LEGIBILITY NOTICE

A major purpose of the Technical Information Center is to provide the broadest dissemination possible of information contained in DOE's Research and Development Reports to business, industry, the academic community, and federal, state and local governments.

Although a small portion of this report is not reproducible, it is being made available to expedite the availability of information on the research discussed herein.

TITLE COMPARISON OF THE ALGEBRAIC RECONSTRUCTION TECHNIQUE WITH
 THE MAXIMUM ENTROPY RECONSTRUCTION TECHNIQUE FOR A VARIETY
 OF DETECTION TASKS

AUTHOR(S) K. M. Hanson
 K. J. Myers

SUBMITTED TO Proceeding SPIE, Volume 1231
 Medical Imaging IV Meeting
 Newport Beach, CA
 February 4-9, 1990

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

By acceptance of this article, the publisher recognizes that the U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes.

The Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy.

Los Alamos

Los Alamos National Laboratory
Los Alamos, New Mexico 87545

MASTER

Comparison of the algebraic reconstruction technique with the maximum entropy reconstruction technique for a variety of detection tasks

K.J. Myers¹ and K.M. Hanson²

¹Center for Devices and Radiological Health, FDA
HFZ-142, 12720 Twinbrook Parkway
Rockville, MD 20857

²Los Alamos National Laboratory, MS P940
Los Alamos, NM 87545

ABSTRACT

A method for comparing reconstruction algorithms is presented based on the ability to perform certain detection tasks on the resulting images. The reconstruction algorithms compared are the algebraic reconstruction technique (ART) and the maximum entropy reconstruction method (MaxEnt). Task performance is assessed through a Monte Carlo simulation of the complete imaging process, including the generation of a set of object scenes, followed by data-taking, reconstruction, and performance of the specified task by a machine observer. For these detection tasks the figure of merit used for comparison is the detectability index, d' . When each algorithm is run with approximately optimized parameters, these studies find comparable values for d' .

1. INTRODUCTION

The overall purpose of a medical imaging system is usually to provide visual information for interpretation by a radiologist or other trained observer. Assessment of the quality of images obtained from a medical imaging system has been a long-standing problem. The merit of an image cannot be determined simply by subjective visual inspection. Rather, image quality must be evaluated objectively based on how well the image provides the information required by the observer to perform a specific task. If a scalar figure of merit can be determined that represents the usefulness of the image for the specified task, many aspects of the imaging chain can be optimized by adjusting the system parameters to increase that measure. Or, the figure of merit can be used to compare two different imaging systems. In this study we shall consider one aspect of the imaging chain, the reconstruction procedure, and consider both the optimization of a single reconstruction algorithm and the comparison of two different (and each optimized) reconstruction algorithms based on the usefulness of the images they provide for performing certain detection tasks.

2. CALCULATION OF TASK PERFORMANCE

In order to evaluate the usefulness of an image we need to first specify some task and then determine the ability of some observer to perform the task using the image. Psychophysical studies are one avenue for doing this; a set of real or simulated images are presented to trained human observers and their ability to perform a stated task is measured. Unfortunately, psychophysical studies can be very time consuming to perform. Optimization of imaging systems based on psychophysical testing would be cumbersome, since many images and observers are required to evaluate a single set of imaging system parameters. The whole psychophysical study would have to be performed many, many times with great resolution in order to determine the optimal system configuration over several system parameters.

An alternative strategy is to calculate the performance of a model observer. There is a considerable literature on the ideal or Bayesian observer, defined as one who has full knowledge of all relevant statistical properties of the images and of the task at hand. The performance of the Bayesian observer has been calculated analytically or numerically for several simple detection and discrimination tasks where the signal

and background are completely specified and the observer's performance is limited by the randomness in the data resulting from the quantum fluctuations in the incoming radiation stream.^{1,2} These fluctuations result in a noisy image that changes for repeated imaging trials even when the object is the same each time. More recently, performance calculations for two model observers (both sub-optimal) were derived for the case where the signal was known exactly but the observer was limited by both image noise and a random, inhomogeneous background.³ The calculation of ideal-observer performance measures in cases where there is some randomness in the parameters defining the signal to be detected or discriminated is the subject of ongoing research, and has been presented at this and other recent SPIE Medical Imaging meetings.⁴⁻⁷ Still, these investigations do not address the degradation in task performance that can occur when the reconstructed image contains artifacts. We define artifacts as those deviations in the reconstructed image from the original object that are present in the absence of measurement noise. Artifacts are the result of inadequate sampling of the object during the measurement process and do not change when the same object is imaged repeatedly. Since artifacts depend on the object in a complicated fashion, we choose a third method for image evaluation based on Monte Carlo simulations of the complete imaging process.

2.1 Monte Carlo Method

The Monte Carlo method is favored for our purpose because we can define a class of objects and thereby obtain an ensemble of images, each one a single realization of the noise and artifacts that result from the imaging process. By performing the visual task on the ensemble of images we obtain a statistically meaningful average of the response of each algorithm to the object class and task. This Monte Carlo method has been used by Hanson to optimize the algebraic reconstruction technique (ART) for a wide variety of detection and estimation tasks.⁸⁻¹¹ The method has also been used recently to optimize the maximum a posteriori reconstruction method.¹²

Before the Monte Carlo method can proceed, the problem must be fully specified. In general, the following steps must be performed:

- 1) First, the object class or classes must be specified, including any variability in the signal and background parameters that are to be modeled.
- 2) The measurement geometry must be specified, including all sources of noise and blur. If the acquired data are to be reconstructed or otherwise processed in some manner, the processing algorithm must be specified.
- 3) The task to be performed must be clearly defined.
- 4) The method of task performance must be designated. The method should be appropriate for the intended application. If a human viewer is to be the image interpreter for the final system, a model of the human observer can be invoked. If a machine reader is to analyze the final images, then a machine can perform the task in the Monte Carlo simulation. Or, the Bayesian observer can be used to determine the best possible task performance given the image and any available information about the prior probabilities.
- 5) A figure of merit that quantifies the usefulness of the images for performing the specified task must be chosen.

Once the steps above have been accomplished, the Monte Carlo simulation proceeds as follows:

- 1) One realization of an object is generated.
- 2) The object is used to generate a detected data set by simulating the imaging process. Any required post-processing or reconstruction of the raw data is performed.
- 3) The task is performed by the simulated observer according to the prescribed rule.
- 4) Steps 1) through 3) are repeated a sufficient number of times so that an accurate estimate of the statistics of the task performance may be obtained.
- 5) Finally, the figure of merit is determined that summarizes the usefulness of the imaging system based on the task performance statistics for the ensemble of objects.

2.2 Example - Comparison of ART and MaxEnt

The specific problem we shall address is the comparison of two reconstruction algorithms, the algebraic reconstruction technique (ART) and the maximum entropy reconstruction method (MaxEnt). The object class shall consist of a set of 10 scenes, each containing many randomly placed, non-overlapping discs on a zero background. Specifically, 10 low-contrast (amplitude=0.1) discs are present in each scene. To probe the effect of object-dependent artifacts, 10 high-contrast discs (amplitude=1.0) are also randomly placed in each scene. The high- and low-contrast discs are all 8 pixels in diameter in an object of 128 pixels in diameter. One sample object from the ensemble is shown in Figure 1. The task shall be the detection of the low-contrast discs, with amplitude, size, and location known. The measurements consist of a number of parallel projections equally spaced over 180° , each containing 128 samples. We shall consider two data-taking situations. The first is the situation where there are many views, 100 in total, and each of the 100 projections contains additive, zero-mean Gaussian noise with rms=8. This first case is intended to investigate the noise-limited regime of the reconstruction problem. The second case we shall consider is the limited-view situation, where the number of projections is small, 8 in all, and no noise is added to the detected data. This scenario is meant to investigate the artifact-limited domain of the reconstruction problem.

The simulated projection data become the input to the reconstruction algorithms, whose job it is to find an estimate of the original object on which the detection task can be performed. The details of the reconstruction algorithms are given in sections below. Both reconstruction algorithms employ a pre-smoothing step that has the effect of mildly blurring the projection data. The pre-smoothing filter is a 3-pixel-wide triangular window that reduces the rms noise in the data by a factor of 0.494. Such a pre-smoothing step is recommended for these algorithms on tomographic problems to avoid over-fitting the high-frequency components in the input data that can result in ringing in the reconstruction. In the results sections that follow, the stated rms noise in the data is the value before pre-smoothing.

2.3 Figures of Merit for Task Performance

To evaluate and compare the reconstructions, a figure of merit summarizing the detectability of the low-contrast discs in the scenes must be obtained. The task we shall consider is the binary detection problem, that is, is the disc present or absent? The observer or decision-maker will be the non-prewhitening matched filter. We choose this model observer because the Bayesian observer requires full knowledge of the probability density functions on the data when the signal is present and absent. These distributions are difficult to obtain because of the scene-dependent nature of the artifacts in the images. The non-prewhitening matched filter is essentially a handicapped Bayesian observer in that it uses all information about the signal parameters perfectly, but does not attempt to correct for any correlations in the background fluctuations from either artifacts or noise post-processing. This model has been found to correlate well with human observer performance in a number of psychophysical investigations.^{13,14}

When the signal and background are assumed to be completely known, the non-prewhitening matched filter forms a template that is a simple disc of radius matched to the low-contrast disc, and counts up all the activity in the region it superimposes. The output of that operation becomes the observer's decision variable. The observer declares the signal to be present, a positive response, if the decision variable is above his decision threshold. By applying this strategy to many locations where the signal is known to be present, and again where the signal is known to be absent, we can derive histograms for the observer's decision variable (plots of the frequency of occurrence of the decision variable) under both the signal-present and signal-absent conditions. The traditional receiver operating characteristic (ROC) curve is obtained by plotting the fraction of true-positive responses versus the fraction of false-positive responses from the histograms as the observer's decision threshold is varied. In our studies the histograms are generated by reconstructing 10 different object scenes, each containing 10 low-contrast discs that are used to form the signal-present decision-variable histogram and 30 locations that are free of the disc signals and are therefore used to form the signal-absent decision-variable histogram. The further apart the histograms, the better the observer can correctly distinguish signal-present cases from those where the signal is absent.

The degree of separation of the two histograms can be characterized by the detectability index d' , given by¹⁵

$$d' = \frac{\bar{\psi}_1 - \bar{\psi}_0}{\sqrt{\frac{\sigma_1^2 + \sigma_0^2}{2}}} \quad (1)$$

where $\bar{\psi}_1$ and σ_1 are the mean and rms deviation of the frequency distribution when the signal is present, and $\bar{\psi}_0$ and σ_0 are the mean and rms deviation of the frequency distribution when the signal is absent. The detectability index is sometimes called the observer's signal-to-noise ratio. This is a meaningful figure of merit provided the distributions on the decision variable under each hypothesis are Gaussian. However, one should be aware that improper conclusions regarding the ranking of observer performances can be obtained when the detectability index is used to describe very non-Gaussian decision-variable frequency distributions.⁷

We could have chosen to plot the entire ROC curve and use the area under the curve as the figure of merit instead of the detectability index d' . Or, an effective detectability index d_A that can be derived from the area might have been used.^{15,16} We choose to use the detectability index d' as our figure of merit because it is the most accurate representation of the overlap of the two histograms, especially for the high d' values we encounter in some of our examples, provided they meet the Gaussian assumption.^{7,11} Hanson has found good agreement between d' and d_A for tasks similar to those considered here.¹¹

A number of other summary measures have been suggested as figures of merit for image quality. One of the best-known of these is the root mean square difference between a reconstructed or estimated object and the actual object, called the L2 norm. The problem with this metric is that the relationship between the L2 norm and the usefulness of the image for task performance is unclear. Measures based on the agreement between the actual data and the estimated data that would be derived from the estimated object, such as the rms residual, can be ill-conditioned or ill-posed. In the approach to system (specifically algorithm) evaluation presented here, we are using a figure of merit that is fundamentally related to the ability of an observer to use the image to perform a specified task.

3. ART

The algebraic reconstruction technique¹⁷ is an iterative algorithm that reconstructs a function from its projections and has been shown to be useful in situations where there are a limited number of projections. Optimization of the algorithm for a variety of visual tasks has been presented by Hanson.⁸⁻¹¹ Results of the optimization of ART for a number of detection tasks shall be summarized briefly below to allow comparison with the MaxEnt algorithm for the same tasks.

3.1 Theory

We assume that N individual projection measurements are acquired of the unknown object $\mathbf{f}$, with the boldface character denoting a vector. The data vector for the i^{th} projection can be written as

$$\mathbf{g}_i = \mathbf{H}_i \mathbf{f} + \mathbf{n}_i \quad , \quad i=1,\dots,N, \quad (2)$$

where $\mathbf{H}_i$ is the corresponding row of the measurement matrix and $\mathbf{n}_i$ is the noise vector for that projection.

The ART algorithm proceeds by forming an initial guess, for example, $\hat{f}^0=0$. Then for each iteration the estimate is updated by iterating on the individual measurements as follows:

$$\hat{f}^{k+1} = \hat{f}^k + \lambda^k H_i^T \left[\frac{g_i - H_i \hat{f}^k}{H_i^T H_i} \right], \quad i=k \cdot \text{mod}(N)+1, \quad (3)$$

where $\hat{f}^k$ is the k^{th} estimate of the object f , λ^k is the relaxation factor for the k^{th} update, and the superscript "T" represents the transpose of a vector or matrix. Constrained ART uses the same iteration approach represented by Eq. (3), with the additional rule that any negative object estimate be set to zero before proceeding to the next update. The relaxation factor is given by

$$\lambda^K = \lambda_0 (r_\lambda)^{K-1}, \quad (4)$$

where the iteration number $K=\text{int}(k/N)$ represents the number of passes through all N projections. Hanson has investigated the optimization of the relaxation parameters λ_0 and r_λ for several detection and discrimination tasks.⁹⁻¹¹

3.2 Results

Table 1 summarizes the results for the optimization of both the unconstrained and constrained ART algorithm for the detection task. From the top half of the table, we see that nominal values for the relaxation parameters in the 100-view case lead to a d' value of 1.99 that degrades slightly with the application of the positivity constraint for 10 iterations of ART. Neither optimization of the relaxation parameters, nor increasing the number of iterations to 100 has any large effect on the detectability index for this case. Figure 2a shows a reconstruction of the object in Figure 1, achieved via the constrained ART algorithm from 100 noisy projections after 100 iterations, corresponding to the $d'=1.84$ row of Table 1.

The second half of Table 1 reviews the results for the optimization of the ART algorithm in the 8-view case. This imaging geometry is found to yield d' values that are greatly improved through optimization of the relaxation parameters and through the application of the positivity constraint for 10 iterations of ART. Constrained ART with relaxation parameters optimized for 100 iterations yields a detectability index of 11.69. Figure 3a is an image of the source shown in Figure 1, reconstructed from 8 noiseless views after 100 iterations of optimized, constrained ART.

4. MAXENT

The maximum entropy algorithm employed here is the "historic" MaxEnt algorithm described by Gull and Skilling¹⁸ for the reconstruction of positive, additive images. A commercial software package, entitled MEMSYS 2(c), was used to obtain the MaxEnt reconstructions.¹⁹

4.1 Theory

We again assume that the data are a set of projection vectors as described by Eq. (2) above. The entropy of the reconstruction is defined to be

$$S(\hat{f}) = - \sum_j p_j \log p_j \quad p_j = \hat{f}_j / \sum \hat{f}_j \quad (5)$$

In the case where there is an underlying model for the object class, the entropy can be rewritten as

$$S(\hat{f}) = \sum_j [\hat{f}_j - m_j - \hat{f}_j \log(\hat{f}_j/m_j)] \quad (6)$$

where m_j is the value of the model in the j^{th} object pixel. The global maximum of $S(\hat{f})$ occurs when $\hat{f}_j = m_j$ for all j , where $S(\hat{f}) = 0$. Hence the entropy term measures the deviation of the object estimate from the assumed model. Since the randomly placed discs in the object scenes in the simulations here are equally likely to occur in any location, the underlying model is a uniform grey level set equal to the average pixel brightness in the object. Note also that the logarithmic term in Eqs. (5) and (6) has the effect of imposing a positivity constraint on the reconstruction values.

The misfit between the reconstruction and the detected data values is measured in the data domain by the chi-squared statistic:

$$\chi^2(\hat{f}) = \sum_i \left[\frac{\left[g_i - \sum_j H_{ij} \hat{f}_j \right]^2}{\sigma^2} \right] \quad (7)$$

where σ^2 specifies the expected variance in the data. Here we assume the noise to be independent of location, so it appears as a constant independent of the summation index. The chi-squared value is a meaningful measure of the agreement between the data and the reconstruction as long as the noise is Gaussian and uncorrelated. Note that the χ^2 value is related to the rms residual of the reconstruction by

$$\text{rms residual} = \left[\chi^2(\hat{f}) \cdot \frac{\sigma^2}{\text{No. of measurements}} \right]^{1/2} \quad (8)$$

One obtains the maximum entropy solution by maximizing $S(\hat{f})$ over $\chi^2(\hat{f}) = \chi_{\text{aim}}^2$, where χ_{aim}^2 sets the desired degree of consistency with the actual data. Formally, this is expressed by

$$\text{maximize } (S(\hat{f}) - \lambda \chi^2(\hat{f})) \text{ with respect to } \hat{f} \quad (9)$$

where λ is a Lagrange multiplier. The solution is a balance between the entropy term and the term that demands some agreement with the detected data. Where the balance is struck is determined by the user-supplied value for χ_{aim}^2 . For our cases, increasing χ_{aim}^2 yields increasingly smooth reconstructions as the solution moves toward the smooth underlying model. Decreasing χ_{aim}^2 drives the algorithm toward a solution that relies more heavily on the data, or what might be called a constrained-least-squares or constrained-maximum-likelihood solution.

4.2 Results

Table 2 summarizes the MaxEnt results for the noisy 100-view case. The table gives the detectability index as a function of χ_{aim}^2 for 15 iterations of the MaxEnt algorithm and an assumed rms noise ("acc." in the table) equal to 8, the rms noise in the detected data before pre-smoothing. The table gives both the

expected rms residual for each value of χ_{aim}^2 (in terms of the assumed rms noise in the data) as well as the rms residual measured from the reconstructions. As χ_{aim}^2 decreases from a starting value of 8000, the detectability improves and reaches a maximum at $\chi_{\text{aim}}^2=2900$. For $\chi_{\text{aim}}^2 \geq 2900$ the expected rms residual is equal to the actual rms residual obtained in the reconstructions. When we decrease χ_{aim}^2 further, the rms residual in the reconstructions is unable to keep up with the expected rms residual, and the detectability index begins to decline as well. We find the maximum d' value to be equal to about 1.86, a value slightly less than the best detectability index obtainable with the ART algorithm for this problem. An example MaxEnt reconstruction for $\chi_{\text{aim}}^2=2900$ is shown in Figure 2b.

Figure 3b shows a sample MaxEnt reconstruction for the optimum χ_{aim}^2 parameter in the noise-free, 8-projection case. The MaxEnt algorithm requires a positive-definite estimate of the noise variance, which is set equal to .1 for this case. For 50 iterations of the MaxEnt algorithm, a detectability index of 14.53 is obtained for $\chi_{\text{aim}}^2=8$. This maximum detectability index is somewhat higher than the best d' value obtained with the ART algorithm, as shown in Table 1.

Figure 4 is a plot of the expected rms residual and the actual rms residual in the data as a function of χ_{aim}^2 . The actual rms residual saturates at a value of 0.012, a value substantially lower than the rms residual obtained in the optimized ART reconstructions for the limited-view case. The Maxent detectability index is plotted as a function of the expected rms residual in Figure 5. The detectability index is found to increase with decreasing χ_{aim}^2 until a saturation occurs when the desired rms residual departs from the actual rms residual obtained in the reconstructions. The decrease in detectability as a function of χ_{aim}^2 can be partially explained by the effect depicted in Figure 6, which is a plot of the average amplitude in the signal-present and signal-absent regions in the reconstruction as a function of χ_{aim}^2 . This figure graphically shows the algorithm moving toward a uniform grey solution - the underlying model - as χ_{aim}^2 increases, so that the pixel values in the reconstruction in the disc locations and background become more and more similar. This has the effect of moving the histograms of the decision variable for the signal-present and signal-absent cases closer together, thereby reducing the detectability index provided the histogram widths stay constant. (Although not depicted on the figure, the histogram widths are unchanged or even broadened.)

5. DISCUSSION

We have shown how algorithms can be optimized and compared based on the ability of a simulated observer to do signal-known-exactly detection tasks. We might expect the MaxEnt algorithm, with its implicit positivity constraint, to give similar results to the constrained ART algorithm. This is indeed found to be true for the noise-limited 100-view case. The detectability index of the constrained ART algorithm and the MaxEnt algorithm are both found to be about 1.9.

The noise-free 8-view case compares the ART and MaxEnt algorithms for very limited data sets. We find that the MaxEnt algorithm gives a peak detectability index of 14.5, while constrained ART gives a peak d' of 11.7. It is clear from the noise-free 8-view results that the positivity constraint inherent in MaxEnt and explicit in constrained ART is a powerful tool for reconstructing limited data. In addition, the MaxEnt algorithm demonstrates a superior ability to reduce the rms residual. The average rms residual in the optimized reconstructions is .012 for MaxEnt and .020 for ART. We believe this to be the primary reason for the improved detectability index determined from the MaxEnt reconstructions. In fact, an rms residual of .020 for the MaxEnt algorithm corresponds to a d' value of about 11.5, based on the plot in Figure 4. Thus, the MaxEnt algorithm yields the same d' as the ART algorithm at the same value for the rms residual. The lower rms residual from MaxEnt could be the result of the multiplicative nature of the algorithm. Or, another possible reason is the additional smoothness condition in the MaxEnt algorithm as we have applied it. A test of this supposition would be to use an underlying model with the same average grey level but with an added zero-mean random component so that the model is no longer smooth.

In optimizing MaxEnt, we find in general that the d' value is a maximum at the lowest χ^2_{aim} value for which the desired rms residual is the actual rms residual achieved in the reconstruction. In the noise-limited 100-view case, values of χ^2_{aim} either above or below that optimum cause marked changes in the character of the reconstruction and the associated detectability index. The noise-free 8-view case shows a saturation in d' when the desired rms residual can no longer be met in the reconstruction. Lower values of χ^2_{aim} do not change the reconstruction or the d' value. For both imaging geometries, a choice of $\chi^2_{\text{aim}} = 1/\text{measurement}$ with the estimated noise rms set approximately equal to the actual rms noise in the detected data gives d' values far from the optimum.

We have performed initial investigations of the ability of a simple matched filter to detect a signal when the background is not known a priori using reconstructions obtained from the MaxEnt and constrained ART algorithms. For these studies the filter is a center-surround template that measures the counts in the region where the presence of a signal is being tested, and subtracts off an estimate of the local background activity in an annular region about the signal. For 10 iterations of constrained ART, the need for background fitting in the noisy 100-view case reduces the d' from the known-background value of 1.82 to a d' of 1.53, a decrease of 16%. These d' values are for nominal relaxation parameters after optimization of the areas of the signal and background regions in the unknown-background template. The MaxEnt algorithm yields similar results in the noise-limited 100-view case. For 15 iterations, the unknown-background d' is found to be 1.56, down from the known-background d' of 1.86 by 16%. These results are obtained using the best χ^2_{aim} from the known-background case for both tasks.

Similar performance penalties for the two algorithms are also found in the noise-free 8-view case. For 100 iterations of constrained ART, the optimized d' value decreases from 11.69 to 5.96 when the background has to be estimated, a decrease of 49%. The MaxEnt algorithm gives reconstructions that lead to a similar performance loss: the background-unknown d' is 7.78, which is only 46% of the known-background d' we find of 14.53. These d' values are obtained after 50 iterations of the MaxEnt algorithm with the same χ^2_{aim} value ($\chi^2_{\text{aim}}=8$) that gives the known-background d' . Further investigation of these algorithms and their optimization for the unknown-background task is the subject of current research.

Several possibilities exist for CPU time comparisons for these algorithms. CPU time per iteration can be misleading, since one iteration of MaxEnt can mean up to four projection-backprojection steps. It is perhaps more meaningful to compare CPU times for different algorithms on the basis of CPU time per d' . If we consider the 100-view case, on a μ VAX 3500 machine, 15 iterations of the MaxEnt algorithm takes about 7.5 hours for 10 trial scenes and 128^2 size images. A similar d' value is obtained for 10 iterations of the constrained ART algorithm in less than 2 hours. These times refer to computer simulations that evaluate the detectability index with constant algorithm parameters; no optimization steps are involved.

6. CONCLUSIONS

We have presented a method for the comparison of reconstruction algorithms on the basis of task performance. The algorithms were first optimized to provide the best images for detection of low-contrast discs by a simple matched filter, as measured by the detectability index d' . In particular, we have shown how the choice of the χ^2_{aim} parameter in the historic MaxEnt algorithm affects the detectability index.

For many noisy views, we found the known-background d' values obtained from the ART and MaxEnt algorithms to be very similar. The MaxEnt algorithm was found to give a slightly lower d' than constrained ART, which has been shown by Hanson to be slightly inferior to unconstrained ART for this imaging geometry. Having to estimate the background level caused a performance degradation of 16% for images obtained from either algorithm.

For a few noise-free views, the algorithm ranking was reversed. In this case the MaxEnt algorithm gave a higher d' value than the constrained ART algorithm, which in turn has been shown to be superior to

the unconstrained ART algorithm. An unknown background caused similar relative performance penalties for the MaxEnt and constrained ART algorithms, so that the need for background estimation did not affect the algorithm ranking.

We have found the CPU time of the MaxEnt algorithm to be 2-4 times slower than the ART algorithm for simulations that yield approximately equal detectability indices.

Finally, aside from the specific results for the algorithms considered here, we emphasize the underlying theme of this work. It is important that images be evaluated based on an objective figure of merit that summarizes their usefulness for whatever task they are intended for.

7. ACKNOWLEDGMENTS

This work was supported in part by the United States Department of Energy under contract number W-7405-ENG-36.

8. REFERENCES

1. K.M. Hanson, "Variations in task and the ideal observer," Proc. SPIE 419, 60-67 (1983).
2. R.F. Wagner and D.G. Brown, "Unified SNR analysis of medical imaging systems," Phys. Med. