



Calhoun: The NPS Institutional Archive
DSpace Repository

Faculty and Researchers

Faculty and Researchers Collection

1995-05

**Analysis of the CLEAN algorithm and
implications for superresolution**

Fried, David L.

<http://hdl.handle.net/10945/44104>

Downloaded from NPS Archive: Calhoun



Calhoun is a project of the Dudley Knox Library at NPS, furthering the precepts and goals of open government and government transparency. All information contained herein has been approved for release by the NPS Public Affairs Officer.

Dudley Knox Library / Naval Postgraduate School
411 Dyer Road / 1 University Circle
Monterey, California USA 93943

<http://www.nps.edu/library>

Analysis of the CLEAN algorithm and implications for superresolution

David L. Fried

Department of Physics, Naval Postgraduate School, Monterey, California 93943-5000

Received January 6, 1994; revised manuscript received November 18, 1994; accepted December 5, 1994

The capability of the CLEAN algorithm, which is able to develop image information corresponding to spatial frequencies for which the imaging system's optical transfer function (OTF) is equal to zero, is shown to be dependent on the limited size of the object being imaged. It is found that this capability is available without a severe signal-to-noise-ratio penalty only for the recovery of a spatial frequency that is sufficiently close to some other spatial frequency for which the OTF is not equal to zero. As used here the term "sufficiently close" means that the magnitude of the separation of the spatial frequencies is less than one half of the inverse of the size of the object being imaged. This represents a limitation of CLEAN's capability deriving from object size. It is suggested that this capability can be thought of in terms of superresolution, with the same limitation in regard to object size.

1. INTRODUCTION AND SUMMARY

The origin of the research reported here lies in an effort to develop an understanding of the CLEAN algorithm^{1,2} (also referred to simply as CLEAN), used in the astronomical community to remove the ringing from images produced by sparse array imaging systems. In removing the ringing, the CLEAN algorithm is apparently restoring to the image information concerning spatial frequencies for which the array's optical transfer function (OTF) is equal to zero. In the accomplishment of this there is a suggestion of superresolution,³ and my colleagues and I were surprised by the seeming ability of superresolution to function as well as it appeared to in this case without requiring an excessively high signal-to-noise ratio, which we had thought was necessary in superresolution.⁴ The results that we have developed in carrying out this research have provided not only an explanation of the ability of the CLEAN algorithm to perform as it does but also an insight into the capabilities and limitations of superresolution.

In order to apply the results to the matter of understanding superresolution, we must first have an adequate definition and understanding of exactly what should and what should not be included in the concept of superresolution. Because presentation of any explanation of superresolution entails a need to go all the way back to a discussion of exactly what constitutes superresolution, and because offering such an explanation would necessarily distract from the presentation of the basic work concerning CLEAN, all of the discussion of superresolution, except for a few words of overview presented in this section, has been placed in Section 5 at the end of this paper.

The CLEAN algorithm functions in accordance with the following five steps:

1. One starts with a copy of the so-called dirty image (which manifests all the ringing-related artifacts that are associated with the severely sidelobe-contaminated point-spread function of the sparse-array imaging sys-

tem), calling this copy the residual dirty image, and finds the brightest pixel in this residual dirty image.

2. A blank image field, representing the clean image, is established with all pixels set equal to zero.

3. The value of the clean-image pixel whose position corresponds to that of the brightest pixel in the dirty image is incremented by some fraction of the value of that brightest pixel in the residual dirty image.

4. With the array's point-spread function providing the pattern, and the location and the magnitude of what has just been added to the clean image providing the pattern's shift and scaling, appropriate values are subtracted from the residual dirty image.

5. One then reexamines the residual dirty image to find what is now its brightest pixel.

The process of adding to the clean image and subtracting from the dirty image, defined by steps 3, 4, and 5, is cyclically repeated, over and over, continuously building up the clean image and reducing the residual dirty image. The process is subject to the constraint that at the start of each repetition of the cycle the sum of the square of the pixel values in the residual dirty image is calculated, and when the value of that sum stops decreasing, the process is stopped and the clean image is considered to be complete (except possibly for the application of some smoothing of the image). The process of selection of the brightest pixel and the decision to terminate the process are inherently nonlinear operations. It is presumably for this reason that thus far it has not proven possible to carry out a closed-form analysis of CLEAN nor to identify the basis for its ability to recover missing spatial-frequency components.

The approach taken to the problem of understanding how and how well CLEAN works starts with consideration of what a simple linear estimator would be able to do. For this purpose I choose to work with a minimum mean-square-discrepancy algorithm. I consider the dirty image to be the product of convolving the object's true image with the array's point-spread function, with some noise

value then added to each pixel of the dirty image. This is a linear process, and accordingly one is able to develop a linear, closed-form solution to answer the question, What set of pixel values, i.e., what (clean) image (intended to serve as an estimate of the true image) when convolved with the array's point-spread function will result in a set of pixel values that will be minimally discordant, in a minimum mean-square-discrepancy sense, with the pixel values of the dirty image? Using this solution one then determines what is the mean-square discrepancy between the pixel values of the true image and the pixel values of the (clean) image generated by this least-square process. This mean-square discrepancy (or a quantity proportional to this discrepancy) was taken as a measure of the quality of the image produced by the algorithm.

Calculation was done for a one-dimensional problem involving an imaging array with holes in its OTF coverage pattern. An array was used with an overall length L and was considered to be operating at a wavelength λ . Figure 1 shows the kind of array used, and Fig. 2 shows the OTF for this particular array. (Although a number of different array layouts were used in the study, the one shown here is indicative of all of them. Since the same basic results were obtained with each of the arrays, in this paper the results presented are restricted to those obtained with the particular array shown in Fig. 1.) Note from Fig. 2 that the OTF is full of small holes, as well as one larger hole, in its coverage. Each of the small holes has a width of the order of $1/17$ of the array's cutoff frequency, i.e., $1/17 L/\lambda$. The larger hole extends from $\sim 13/17$ to $\sim 16/17$ of the array's cutoff spatial frequency, i.e., from $13/17 L/\lambda$ to $16/17 L/\lambda$, and has a width of $\sim 3/17 L/\lambda$.

