Zernike-based PSF parameterization is used than for pixel-based PSF parameterization (2820 parameters).

The next two sets of CRBs, shown in Figs. 5 and 6, are plots of the sums of the image-domain object CRBs as a function of the number of measurement frames used in the image restoration process. These plots correspond to all the sets of values used for the results displayed in Figs. 2-4. The plots in Fig. 5 are for read noise and in Fig. 6 for photon noise. It can be seen that the CRBs are a decreasing function of the number of measurement frames included in the image restoration process. This is not unexpected, since the noises in each measurement frame are statistically independent and thus the use of multiple measurement frames helps average out the noise. In addition, the properties of the atmospheric blurring functions produce greater-than-expected noise reduction when including just a few measurement frames. A detailed discussion of this phenomena can be found in [5]. It can also be seen that the CRB plots do not cross as the number of measurement frames included in the estimation process is increased. This demonstrates that the properties of MFBD CRBs as discussed in the preceding figures for a single measurement frame CRBs hold for multiple measurement frames.

Fig. 4. Image-domain object CRBs for (a) read noise and (b) photon noise when the true object support is used as a constraint and the PSFs are estimated using a Zernike decomposition of the pupil phase.


Fig. 5. Plots of the sums of image-domain object CRBs, for read noise, as a function of the number of measurement frames included in the restoration process. The line with asterisks is for pixel-based PSF parameterization with a circular object support. The line with plusses is for pixel-based PSF parameterization with a true object support. The line with diamonds is for Zernike-based PSF parameterization with a true object support.

Fig. 6. Plots of the sums of image-domain object CRBs, for photon noise, as a function of the number of measurement frames included in the restoration process. The line with asterisks is for pixel-based PSF parameterization with a circular object support. The line with plusses is for pixel-based PSF parameterization with a true object support. The line with diamonds is for Zernike-based PSF parameterization with a true object support.

All of the CRBs presented and discussed previously were lower bounds to the variances of estimates of the object pixel intensities. These CRBs are useful for understanding the accuracy with which object intensities can be estimated; however, they provide no information on the achievable spatial resolution possible with MFBD. For this information, it is necessary to obtain CRBs in the Fourier domain. For our analysis, we calculated energy spectrum SNRs in the same way that it is done in speckle imaging theory [2]; that is, the energy-spectrum SNR at a given spatial frequency is equal to the true object energy spectrum divided by the square root of the energy spectrum CRB at that spatial frequency. Because we are using CRBs instead of variances, the energy spectrum SNRs to be presented are upper bounds to the achievable SNRs.

The SNRs of an object's estimated energy spectrum can be used to determine the achievable spatial resolution in the object estimate. Qualitatively, we desire to include in the image restoration all of the estimated object's Fourier spectrum whose energy spectrum SNRs are greater than a threshold level and exclude the rest. A useful (but somewhat arbitrary) SNR threshold is one, since, for this value, the signal and noise are roughly comparable in magnitude.

Energy spectrum SNRs for read noise for both pixel-based and Zernike-based PSF parameterization are plotted versus spatial frequency magnitude in Fig. 7. Each point on the plot is an average of all the two-dimensional energy spectrum SNRs for that spatial frequency magnitude. For these plots, the use of the true object support as a constraint was modeled. It can be seen that the SNRs for Zernike-based PSF parameterization are higher than for pixel-based PSF parameterization. In terms of the SNR = 1 metric, the use of Zernike-based PSF parameterization generates an achievable resolution that is approximately 30% higher than for pixel-based PSF parameterization.

In Fig. 8, energy spectrum SNRs are plotted for the photon noise case assuming that the true object support was used as a constraint. As for Fig. 7, plots for both pixel-based and Zernike-based PSF parameterization are displayed. Although the Zernike-based PSF parameterization SNRs are higher than the pixel-based PSF parameterization SNRs, they are not much higher, unlike for the pixel-based PSF parameterization case.
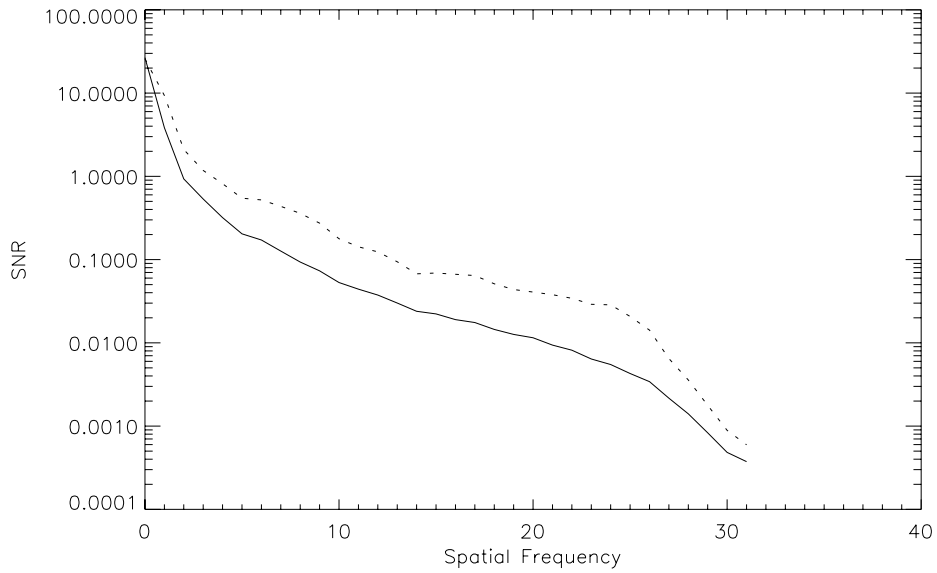
Fig. 7. Energy spectra CRBs for read noise when the true object support is used as a constraint. The solid and dashed lines correspond to pixel-based and Zernike-based PSF parameterization, respectively. The horizontal axis is spatial frequency magnitude.


Fig. 8. Energy spectra CRBs for photon noise when the true object support is used as a constraint. The solid and dashed lines correspond to pixel-based and Zernike-based PSF parameterization, respectively. The horizontal axis is spatial frequency magnitude.

5.   CONCLUSIONS

In this paper, we presented MFBD CRBs and SNRs for a number of scenarios. An analysis of these CRBs indicates that MFBD image restorations have several interesting properties. The first property is that image restorations generated using a Zernike-based PSF parameterization are less noisy and have higher spatial resolution than do image restorations generated using pixel-based PSF parameterization. The benefits of using a Zernike parameterization are greater for read noise than for photon noise. A second property is that the most noise reduction tends to occur near the edges of the true object support, even when a larger support region is used as a constraint, unlike for the non-blind case.

6.   ACKNOWLEDGEMENTS

7. REFERENCES

[1] C. L. Matson, C. C. Beckner, Jr., K. Borelli, S. M. Jefferies, E. K. Hege, and M. Lloyd-Hart, "A fast and optimal multi-frame blind deconvolution algorithm for high-resolution ground-based imaging of space objects," in preparation.

[2] M. C. Roggemann and B. Welsh, *Imaging Through Turbulence*, CRC Press, Boca Raton, FL 1996.

[3] B. Porat, *Digital Processing of Random Signals -Theory and Methods*, (Prentice-Hall, Englewood Cliffs, 1994), Ch. 3.

[4] C. L. Matson, C. C. Beckner, and K. J. Schulze, "Fundamental limits to noise reduction in images using support - benefits from deconvolution," *Proceedings of the SPIE*, vol. 5562 (2004)

[5] A. Haji, "Cramér-Rao lower bound analysis of multi-frame blind deconvolution," masters thesis, the University of New Mexico (2007).