

Achievability of Cramér-Rao Lower Bounds by Multi-Frame Blind Deconvolution Algorithms, Part II: PSF Estimation

Charles L. Matson and Charles C. Beckner, Jr.

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

Michael Flanagan

SAIC, Kirtland AFB, NM 87117

ABSTRACT

Cramér-Rao lower bound (CRB) theory can be used to calculate an algorithm-independent lower bound to the variance of any unbiased estimate of an unknown parameter. The theory also applies to joint estimation of multiple unknown parameters, to functions of estimates, and to estimates that have known bias gradients. CRBs are guaranteed to be lower bounds, but may not be achievable in practice. In particular, it is well known that algorithms that minimize cost functions to generate estimates have difficulty in achieving the CRBs for low signal-to-noise ratios. Our interest is in the achievability of the CRBs by algorithms that use a multi-frame blind deconvolution (MFBD) framework. In previous work, we analyzed the achievability of CRBs for MFBD-based estimation of object energy spectra. Here, we present initial results from our extension of this previous work to the analysis of the achievability of CRBs for MFBD-based estimation of point spread function (PSF) energy spectra that are estimated jointly with the object energy spectra.

1. INTRODUCTION

A capability of ongoing interest to the Space Situational Awareness (SSA) community is that of obtaining high-resolution pictures of satellites using ground-based telescopes. More generally, we consider the situation of taking pictures of an object (such as a satellite) through atmospheric turbulence with a noisy detector, producing images that are blurred and noisy. Although the diameter of a ground-based telescope imposes a hard limit to the resolution in images collected by that telescope, the dominant limitation for telescopes with diameters of a meter or larger (for visible light) is atmospheric turbulence. Research results published during the past decade [1] have shown that multi-frame blind deconvolution (MFBD) algorithms are effective in generating a high-resolution estimate of an object using multiple atmospherically-blurred (low-resolution) and noisy images of the object, where the blurring and noise realizations in any given image are distinct from those in any other image, while the object is the same for all the images. What is less well known is how close to the fundamental limits to image quality are these MFBD-based object estimates. There are many measures of image quality, most of which are subjective in nature. We have chosen to use an information-theory-based measure of image quality. In particular, we have chosen to look at fundamental limits in terms of the highest-possible signal-to-noise ratios (SNRs) in estimates of appropriate object parameters. To do so, it is necessary to look at fundamental limits to the variances of estimates of appropriate object parameters. We have chosen to use the Cramér-Rao lower bound (CRB) theory [2] to generate lower bounds to the variances of unbiased estimates of object energy spectra. We chose CRB theory because the lower bounds are straightforward to calculate and because they are achievable in many situations. We chose to look at variances of estimates of object energy spectra because, as discussed in a previous publication [3], analysis of the achievability of the CRBs is best undertaken in the Fourier domain, and the energy spectrum is a convenient way to do so. Furthermore, SNRs in energy spectra can be related in a simple way to resolution in the image domain. We chose to look at unbiased estimates because the quantities required to generate CRBs for biased estimates cannot be calculated, in general, due to the lack of closed-form expressions for biases.

In previous work [4], we presented an analysis of the CRB-based object energy spectrum SNRs for MFBD algorithms and compared the results to the corresponding speckle imaging object energy spectrum SNRs [9]. We showed that the MFBD SNRs were superior to the speckle imaging SNRs in virtually every case. It wasn't clear, however, that these MFBD CRB-based object energy spectrum SNRs reflect actual achievable SNRs, since CRBs are not guaranteed to be achievable. If the MFBD CRBs are not achievable, then the actual MFBD SNRs are smaller, and the results of a comparison to speckle imaging SNRs might produce different answers. For this reason, we analyzed the achievability of these object energy spectrum CRBs using our MFBD algorithm called the Physically-Constrained Iterative Deconvolution (PCID) algorithm. We showed [3] that the CRBs are indeed

achievable for SNRs on the order of one and larger when perfect support constraints are used on the object and all the PSFs (for this analysis, we generated finite-support approximations to atmospheric PSFs). The achievability became more difficult as the support constraints became less accurate.