When an extensive set of results obtained with this minimum mean-square-discrepancy algorithm was studied, it was observed that the rms error in the (clean) image produced by the algorithm was modest as long as the size of the object was small enough that its inverse was greater than approximately the width of the larger hole in the OTF. It was found that if the object was larger than that, the error could still be kept small if a sharp-cutoff low-pass filter was applied to the final image, with the filter's cutoff set so that all spatial frequencies larger than $13/17 L/\lambda$ were excluded from the final image; the error would be small provided that the object was not larger than the inverse of the width of the small holes. As long as the object size was less than approximately the inverse of the size of the small holes, the missing components for all the smaller holes in the OTF were satisfactorily estimated. The missing components for the larger hole could not be estimated without introduction of considerable additional noise into the result if the object was larger than the inverse of the size of the larger hole. If the object was larger than the inverse of the size of the smaller holes, then even the missing spatial-frequency components for the small holes could not be satisfactorily estimated and no satisfactory image could be obtained, without a severe noise penalty.

The minimum mean-square-discrepancy algorithm's performance in cleaning up the image seemed to be quite good; i.e., there was no excessive noise sensitivity, provided that the object's size was small enough. Surprisingly, for small-enough objects the algorithm was apparently developing good estimates of spatial-frequency

components for spatial frequencies for which the array's OTF was equal to zero.

Observation of this fact led us to consider the possibility that for small objects there is some unavoidable correlation between spatial-frequency components for spatial frequencies that are sufficiently close together, a correlation that is unavoidably present simply because the object is small. (It seemed that there was little else that an algorithm as simple as a linear estimator, which is what the minimum mean-square-discrepancy algorithm is, could be using to produce, in effect, the estimates of the missing spatial-frequency components. If this correlation were present then the algorithm could be forming an estimate of the unmeasured spatial-frequency components from observation of the value of the measured components.) The results of an analysis of this matter, as presented below, show that this correlation is present and that the correlation function, which is dependent on the size of the object being imaged, has a range of the right size to explain the observed algorithm performance.

These results, of course, were for the minimum-square-discrepancy algorithm, and although they were suggestive for CLEAN, they were not firmly linked to CLEAN. To tie the results more directly to the CLEAN algorithm, we wrote a computer program that simulated the operation of CLEAN, and, operating it in a Monte Carlo mode, we used the program to develop mean-square image error results for targets of various sizes. In each case the resultant image (namely, the clean image) was subject to low-pass filtering with a sharp cutoff at $13/17 L/\lambda$. The mean-square image error was calculated from the difference between the clean image and the image that we obtained by similarly low-pass filtering the true image. Applying this low-pass filter to the clean image meant that the only missing spatial-frequency components for which the algorithm had to develop estimates were those whose spatial frequency fell within one of the small holes of the array's OTF. We found that the mean-square error in the clean image was fairly constant, independent of the size of the object being imaged, provided that the object size was smaller than the inverse of the size of the smaller

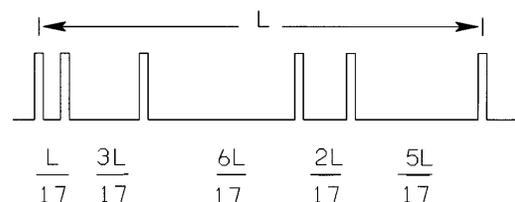


Fig. 1. Array pattern for a six-element, one-dimensional, nonredundant array.

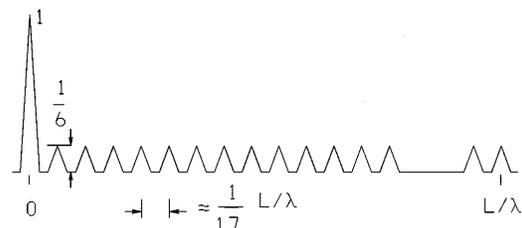


Fig. 2. OTF for the six-element, one-dimensional, nonredundant array shown in Fig. 1.

holes in the array's OTF. This was in good agreement with what we obtained from the minimum-mean-square-discrepancy algorithm.

It was concluded from this that the operation of CLEAN is based on the same consideration as is the operation of the minimum-mean-square-discrepancy algorithm, that both exploit the correlation that necessarily exists for small objects: the correlation between any two spatial-frequency components when the spatial frequencies are sufficiently close to each other. For both algorithms the value of the missing spatial-frequency component apparently is in effect being estimated on the basis of its correlation with the value of some other, directly measurable, spatial-frequency component. We believe that for both algorithms the key feature is the finite size of the object, which is responsible for the existence of the correlation. We consider that exploitation of the size of the object is (essentially without our noticing it) built into the algorithm in the algorithm's initial formulation; as the algorithms are formulated, the measured dirty image is to be explained in terms of a finite-size array of object pixels. The minimum-mean-square-discrepancy approach formulates a set of simultaneous equations to obtain a clean image for which the discrepancy with respect to the dirty image is minimum (and then solves this set of simultaneous equations directly, potentially with matrix inversion methods). Because the problem was formulated with a finite number of pixels and a minimum-mean-square-discrepancy solution was sought, the finite size of the object and the exploitation of the consequent correlation were incorporated automatically.

It is my opinion that in effect the CLEAN algorithm approaches the problem with the same tools. The finite size of the object is definitely built into CLEAN. I view the steps of the CLEAN algorithm as representing an indirect approach to the solution of the same minimum-mean-square-discrepancy problem, a solution by a numerical relaxation method rather than by the more direct approach of matrix inversion and matrix multiplication. Apparently CLEAN was formulated with the attitude that for real images the number of simultaneous equations that would have to be solved for the discrepancy to be minimized is prohibitively large, and, accordingly, CLEAN attempts to minimize the discrepancy by what is in effect a numerical relaxation method.