In this paper, we present preliminary results of extending our analysis of the achievability of the CRBs from object energy spectra estimates to PSF energy spectra estimates. For many applications, the PSF variables are nuisance variables in the sense that they aren't of particular interest, but they have to be estimated in order to estimate the object quantities. Since we do have their estimates available, we decided to look to see how well we are able to estimate them in terms of achieving the CRBs to see if we could gain additional insight into the achievability of the object energy spectra CRBs. More specifically, we wanted to see how the achievability of the PSF energy spectra CRBs degrade with respect to the SNRs in the reconstructed energy spectra and with the accuracy of the support constraints. We also wanted to investigate how the achievability of the PSF energy spectra CRBs compares to the achievability of the object energy spectra CRBs.

The paper is laid out as follows: we provide background information on the imaging model, our MFBD algorithm, CRB theory, sample statistics calculations, and CRB achievability issues in Section 2; in Section 3, we present our results and discuss them, and give conclusions in Section 4.

2. BACKGROUND

We give the necessary background information in this section to understand the results presented in Section 3. In Section 2.1 we give the imaging model that describes the image formation process. A brief description of our MFBD algorithm, PCID, is provided in Section 2.2. We present a summary of the necessary CRB theory in Section 2.3, while in Section 2.4 we discuss the generation of sample statistics from PCID reconstructions. Finally, in general terms, we briefly review the achievability of CRBs when using iterative algorithms.

2.1 Imaging model

The linear space-invariant imaging equation used for the results in this paper is given by

$$i_m(\mathbf{x}) = o(\mathbf{x}) * h_m(\mathbf{x}) + n_m(\mathbf{x}); \quad m=1, \dots, M \quad (1)$$

where $i_m(\mathbf{x})$ is the m^{th} short-exposure image, $o(\mathbf{x})$ is the object, $h_m(\mathbf{x})$ is the m^{th} point spread function (PSF), $n_m(\mathbf{x})$ is the (zero-mean) m^{th} noise realization, M is the number of images used to reconstruct a single object estimate, \mathbf{x} is a two-dimensional spatial location vector, bold-face type indicates vector and matrix quantities, and $*$ denotes convolution. The noise term incorporates both Gaussian and Poisson noises with a combined variance of $\sigma_m^2(\mathbf{x}_n)$ in the n^{th} pixel of the m^{th} image. We call a PSF invertible if its Fourier transforms has no zeros.

2.2 MFBD algorithm

A detailed description of the PCID algorithm can be found in Ref.[1]. In this section we provide a brief overview of the version that we used for this paper's results. The PCID algorithm finds the (unbiased) estimate $\hat{o}(\mathbf{x})$ of $o(\mathbf{x})$ and the (unbiased) estimates $\{\hat{h}_m(\mathbf{x})\}$ of $\{h_m(\mathbf{x})\}$ that minimize the cost $J[\hat{o}(\mathbf{x}), \{\hat{h}_m(\mathbf{x})\}]$ defined by

$$J[\hat{o}(\mathbf{x}), \{\hat{h}_m(\mathbf{x})\}] = \sum_{m=1}^M \sum_{n=1}^{N_i} \frac{1}{\sigma_m^2(\mathbf{x}_n)} [i_m(\mathbf{x}_n) - \hat{i}_m(\mathbf{x}_n)]^2 \quad (2)$$

where $\hat{i}_m(\mathbf{x}) = \hat{o}(\mathbf{x}) * \hat{h}_m(\mathbf{x})$ is the model of the m^{th} measurement, N_i is the linear dimension of the square arrays containing the object, the blurring functions, and the images, and \mathbf{x}_n is the n^{th} pixel in each of these arrays. The minimization of Eq.(2) is carried out with respect to the parameters that generate $\hat{o}(\mathbf{x})$ and $\{\hat{h}_m(\mathbf{x})\}$, which are the image-domain pixel values for the results in this paper. Although the PCID algorithm has several methods of regularization that can be used to minimize the impact of Fourier-domain noise on the reconstructions, no regularization was used for our results due to CRB calculation considerations as described in Section 2.3.

There is no unique solution to the inverse problem of estimating $o(\mathbf{x})$ and $\{h_m(\mathbf{x})\}$ from the measurements $\{i_m(\mathbf{x})\}$ based on Eq.(1) alone. Additional information must be added to the inverse problem to provide the necessary uniqueness. Also, the additional information must be able to be incorporated into the CRB calculations to enable

comparisons between sample variances and the associated CRBs. For this reason, we use knowledge of the supports of $o(\mathbf{x})$ and $\{h_m(\mathbf{x})\}$ as our additional information. To guarantee uniqueness using only the supports as additional knowledge, either the knowledge of the support of $o(\mathbf{x})$ must be perfect or knowledge of all the supports of $\{h_m(\mathbf{x})\}$ must be perfect [5].

2.3 CRB theory

As stated in the introduction, we desire to calculate CRBs associated with PSF and object energy spectra estimates in order to compare sample variances to CRBs. We also need these CRBs to calculate PSF and object energy spectra SNRs. A description of how to calculate energy spectrum CRBs is given in Ref. [4]. Here we just summarize that description.

To calculate CRBs for a set of parameters, the appropriate Fisher Information Matrix (FIM) must be created. To calculate the CRBs using only a FIM, we require a set of random variables whose joint probability density function (PDF) can be written as functionally dependent upon the parameters for which we desire to calculate the CRBs. Once the FIM is generated in the standard way from the PDF, the diagonal of its inverse contains the desired CRBs.

Equation (1) provides us a set of random variables, $\{i_m(\mathbf{x})\}$, whose means are given by $\{o(\mathbf{x}) * h_m(\mathbf{x})\}$ and whose PDF functional form is given by the PDF of $\{n_m(\mathbf{x})\}$. The FIM generated using this PDF is given by

$$\mathbf{F} = \begin{bmatrix} \mathbf{F}_{11} & \cdot & \cdot & \cdot & \mathbf{F}_{1(M+1)} \\ \cdot & \cdot & \cdot & \mathbf{0} & \\ \cdot & \cdot & \cdot & & \\ \cdot & \mathbf{0} & \cdot & & \\ \mathbf{F}_{(M+1)1} & & & & \mathbf{F}_{(M+1)(M+1)} \end{bmatrix} \quad (3)$$

\mathbf{F}_{11} is the block of \mathbf{F} associated with the object pixel values, \mathbf{F}_{p1} and \mathbf{F}_{1p} for $p > 1$ are the blocks of \mathbf{F} that contain the cross-terms associated with the object pixel values and the $(p-1)^{\text{th}}$ PSF pixel values, and \mathbf{F}_{pp} for $p > 1$ is the block of \mathbf{F} associated with the $(p-1)^{\text{th}}$ PSF pixel values. Unfortunately, the diagonal of the inverse of \mathbf{F} contains the CRBs corresponding to unbiased estimates of $o(\mathbf{x})$ and $\{h_m(\mathbf{x})\}$, not unbiased estimates of their energy spectra. However, since we can write down mathematically how to transform $o(\mathbf{x})$ and $\{h_m(\mathbf{x})\}$ into their energy spectra, we can use the inverse of \mathbf{F} , along with the Jacobian of these transformations, to create a matrix whose diagonal contains the energy spectra CRBs [2]. We denote the energy-spectrum CRBs of $O(\mathbf{f})$ and $H_m(\mathbf{f})$ at a frequency \mathbf{f} as $\text{CRB}_{\text{ES}}[O(\mathbf{f})]$ and $\text{CRB}_{\text{ES}}[H_m(\mathbf{f})]$, respectively. $O(\mathbf{f})$ and $H(\mathbf{f})$ are the Fourier transforms of $o(\mathbf{x})$ and $h_m(\mathbf{x})$, respectively.

To take the inverse of \mathbf{F} , it must be invertible. Because we are modeling the object as having finite support, \mathbf{F} is theoretically invertible even with non-invertible PSFs; however, the inverse is discontinuous. This results in a condition number that is arbitrarily large. The use of a pseudo-inverse approach to avoid numerical instability significantly corrupts the accuracy of the calculated CRBs due to the truncation of non-zero singular values. The use of a true inverse approach to avoid truncation errors results in noise significantly corrupting the accuracy of the calculated CRBs. We are still exploring clever ways around this problem; for the results in this paper, we have chosen to make the PSFs invertible to make the condition number of \mathbf{F} manageable. We describe in Section 3 how we generate reasonably realistic PSFs that are also invertible.

2.4 Sample statistics calculations

The results presented in Section 3 require the calculation of CRB-based PSF and object energy spectra SNRs. They also require calculation of sample variances of energy spectrum estimates from PCID reconstructions. In this section, we describe these calculations and discuss important issues underlying the calculations.

The means to calculate the necessary CRBs are given in Section 2.3. These CRBs can be used to calculate the required SNRs as follows. We define the object energy spectrum SNR at a frequency \mathbf{f} , $\text{SNR}_{\text{ES}}[O(\mathbf{f})]$, by

$$\text{SNR}_{\text{ES}}[O(\mathbf{f})] = \frac{|O(\mathbf{f})|^2}{\{\text{CRB}_{\text{ES}}[O(\mathbf{f})]\}^{1/2}} \quad (4)$$

The PSF energy spectrum SNR is defined similarly as

$$\text{SNR}_{\text{ES}}[\mathbf{H}_m(\mathbf{f})] = \frac{|\mathbf{H}_m(\mathbf{f})|^2}{\{\text{CRB}_{\text{ES}}[\mathbf{H}_m(\mathbf{f})]\}^{1/2}} \quad (5)$$

We use the CRBs and the true values for $\mathbf{O}(\mathbf{f})$ and $\mathbf{H}_m(\mathbf{f})$ in Eqs. (4) and (5) so that the inaccuracies of estimating these quantities from the PCID reconstructions do not corrupt the SNR values. When the CRBs are achievable, these are the obtainable SNRs; when they are not, these are upper limits to the obtainable SNRs.

Understanding the sample variance calculations from the PCID reconstructions requires understanding how the reconstructions were generated. For each set of parameters for a given scenario (i.e., keeping everything fixed in Eq.(1) except for actual noise realizations), we generated 500 realizations of each set of images $\{i_m(\mathbf{x})\}$, $m=1, \dots, M$. We used PCID to generate 500 reconstructions for these 500 sets of ten image realizations. We calculated unbiased estimates of the energy-spectrum variances using standard methods. The use of 500 realizations produced sample variance estimates with reasonable accuracy.

Even though the sample variance calculations described in the previous paragraph produce unbiased estimates of the underlying variances, the parameter estimates themselves may be biased. In particular, if a parameter estimate is biased, and that bias is a function of the parameter value, the unbiased CRB expression is no longer valid. Although a biased CRB can be calculated, it requires an analytic expression for the bias gradient, which cannot be calculated except for special cases in MFBD estimation, none of which are relevant for the answers we are seeking. The key issue is not the biases in the parameter estimates; rather, it is the dependence of the biases on the parameter values. The bias gradient calculations are intractable for such a large number of parameters being estimated simultaneously; therefore, we looked at only the unbiased CRBs, recognizing that some of our conclusions may be limited by that choice.

2.5 Single-parameter SNR threshold insights

In this section we present insights gleaned from previously published results [6,7,8] on the SNR dependence of the achievability of the CRBs when estimating a single parameter using iterative algorithms. In these previous published results, the plots are of mean square errors as functions of the estimated parameter's SNRs. Since our interest is in variances, we do not include biases in our discussion here or in our results. Figure 1 shows a notional plot of the ratios of the sample variances of estimates of a single parameter to their associated CRBs as a function of the SNRs of the parameter estimates. Three distinct regions can be observed in this plot. In the high SNR region (SNRs greater than SNR_1), the CRBs are achievable and the ratios are equal to one. In the transition region between SNR_1 and SNR_2 , the CRBs are no longer achievable and the ratios increase as the SNRs decrease. In this region the parameter estimates are typically becoming biased. In the low SNR region (SNRs less than SNR_2), the ratios decrease as the SNRs decrease. In this region the parameter estimates approach a bias determined by the problem formulation and thus the variances decrease and can be less than the unbiased CRBs.

The behavior shown in Fig. 1 is for a single parameter estimate; i.e., each point on the curve corresponds to a single estimation problem. This means that the lack of achievability of the CRBs for SNRs less than SNR_1 has no impact on the achievability of the CRBs for SNRs greater than SNR_1 . For the results in Section 3, thousands of parameter are estimated jointly, and the range of their SNRs is orders of magnitude. In addition to the questions given in Section 1, we also wish to determine if the curve in Fig. 1 is applicable in the case of joint parameter estimation.




Fig. 1. Notional plot of the ratios of sample variances to CRBs as a function of SNR for a single parameter estimate

3. RESULTS

In this section, we present some preliminary answers to the questions we posed in Section 1. We begin by describing in Section 3.1 how we generated the measurements and the range of parameters used for their generation. Section 3.2 contains the results when we used perfect support constraints to reconstruct the object and all the PSFs. In Section 3.3, we show results for non-perfect object support, and in Section 3.4 we show results for non-perfect PSF support.

3.1 Parameters

Each set of PCID reconstructions of an object and PSFs was obtained from a set of ten noisy measurement frames and a fixed set of support constraints. Each noise-free measurement frame was generated by convolving a true object with a PSF. A set of ten noise-free measurement frames was produced with the use of ten distinct PSFs and a single true object. Five hundred sets of ten noisy measurements each were obtained by first setting a photon level and a read noise variance, and then adding unique noise realizations to each noise-free measurement. The PCID software was used to generate an object reconstruction and reconstructions of all ten PSFs from each set of ten noisy measurements. Each reconstruction's energy spectrum was calculated and sample variances were calculated.

The OCNR satellite model, shown in Fig. 2(a), was used as the true object for the results reported here. We also used the two-circ model [3]; however, the results were similar to the results from the OCNR model, so we report only on the OCNR results.

We used one set of ten PSFs to generate noise-free measurements. A representative PSF from the set is shown in Fig. 2(b). The PSFs were generated by first creating a set of ten PSFs using standard methods [9] that correspond to a telescope of diameter D looking vertically through an atmosphere with a spatial coherence scale r_o of $D/8$. These PSFs were made invertible by replacing their values outside a circle of finite radius with zeros. The circle's radius was chosen so that at least 99% of the PSF's energy was inside the circle. The tradeoffs associated with increasing the circle's radius to include a higher fraction of energy, which increases the fidelity of the model to real-world conditions but also increases the condition number of the FIM, are under investigation.



Fig. 2. True object and PSFs used for the results in Section 3: (a) OCNR object, (b) a PSF from the set of ten PSFs

We varied the photon level in the measurements from 10^6 to 10^8 in order to generate SNR levels in the reconstructed object and PSF energy spectra that encompassed the SNR_1 value in Fig. 1. Because the conclusions we draw in this paper do not depend on the photon level, we show only the results for 10^7 photons. The support constraints used in the PCID reconstructions are described in the sections where they were used. All results in Section 3 have the read noise variance equal to 100.

3.2 Results – Perfect Support Constraints

In this section, we present results showing how the achievability of the CRBs for PSF energy spectra depends upon the SNRs in the reconstructed PSF energy spectra when perfect support constraints are used. We also compare these results to the corresponding object energy spectra results that are estimated jointly with the PSF energy spectra.

The PSF energy spectra results for one of the ten reconstructed PSFs are shown in Fig. 3, while in Fig. 4 the corresponding object energy spectra results are displayed. In Fig. 3(a), the ratios of the reconstructed PSF energy spectra sample variances to their corresponding CRBs are plotted as a function of the SNRs of the reconstructed PSF energy spectra. Recall that the ratios will be equal to one if the CRBs are achieved by the estimates. It can be seen that the ratios are approximately equal to one for all SNRs greater than one, but increase quickly as the SNRs become less than one (i.e., $\text{SNR}_1 = 1$). This is the same result, qualitatively, as for the object energy spectra

reconstructions as can be seen in Fig. 4(a); however, the SNR_1 value for the object is a little larger than the value for the PSF, indicating that the PSF energy spectra CRBs are achieved for slightly lower SNRs. Figures 3(b) and 4(b) display the ratios from Figs. 3(a) and 4(a), respectively, as a function of spatial frequency location where zero spatial frequency is in the center of the figure. Notice that the CRBs are achieved for lower spatial frequencies but not for higher spatial frequencies, which is just a restatement of the information shown in the plots in Figs. 3(c) and 4(c) that show that higher SNRs are found at lower spatial frequencies. The minimal amount of structure seen in Figs. 3(b) and 4(b) at lower spatial frequencies is another way to see that the CRBs are achieved for all higher SNRs. The “whiteness” of the ratios in Fig. 3(b) at high spatial frequencies as compared to the object results in Fig. 4(b) are a result of the SNRs of the object spectra at higher spatial frequencies being larger than the corresponding PSF SNRs. This can be seen from the SNR plots in Figs. 3(c) and 4(c).




Fig. 3. One of ten PSF energy-spectrum results for perfect support constraints: (a) a scatter plot of the ratios of the sample variances to the CRBs versus the SNRs of the reconstructed PSF energy spectra, (b) a 2-D display of these ratios, where the horizontal and vertical axes are the two spatial frequency axes and the center of the array is the location of the zero spatial frequency value, (c) a plot of the CRB-based reconstructed PSF energy spectrum SNRs as a function of spatial frequency magnitude normalized to one at the largest value. In order to reveal detail at lower sample-variance-to-CRB-ratio values, all values in (b) greater than four are truncated to four.




Fig. 4. Object energy-spectrum results for perfect support constraints: (a) a scatter plot of the ratios of the sample variances to the CRBs versus the SNRs of the reconstructed object energy spectra, (b) a 2-D display of these ratios, where the horizontal and vertical axes are the two spatial frequency axes and the center of the array is the location of the zero spatial frequency value, (c) a plot of the CRB-based reconstructed object energy spectrum SNRs as a function of spatial frequency magnitude normalized to one at the largest value. In order to reveal detail at lower sample-variance-to-CRB-ratio values, all values in (b) greater than four are truncated to four.

3.3 Results – Perfect PSF Support Constraints, Loose Object Support Constraint

Next, we show results where the PSF support constraints used in the PCID algorithm are perfect, but the object support constraint is not. We used the smallest circle that contains the object (roughly four times the area of the true object support) as the object support constraint. We compare the PSF and object energy spectra results to each other and to the perfect-support case results in Section 3.2.

The PSF energy spectra results for one of the ten reconstructed PSFs are shown in Fig. 5, while in Fig. 6 the corresponding object energy spectra results are displayed. In Fig. 5(a), the ratios of the reconstructed PSF energy

spectra sample variances to their corresponding CRBs are plotted as a function of the SNRs of the reconstructed PSF energy spectra. In addition, the sample variance/CRB ratios for perfect support constraints are plotted in red to assist in comparing the perfect-support-constraint results to these results. Notice that the SNR_1 value for the non-perfect-object-support case is larger than for the perfect support case, indicating that the reconstructed PSF energy spectra CRBs are less achievable in this case. In contrast, however, the results in Fig. 6 show that the SNR_1 value for the reconstructed object energy spectra is the same as in the perfect support constraint case. This result is surprising, because the object support constraint is loose and the PSF support constraints are perfect; thus, we were expecting to see that the object CRBs were less obtainable than the PSFs. Also, it can be seen in Figs. 5(b) and 6(b) that there is some structure emerging in the lower spatial frequencies locations where, according to Figs. 5(c) and 6(c), the SNRs are well above one. This means that the CRBs are becoming less obtainable even at the higher SNRs. This can also be seen in Figs. 5(a) and 6(a), where at the higher SNR values the black trace has values higher than the red trace.




Fig. 5. One of ten PSF energy-spectrum results for perfect PSF support constraints and a loose object support constraint: (a) a scatter plot of the ratios of the sample variances to the CRBs versus the SNRs of the reconstructed PSF energy spectra, (b) a 2-D display of these ratios, where the horizontal and vertical axes are the two spatial frequency axes and the center of the array is the location of the zero spatial frequency value, (c) a plot of the CRB-based reconstructed PSF energy spectrum SNRs as a function of spatial frequency magnitude normalized to one at the largest value. In order to reveal detail at lower sample-variance-to-CRB-ratio values, all values in (b) greater than four are truncated to four. The red trace in (a) contains the sample variance/CRB ratios from Fig. 3 to assist in evaluating the impact of non-perfect support constraints.




Fig. 6. Object energy-spectrum results for perfect PSF support constraints and a loose object support constraint: (a) a scatter plot of the ratios of the sample variances to the CRBs versus the SNRs of the reconstructed object energy spectra, (b) a 2-D display of these ratios, where the horizontal and vertical axes are the two spatial frequency axes and the center of the array is the location of the zero spatial frequency value, (c) a plot of the CRB-based reconstructed object energy spectrum SNRs as a function of spatial frequency magnitude normalized to one at the largest value. In order to reveal detail at lower sample-variance-to-CRB-ratio values, all values in (b) greater than four are truncated to four. The red trace in (a) contains the variance/CRB ratios from Fig. 4 to assist in evaluating the impact of non-perfect support constraints.

3.4 Results – Perfect Object Support Constraint, Loose PSF Support Constraints

Finally, we show results where the object support constraint is perfect, but the PSF support constraints are not. We used as PSF support constraints circles whose centers are at the same location as the center of the perfect-support circles, but have an area that is four times as large. We compare the PSF and object energy spectra results to each other, to the perfect-support results in Section 3.2, and to the non-perfect-object-support results in Section 3.3.

The PSF energy spectra results for one of the ten reconstructed PSFs are shown in Fig. 7, while in Fig. 8 the corresponding object energy spectra results are displayed. In Fig. 7(a), the ratios of the reconstructed PSF energy spectra sample variances to their corresponding CRBs are plotted as a function of the SNRs of the reconstructed PSF energy spectra. In addition, the sample variance/CRB ratios for perfect support constraints are plotted in red. Notice that the SNR_1 value for the non-perfect-PSF-support case is the same as the SNR_1 value for the perfect support case. By contrast, the object energy spectra results in Fig. 8(a) show that its SNR_1 value is greater for non-perfect-PSF-support case than for perfect supports. As for the results in Section 3.3, this is unexpected, but it is consistent with the Section 3.3 results.

Notice also that the non-perfect-PSF-support ratios aren't as close to one for higher SNR values as are the ratios for perfect support constraints. In Figs. 7(b) and 8(b), with the help of Figs. 7(c) and 8(c), respectively, it can be seen that there are obvious regions in Fourier space where the SNRs are well above one but the ratios are greater than one. The presence of these regions in Figs. 7(a) and 8(a) are indicated by the humps in the black traces for high SNR values. The sizes of these areas and their values show that the CRBs are much less achievable at high SNRs than for the cases of either perfect supports, or perfect PSF supports with imperfect object support. This, at least, seems intuitively correct, because having imperfect PSF support constraints means that ten of the eleven support constraints used in the PCID algorithm are not perfect. An imperfect object support constraint with perfect PSF support constraints means that only one of the eleven support constraints are imperfect.




Fig. 7. One of ten PSF energy-spectrum results for a perfect object support constraint and loose PSF support constraints: (a) a scatter plot of the ratios of the sample variances to the CRBs versus the SNRs of the reconstructed PSF energy spectra, (b) a 2-D display of these ratios, where the horizontal and vertical axes are the two spatial frequency axes and the center of the array is the location of the zero spatial frequency value, (c) a plot of the CRB-based reconstructed PSF energy spectrum SNRs as a function of spatial frequency magnitude normalized to one at the largest value. In order to reveal detail at lower sample-variance-to-CRB-ratio values, all values in (b) greater than four are truncated to four. The red trace in (a) contains the sample variance/CRB ratios from Fig. 3 to assist in evaluating the impact of non-perfect support constraints.




Fig. 8. Object energy-spectrum results for a perfect object support constraint and loose PSF support constraints: (a) a scatter plot of the ratios of the sample variances to the CRBs versus the SNRs of the reconstructed object energy spectra, (b) a 2-D display of these ratios, where the horizontal and vertical axes are the two spatial frequency axes and the center of the array is the location of the zero spatial frequency value, (c) a plot of the CRB-based reconstructed object energy spectrum SNRs as a function of spatial frequency magnitude normalized to one at the largest value. In order to reveal detail at lower sample-variance-to-CRB-ratio values, all values in (b) greater than four are truncated to four. The red trace in (a) contains the variance/CRB ratios from Fig. 4 to assist in evaluating the impact of non-perfect support constraints.

4. CONCLUSIONS

We have presented preliminary results that demonstrate some properties of the achievability of PSF energy spectra CRBs with the use of MFBD algorithms, as a function of the SNRs in the reconstructed spectra and the accuracy of the support constraints. We compared these results with previous results for object energy spectra. We showed that for reconstructed energy spectra SNRs of one or greater, the CRBs are achievable for perfect support constraints. When either the object or PSF support constraints are not perfect, two important behaviors were seen. First, the CRBs were unachievable for some spatial frequency regions where the reconstructed PSF energy spectra SNRs were well above the SNR_1 value, and the size of these regions increased as the number of non-perfect support constraints increased. Second, counter-intuitively, when the object support constraint is loosened while the PSF support constraints remain perfect, it is the PSF SNR_1 value that increases, not the object SNR_1 value. Conversely, when the PSF support constraints are loosened while the object support constraint remains perfect, it is the object SNR_1 value that increases.

In future work, we will investigate the non-intuitive behavior of the SNR_1 value dependence on object and PSF supports. We will also investigate the impact of increasing the percentage of energy included in the support-constrained versions of the non-invertible PSFs that make them invertible. We also plan to look into the properties of the regions in Fourier space where the SNRs are high and yet the CRBs are not obtained when the support constraints are loosened. It appears that these regions are where the SNRs of the energy spectra are particularly high, not low – another non-intuitive result.

5. ACKNOWLEDGEMENTS

The authors wish to thank the Air Force Office of Scientific Research and the Air Force Research Laboratory for their support that made this research possible.

6. REFERENCES

1. C. L. Matson, K. Borelli, S. Jefferies, C. C. Beckner Jr., E. K. Hege, and M. Lloyd-Hart, "Fast and optimal multi-frame blind deconvolution algorithm for high-resolution ground-based imaging of space objects," *Applied Optics* **48**, A75-A92 (2009).
2. B. Porat, *Digital Processing of Random Signals -Theory and Methods*, (Prentice-Hall, Englewood Cliffs, 1994).
3. C. L. Matson, M. Flanagan, and R. A. Vincent, "The impact of low signal-to-noise ratio values on the achievability of Cramér-Rao lower bounds with multi-frame blind deconvolution algorithms," *Proc. SPIE* **7828** (2010).
4. C. L. Matson, "A comparison of multi-frame blind deconvolution and speckle imaging energy spectrum signal-to-noise ratios," *Proc. SPIE* **7108** (2008).
5. R. G. Lane and R. H. T. Bates, "Automatic multidimensional deconvolution," *J. Opt. Soc. Am. A* **4**, 180-188 (1987).
6. W. Xu, "Cramér-Rao bound for bearing estimation with bias correction," *Proc. MTS/IEEE Oceans*, 1-5 (2007).
7. E. Chaumette, J. Galy, A. Quinlan, and P. Larzabal, "A new Barankin bound approximation for the prediction of the threshold region performance of maximum likelihood estimators," *IEEE Trans. Sig. Proc.* **56**, 5319-5333 (2008).
8. A. Renaux, L. Najjar-Atallah, P. Forster, and P. Larzabal, "A useful form of the Abel bound and its application to estimator threshold prediction," *IEEE Trans. Sig. Proc.* **55**, 2365-2369 (2007).
9. M. C. Roggemann and B. Welsh, *Imaging Through Turbulence*, (CRC Press, Boca Raton, 1996).