As remarked above, the development of an estimate of the value of spatial-frequency components of the object's image, components for which the imaging system's OTF is equal to zero, is seen as a manifestation of superresolution. Admittedly, all the spatial frequencies involved in the above discussion of results are lower than the upper-limit cutoff frequency of the arrays, and it may accordingly be argued that these results ought not to be taken as having to do with superresolution. However I believe that these results should be considered to be an example of superresolution—that developing information for spatial frequencies for which the imaging system's OTF is equal to zero is superresolution no matter whether the spatial frequencies are above or below the system's diffraction-limited cutoff frequency—and in Section 5 below I present thoughts as to just what constitutes superresolution and why these results should be considered to be a manifestation of it. Moreover, we have found

evidence in our numerical results for CLEAN that image processing is able to recover spatial-frequency components for frequencies that are a bit above the imaging system's upper-limit cutoff, without much of a signal-to-noise-ratio penalty. Just how far above the cutoff components can be recovered depends on the object size; apparently how far corresponds to a distance, in frequency space, equal to approximately one half of the inverse of the object size. The analytic and the numerical results to be presented provide some direct support for the conclusion that spatial-frequency components can be recovered for spatial frequencies that are greater than the imaging system's cutoff frequency—greater by an amount (i.e., a frequency increment) equal to approximately one half of the inverse of the size of the object being imaged.

On the basis of the above, I conjecture that for a filled circular aperture of diameter D imaging at a wavelength λ , for which the (diffraction-limited) spatial-frequency cutoff is D/λ (cycles/rad_{fov}) if we are to be able to develop a superresolution estimate of the value of the image components for a spatial frequency of the order of $1.5D/\lambda$, i.e., for just one and one half times the diffraction-limited cutoff (without requiring an excessive signal-to-noise ratio), then the object can be no longer than λ/D rad_{fov} if substantial noise penalties are to be avoided. [The notation rad_{fov} is used to make clear that it is radians of field of view (fov) and not radians in a cycle.]

Section 2 presents the analysis leading to the conclusion that there is a correlation among different spatial-frequency components of a small object. Section 3 is a discussion of the numerical results obtained from the closed-form analytic results for the minimum-mean-square-discrepancy linear algorithm. Section 4 presents the results obtained with the CLEAN simulation run in the Monte Carlo mode, along with conclusions. Section 5 is a discussion of superresolution; an attempt is made there to establish a link between the CLEAN-related results and superresolution.

2. CORRELATION OF SPATIAL-FREQUENCY COMPONENTS

As stated above, we believe that the fundamental fact explaining why either of these two algorithms works is that for a finite-sized object (on a dark background) the spatial-frequency components are necessarily correlated if they are close together; they are strongly correlated if the magnitude of the difference of the two spatial frequencies is less than approximately one half of the inverse of the object size. This correlation is an unavoidable consequence of the finite size of the object.

To see that this is the case, consider an infinite-extent random pattern $f(\mathbf{r})$, with different spatial-frequency components $\tilde{f}(\boldsymbol{\kappa})$ being completely uncorrelated, where

$$\tilde{f}(\boldsymbol{\kappa}) = \int d\mathbf{r} \exp(-2\pi i \boldsymbol{\kappa} \cdot \mathbf{r}) f(\mathbf{r}) \quad (1)$$

and $\boldsymbol{\kappa}$ denotes a spatial-frequency vector. We cut a finite-sized circular pattern of diameter D from this infinite pattern by multiplying by $W(\mathbf{r})$, where

$$W(\mathbf{r}) = \begin{cases} 1 & \text{if } \mathbf{r} \leq 1/2 D, \\ 0 & \text{otherwise} \end{cases}, \quad (2)$$

to obtain

$$f_W(\mathbf{r}) = f(\mathbf{r})W(\mathbf{r}). \quad (3)$$

We are interested in the correlation between two spatial-frequency components, $\langle \tilde{f}_W(\boldsymbol{\kappa})\tilde{f}_W^*(\boldsymbol{\kappa}') \rangle$, where the angle brackets denote an ensemble average and

$$\tilde{f}_W(\boldsymbol{\kappa}) = \int d\mathbf{r} \exp(-2\pi i \boldsymbol{\kappa} \cdot \mathbf{r}) f_W(\mathbf{r}), \quad (4)$$

is the component for spatial frequency $\boldsymbol{\kappa}$. With some rather straightforward manipulation of multiple integrals we can show that

$$\begin{aligned} \langle \tilde{f}_W(\boldsymbol{\kappa})\tilde{f}_W^*(\boldsymbol{\kappa}') \rangle &= \int d\mathbf{x} \Phi[1/2(\boldsymbol{\kappa} + \boldsymbol{\kappa}') + \mathbf{x}] \\ &\quad \times \widetilde{W}[\mathbf{x} - 1/2(\boldsymbol{\kappa} - \boldsymbol{\kappa}')] \\ &\quad \times \widetilde{W}[\mathbf{x} + 1/2(\boldsymbol{\kappa} - \boldsymbol{\kappa}')], \end{aligned} \quad (5)$$

where

$$\begin{aligned} \widetilde{W}(\boldsymbol{\kappa}) &= \int d\mathbf{r} \exp(-2\pi i \boldsymbol{\kappa} \cdot \mathbf{r}) W(\mathbf{r}) \\ &= (1/4\pi D^2) \frac{2J_1(\pi \kappa D)}{\pi \kappa D} \end{aligned} \quad (6)$$

and $\Phi(\boldsymbol{\kappa})$ is the power spectral density associated with the infinite extent pattern $f(\mathbf{r})$ and is defined by the expression

$$\Phi(\boldsymbol{\kappa}') = \int d\mathbf{r} \exp(-2\pi i \boldsymbol{\kappa}' \cdot \mathbf{r}) \langle f(\mathbf{r}' + 1/2\mathbf{r}) f(\mathbf{r}' - 1/2\mathbf{r}) \rangle. \quad (7)$$

Equation (5) is the basic result, indicating the correlation's value. From consideration of the width (full width at half-maximum) of the main lobe of $\widetilde{W}(\boldsymbol{\kappa})$ —a width approximately equal to D^{-1} —along with consideration of Eq. (5), we can see that there will be a significant correlation between $\tilde{f}_W(\boldsymbol{\kappa})$ and $\tilde{f}_W(\boldsymbol{\kappa}')$ provided that

$$|\boldsymbol{\kappa} - \boldsymbol{\kappa}'| < 1/2 D^{-1}, \quad (8)$$

this being the condition for there to be a significant correlation of components.

We attribute to this correlation the ability of CLEAN and of the minimum-mean-square-discrepancy algorithm to develop useful estimates of missing or unmeasured spatial-frequency components of an image. We consider the ability of these algorithms to form an estimate of the amplitude of some spatial-frequency component for which the OTF is equal to zero to be based (in an indirect and therefore not immediately noticeable way) on the fact that the unmeasured amplitude, as Eq. (5) demonstrates, may be strongly correlated with an amplitude that is measured, i.e., for which the OTF is not equal to zero, if the object size is small enough. I introduce the existence of this correlation into the formulation of each of the algorithms, without any notice being called to it, by restricting the solution space to a finite number of pixels and thereby imposing the assumption of a finite object size. Thereafter the algorithm proceeds to find an object pixel pattern (the clean image) that can be reconciled with the array's transfer function and with the original measurement values (the dirty image), and in doing so it takes account (essentially without noting it) of the fact that the object has some particular finite size and thus of the above-noted correlation.

3. ANALYSIS AND RESULTS FOR THE MINIMUM-MEAN-SQUARE-DISCREPANCY ALGORITHM

The minimum-mean-square-discrepancy estimation algorithm has been formulated here in a manner that takes account of the point-spread function of the imaging system and finds the finite-sized object pattern that best explains the measured image pattern. For analytic and computational convenience the problem was formulated in a spatially quantized, i.e., a pixel-type, form. The analysis was carried out with matrix notation. Written with this notation, this process of developing the algorithm can be expressed as follows.

Letting \mathbf{m} represent a column vector of measured pixel values, \mathbf{o} a column vector of object pixel values, and \mathbf{n} a column vector of noise values (associated with the measurement values in \mathbf{m} and with statistical independence among the elements of the noise column vector), we have

$$\mathbf{m} = \mathbf{S}\mathbf{o} + \mathbf{n}, \quad (9)$$

where \mathbf{S} is a matrix representing the imaging array's point-spread function. The minimum-mean-square-discrepancy solution for \mathbf{o} given \mathbf{m} is known to be given by the equation

$$\hat{\mathbf{o}} = \mathbf{A}\mathbf{m}, \quad (10)$$

where

$$\mathbf{A} = (\mathbf{S}^T \mathbf{S})^{-1} \mathbf{S}^T. \quad (11)$$

Strictly speaking, that is all there is to the formulation of the algorithm; but we want to be able to evaluate the mean-square error in the results when various forms of filtering are applied to the image. For this we introduce the Fourier-transform and inverse Fourier-transform operators.

The Fourier transform process can be represented as a multiplication by a matrix, as can its inverse; the two matrices are denoted \mathbf{F} and \mathbf{F}_{inv} . We use the notation $\mathbf{C}(\kappa)$ to represent a low-pass filter. The matrix acts on the Fourier transform to set all components for spatial frequencies above κ to zero. Thus the recovered image, processed to suppress all frequencies above κ , would be evaluated with the expression $\mathbf{F}_{\text{inv}} \mathbf{C}(\kappa) \mathbf{F} \mathbf{A} \mathbf{m}$. Of particular interest to us is the column vector $\mathbf{e}(\kappa)$, representing the error in the resultant image. Its value is given by the equation

$$\mathbf{e}(\kappa) = \mathbf{F}_{\text{inv}} \mathbf{C}(\kappa) \mathbf{F} \mathbf{A} \mathbf{m} - \mathbf{F}_{\text{inv}} \mathbf{C}(\kappa) \mathbf{F} \mathbf{o}. \quad (12)$$

Noting, from consideration of Eq. (11), that $\mathbf{A} \mathbf{S} = \mathbf{I}$, where \mathbf{I} denotes the identity matrix, and making use of Eq. (9) we see that Eq. (12) can be rewritten as

$$\mathbf{e}(\kappa) = \mathbf{F}_{\text{inv}} \mathbf{C}(\kappa) \mathbf{F} \mathbf{A} \mathbf{n}. \quad (13)$$

Using the fact that the inner product is equal to the trace of the outer product, we can show that the mean-square error,

$$\mathcal{E}(\kappa) = \langle [\mathbf{e}(\kappa)]^T \mathbf{e}(\kappa) \rangle, \quad (14)$$

can be shown to have a value given by the equation

$$\mathcal{E}(\kappa) = \text{tr}\{\mathbf{F}_{\text{inv}}\mathbf{C}(\kappa)\mathbf{F}\mathbf{A}[\mathbf{F}_{\text{inv}}\mathbf{C}(\kappa)\mathbf{F}\mathbf{A}]^T\}\sigma^2, \quad (15)$$

where σ^2 is the mean-square raw measurement noise for each element of the measurement array \mathbf{m} . The notation $\text{tr}\{\dots\}$ is used to denote the trace operation.

Correspondingly, for the mean-square error associated with the estimated value of the Fourier-transform component for spatial frequency κ , we can write that

$$E_e(\kappa) = \text{tr}\{\mathbf{D}(\kappa)\mathbf{F}\mathbf{A}[\mathbf{D}(\kappa)\mathbf{F}\mathbf{A}]^T\}\sigma^2, \quad (16)$$

where matrix $\mathbf{D}(\kappa)$ is a diagonal matrix with all of its diagonal elements equal to zero except for one that is equal to unity. The matrix serves to select the single element in the Fourier transform corresponding to the spatial frequency κ , setting all the other components to zero.

With the use of Eqs. (15) and (16), it has been possible to evaluate the errors associated with cases covering a wide range of problem parameters. The computations proceeded sufficiently rapidly that we were able to explore problem parameter space quite extensively. (In fact it was in this exploration that we first became aware of the significance of the size of the object.) A small sample of the extensive results obtained are presented in Figs. 3 and 4.

All the results that are presented here were developed for a one-dimensional object, image, and imaging system array, with a pixel size of $1/4\lambda/L$. As already noted, all the results presented here were developed for the imaging array defined by Figs. 1 and 2. Using the point-spread function calculated for this array, we formulated the \mathbf{S} matrix and thereupon were able to proceed directly with the calculation of \mathbf{A} from Eq. (11) and from that were able to proceed with the calculation of $\mathcal{E}(\kappa)$ and $E_e(\kappa)$, using Eqs. (15) and (16), respectively. In Fig. 3 is shown $E_e(\kappa)$ normalized by being divided by $E_x(\kappa)$, where $E_x(\kappa)$ denotes the mean-square value of the spatial-frequency components of the object being viewed, i.e., the true value (free of noise or imaging effects). [Results are presented with $E_x(\kappa)$ treated as a constant.] Results are shown in Fig. 3 for object size equal to $3\lambda/L$, $6\lambda/L$, $9\lambda/L$, ..., $21\lambda/L$. (Recall that L denotes the overall length of the array and λ denotes the wavelength at which the array is operating.) It is obvious that for object sizes as high as $15\lambda/L$ (but not $18\lambda/L$ or $21\lambda/L$) there is no problem in filling in the missing spatial-frequency content up to a spatial frequency of $\sim 13/17\lambda/L$. I interpret the occurrence of this transition for an object size between $15\lambda/L$ and $18\lambda/L$ as being related to the fact that the width of the smaller holes in the array's OTF pattern (see Fig. 2) is equal to $1/17L/\lambda$, with the inverse, $17\lambda/L$, being just at the transition. In considering Fig. 3 further (in preparation for the discussion of superresolution in Section 5), also note that for the smallest object, with a size $3\lambda/L$, the values of spatial-frequency components are being estimated with tolerable precision for spatial frequencies definitely beyond the array's (diffraction-limited) cutoff frequency of L/λ , out to frequencies $\sim 1.15L/\lambda$. Note that this increment of spatial-frequency coverage, an increment equal to

$0.15L/\lambda$, is nearly equal to one half of the inverse of the object's size, which is $1/2(3\lambda/L)^{-1} \approx 0.1667L/\lambda$.

The mean-square error in the estimated pixel value of the object was also calculated when the estimated image was filtered so as to exclude spatial-frequency components for spatial frequencies above a cutoff frequency of κ_0 . [In accordance with Eq. (12) the error in the filtered image was taken to correspond to the deviation from the correspondingly filtered version of the original object.] Figure 4 shows this error in normalized form, i.e., an $\mathcal{E}(\kappa_0)$ divided by the mean-square variability (i.e., the mean-square variations from realization to realization) of the true pixel values in the spatially filtered object. For obvious reasons this normalized

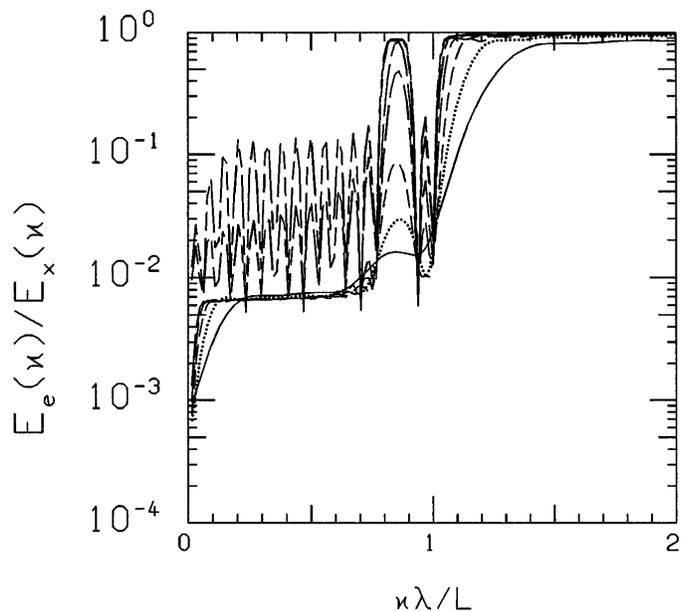


Fig. 3. Spatial-frequency-component estimation error.

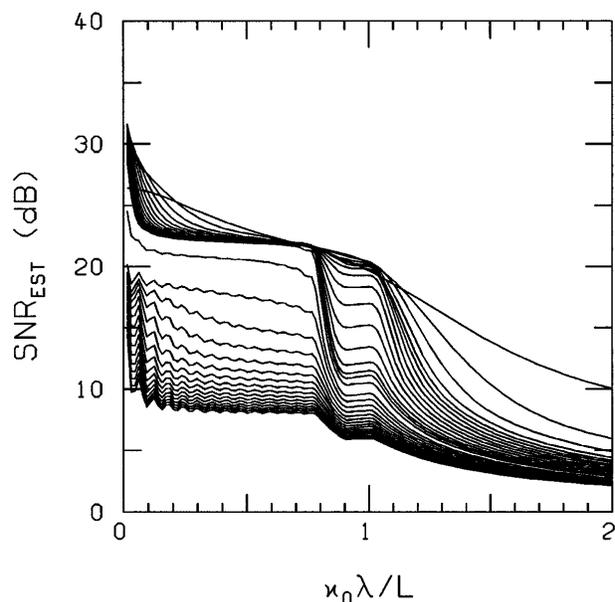


Fig. 4. Enhanced and filtered image signal-to-noise ratio as a function of spatial-frequency cutoff.

form of $\mathcal{E}(\kappa_0)$ is referred to as the signal-to-noise ratio, SNR_{EST} . Results are shown for SNR_{EST} as a function of the filter cutoff frequency, κ_0 , for object sizes of $1\lambda/L$, $2\lambda/L$, $3\lambda/L$, ..., $32\lambda/L$. The break in SNR_{EST} values just past $\kappa_0 = {}^{13/17}\lambda/L$ (for all except perhaps the six smallest objects, for which the sizes are $1\lambda/L$ to $6\lambda/L$) is apparent. The relationship between the size of these objects and the value of the inverse of the larger hole in the OTF, a hole size of $\sim {}^{3/17}L/\lambda = (5.667\lambda/L)^{-1}$, is clear.

We also note from Fig. 4 that for the 16 largest targets, the sizes between $17\lambda/L$ and $32\lambda/L$, there is a noisiness problem for all filter cutoff frequencies. This noisiness can be associated with the fact that the size of the smaller holes in the OTF is ${}^{1/17}L/\lambda$. Here again we see the close tie between the inverse of the size of the object being imaged and the width of the holes in the OTF shown in Fig. 2.

In all cases we run into an algorithm performance problem when the object size is greater than the inverse of the size of the hole in the OTF—or, put somewhat differently, when one half of the inverse of the size of the object is less than the distance (in frequency space) from the frequency at the center of the hole to the frequency at the edge of the hole. The transition from no problem to problem is apparently soft, but there definitely is a transition.

4. CLEAN SIMULATION RESULTS

Once we recognized the significance of the object size and had results showing what performance could be obtained with a simple minimum-mean-square-discrepancy algorithm, we undertook the development of results, using a CLEAN-type algorithm to see if the same sort of size dependence was present. Writing a computer program to simulate CLEAN was quite straightforward. The program steps are as follows:

1. The CLEAN simulation program starts with a zero value for each pixel or element of the object estimate, $\hat{\mathbf{o}}$.
2. The initial value of \mathbf{m}_{res} is taken as \mathbf{m} .
3. Using one half of half the value of the current brightest pixel in \mathbf{m}_{res} , the program adds that amount to the corresponding pixel of $\hat{\mathbf{o}}$.
4. The CLEAN simulation program then reduces the value of every pixel in \mathbf{m}_{res} by an amount calculated to be equal to the increment that was applied to $\hat{\mathbf{o}}$ times the appropriate values from the instrument's point-spread function, \mathbf{S} , but it never reduces the value of any pixel in \mathbf{m}_{res} to less than zero.

The program repeats this process over and over, each time checking to ensure that the sum of the squares of the residual values in \mathbf{m}_{res} is less than that for the previous iteration. When the reduction stops, the iterations are terminated and the value of $\hat{\mathbf{o}}$ is taken as CLEAN's estimate of the object.

Because this algorithm is iterative and because it had to run many times with the same set of problem parameter values to yield statistically reliable Monte Carlo results, the range of problem parameters that could be studied was limited. Accordingly, the minimum-mean-square-discrepancy algorithm was used for general exploration and then, having found a particularly inter-

esting and explanation-laden case, we used the CLEAN simulation to develop results for comparison. Figure 5 shows SNR_{EST} results for the case in which the low-pass filter was set to have a cutoff at a spatial frequency of $\kappa_0 = (13/17)\lambda/L$. The results are shown as a function of object size. Note that with this filter cutoff setting the only relevant holes in the OTF (see Fig. 2) are the smaller holes, for which the width is ${}^{1/17}L/\lambda$. It is clear that as long as the object size is less than the inverse of the size of the relevant holes in the OTF, i.e., the object's size is no larger than $({}^{1/17}L/\lambda)^{-1} = 17\lambda/L$, we have no problem in filling in the missing spatial frequencies. For larger-sized objects there is a significant noise penalty.

It is also of great significance that the results shown in Fig. 5 are in good quantitative agreement with corresponding results obtained with the minimum-mean-square-discrepancy, linear algorithm developed in Section 3. This, along with the other aspects of Fig. 5 noted above, confirms the supposition that, in view of the intention to apply it to large-sized image fields, the CLEAN algorithm may be considered to be a practical but also a somewhat covert implementation of the minimum-mean-square-discrepancy approach to the problem, or at least a near approximation to that approach. Its performance appears to depend on object size in the same way. It seems that the ability of CLEAN to develop estimates for missing spatial-frequency components is dependent on the correlation between components whose spatial frequencies are close together.

To the extent that the ability to develop an estimate of the component's value when the imaging system's OTF is equal to zero for the component's spatial frequency can be considered to be a manifestation of superresolution, it would appear that superresolution is critically dependent on object size. This matter is discussed further in Section 5.

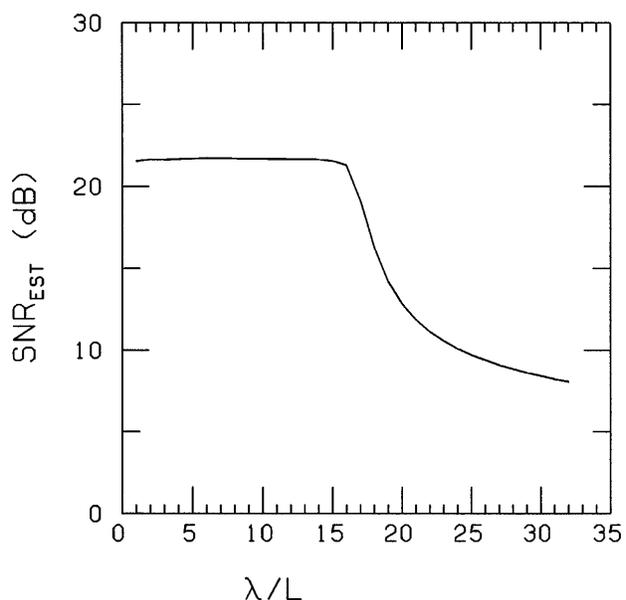


Fig. 5. Enhanced image signal-to-noise ratio as a function of object size.

5. SUPERRESOLUTION AND THE RAYLEIGH LIMIT

I believe that the results that we have obtained in trying to understand how the CLEAN algorithm can achieve the performance that it does, in particular the observation that it (and also the minimum-mean-square-discrepancy algorithm) can develop an estimate of the value of a spatial-frequency component for a frequency for which the imaging system's OTF is equal to zero without introducing any significant signal-to-noise-ratio requirement, are relevant to the development of an understanding of the subject of superresolution. In fact, I would argue that what the CLEAN algorithm does *is* an example of superresolution. I believe that what has been found concerning the importance of the size of the object being imaged and the correlation between spatial-frequency components of the object not only helps in understanding why CLEAN is able to perform as well as it does and what the algorithm's limitations are but also sheds comparable light on the subject of superresolution: why superresolution is possible and what its limitations are. In the following discussion my intention is to apply the results developed above to the matter of superresolution. The first matter to be considered is the relationship of superresolution and the Rayleigh limit to the definition of superresolution.

The concept of superresolution has an allure that continues to attract researchers with its promise of achieving better resolution than the limit that is nominally associated with an instrument,⁵⁻¹⁴ i.e., better than that given by the Rayleigh criterion (or Rayleigh limit). This limit is conventionally expressed by the statement that for imaging at a wavelength λ with an aperture whose maximum dimension is D , two point sources cannot be resolved if their apparent angular separation is less than λ/D . [In some instances,¹⁴ for microscopy in particular, the instrument's limit is expressed more directly as a distance λ (or $\lambda/2$) on the plane of the object being viewed.] It takes only a small amount of experience in image processing to develop a sense that this is not a firm limit, and it is quite likely because of this that the concept of superresolution seems so promising and plausible to many scientists.

In part, the apparent possibility of achieving resolution better than that which the Rayleigh limit seems to allow may be attributed to the fact that the Rayleigh limit is not fully defined; we do not have a firm criterion defining what it means to resolve two points—a criterion providing a basis for saying in any particular instance that the two points are or are not being resolved. For some researchers⁹ the limit is softly enough defined that it is considered to be breached simply by achievement of point-source position measurement precision significantly finer than λ/D . The vagueness as to what constitutes superresolution is characteristic when the matter of resolution in the spatial domain is being considered.

Generally, however, the concept of resolution in imaging is taken to refer to the quality, the sharpness, the definition, or simply the resolution (in the spatial-frequency-content sense) of the final image, and we ought to consider the subject of superresolution in a manner related to the spatial-frequency-content sense of the concept of resolution. In the spatial-frequency domain there is

an absolute cutoff in response. It is possible to formulate an unambiguous and unyielding definition of what constitutes superresolution based on the cutoff, a definition with an obvious (though soft) relationship to the Rayleigh limit.

I argue that we should consider the matter in the spatial-frequency domain, that we should define superresolution as the ability to form an image with meaningful content at spatial frequencies for which the imaging system has an OTF equal to zero. (The term meaningful is used here to imply that for the spatial frequency in question there is a positive correlation between the complex amplitude of that component of the image and the complex amplitude of the corresponding component of the object.) The fact that the OTF of any imaging system does not just become very small but rather goes fully to zero means that simple high-pass image filtering cannot constitute superresolution. This means that one would expect superresolution to be difficult to achieve. It is not just a matter of relying on a high signal-to-noise ratio to allow a source's position to be determined precisely or a matter of simple filtering to enhance components that the imaging process have attenuated. Superresolution calls for the recovery of spatial-frequency components that have been completely eliminated by the image formation process. But it is also reasonable to ask that this definition of superresolution have a relationship to the Rayleigh limit.

This definition of superresolution does incorporate the basic sense of the Rayleigh criterion of resolving two point sources. To see that this is true it is sufficient that one consider the case of an object (to be imaged) that consists of two points that are closer together than λ/D . The spatial-frequency components of the pattern defined by this two-point object, which are the basic carriers of the information that there are two distinct points and not just a single elongated element making up the object, are components whose spatial frequencies are greater than D/λ . But we know that such frequency components are missing entirely from any image produced by an imaging system with aperture size D operating at wavelength λ . The nature of diffraction-limited imaging is such that the components that one presumably needs to be able to tell that the object consists of two points—the term “tell” conveying, in the Rayleigh sense, the concept of resolving the two points—are missing from the image. If we can tell that there are two points (i.e., if we can estimate the two points' parameters), we may consider that we have measured the missing spatial-frequency components, and, conversely, if we can estimate the missing spatial-frequency components, we may consider that we can tell that there are two points there and that we have measured their parameters. It is in this sense that we take the concept of superresolution in the Rayleigh-limit sense as implying that it allows such missing spatial-frequency components to be meaningfully estimated; if we say that we can resolve the two points, we are in effect saying that we can estimate the value of spatial-frequency components for which the imaging system's OTF is equal to zero. We generalize from the two-point problem to consideration of an arbitrary object, requiring of superresolution that it recover in the final image information about spatial-frequency components for which the instrument's OTF is equal to zero.

The definition of superresolution as the recovery of missing spatial-frequency components has the dual virtues of subsuming what we may think of as the concept of the Rayleigh-resolution criterion and also of being quantitatively testable and evaluatable. It transforms the loosely defined concept of being able (or not being able) to resolve into the absolute criterion that for certain (higher) spatial frequencies, the instrument's OTF being equal to zero, the amplitude of that spatial-frequency component can (or cannot) be usefully estimated. It casts superresolution as the task of getting around that limitation and recovering information about the amplitude of spatial frequencies for which the instrument's OTF is equal to zero—without any *a priori* information about the object being imaged other than its size. This certainly seems to be a formidable task (or would seem to be so except in light of what we have learned in studying the CLEAN algorithm).

Note that this definition of superresolution brings with it coverage of a case that might not ordinarily be thought of as being related to superresolution. For an imaging array there can be spatial frequencies whose magnitude is smaller than the size of the array, L , divided by the wavelength, λ , for which spatial frequencies the array's OTF is equal to zero. Should processing the image so as to recover those missing spatial frequencies be considered as constituting superresolution, as the definition of superresolution given above indicates? I argue that it should be so considered, that developing information about spatial-frequency components regarding which the image system reports nothing is an equally remarkable task whether the spatial frequency is greater than or less than L/λ . Moreover, the results that have been developed (see the discussion associated with Fig. 3) indicate that the same achievability criterion applies to our ability to fill in missing spatial-frequency content whether the frequency is above or below the L/λ cutoff frequency. Since in both cases the accomplishment of the task is equally remarkable, and since the same achievability criterion appears to apply in both cases, I see no reason that the term superresolution should not be considered to apply equally well whether the spatial frequency for the component whose missing amplitude is being estimated lies above or below the array's cutoff frequency, L/λ .

The findings presented here indicate that the criterion for achievability of superresolution performance in any situation appears to be that the missing spatial-frequency component that is being estimated should be a distance (in the spatial-frequency domain) no greater than one half of the inverse of the size of the object from a component that is measured (i.e., from a component for which the imaging system's OTF is not equal to zero; otherwise, at the input to the algorithm a considerably higher signal-to-noise ratio will be required. This condition applies whether we consider filling in holes in the OTF or pushing beyond the outer bounds of the OTF).

An interesting point regarding the problem of pushing beyond the outer bounds of the OTF has to do with the product of object size and the expanded spatial-frequency bandwidth, which product is denoted by \mathcal{B} , for space-bandwidth product. This product is equal to the number of cycles of effective image bandwidth across

the diameter of the object. Of interest is what could be achieved if we were to utilize the superresolution algorithm but were subject to the limitation that no excess signal-to-noise ratio is available. If the object's size is $N(\lambda/D)$, then according to the findings presented here the spatial-frequency range could be extended by use of superresolution (from D/λ to $(1 + 1/2N)(D/\lambda)$, and accordingly the product's value would be extended (from $\mathcal{B} = N$) to $\mathcal{B} = N + 1/2$. Considering that \mathcal{B} represents the space-bandwidth product of the final image (in essence one half of the number of useful pixels across the image of the object), it would seem that little increase in image enhancement or resolution would be available through the use of superresolution without production of some excess signal-to-noise ratio that we can afford to let the superresolution process use up. This suggests that an appropriate step in the development of superresolution would be to extend the computational approach of Section 3 to form an estimate of how achievable resolution enhancement depends on the available signal-to-noise ratio that we can afford to have used up by the superresolution algorithm.

ACKNOWLEDGMENT

I acknowledge the assistance of Douglass Sherwood in the preparation of the computer programs used in this work.

REFERENCES

1. A. R. Thompson, J. M. Moran, and G. W. Swenson, Jr., *Interferometry and Synthesis in Radio Astronomy* (Wiley, New York, 1986).
2. J. A. Roberts, ed., *Indirect Imaging*, Proceedings of an International Symposium, Sydney Australia, 30 August to 2 September 1983 (Cambridge U. Press, London, 1984).
3. J. L. Harris, "Diffraction and resolving power," *J. Opt. Soc. Am.* **54**, 931–936 (1964).
4. C. K. Rushforth and R. W. Harris, "Restoration, resolution, and noise," *J. Opt. Soc. Am.* **58**, 539–545 (1968).
5. W. Lukosz, "Optical systems with resolving power exceeding the classical limit," *J. Opt. Soc. Am.* **56**, 1463–1472 (1966).
6. C. W. Barnes, "Object restoration in a diffraction-limited imaging system," *J. Opt. Soc. Am.* **56**, 575–578 (1966).
7. B. R. Frieden, "Band-unlimited reconstruction of optical objects and spectra," *J. Opt. Soc. Am.* **57**, 1013–1019 (1967).
8. B. R. Frieden, "Evaluation, design and extrapolation methods for optical signals, based on use of the prolate functions," in *Progress in Optics*, E. Wolf, ed. (North-Holland, Amsterdam, 1971), Vol. 19, pp. 311–407.
9. J. P. Fillard, H. M'timet, J. M. Lussert, and M. Castagne, "Computer simulation of super-resolution point source image detection," *Opt. Eng.* **32**, 2936–2944 (1993).
10. E. Noel, R. R. Khan, and H. S. Dhadwal, "Optical implementation of a regularized Gerchberg iterative algorithm for super-resolution," *Opt. Eng.* **32**, 2866–2871 (1993).
11. J. G. Walker, E. R. Pike, R. E. Davies, M. R. Young, G. J. Brakenhoff, and M. Bertero, "Superresolving scanning optical microscopy using holographic optical processing," *J. Opt. Soc. Am. A* **10**, 59–64 (1993).
12. E. Scales and G. A. Viano, "Resolving power and information theory in signal recovery," *J. Opt. Soc. Am. A* **10**, 991–996 (1993).
13. P. C. Sun and E. N. Leith, "Superresolution by spatial-temporal encoding methods," *Appl. Opt.* **31**, 4857–4862 (1992).
14. J. M. Vigoureux and D. Courjon, "Detection of nonradiative fields in light of the Heisenberg uncertainty principle and the Rayleigh criterion," *Appl. Opt.* **31**, 3170–3177 (1992).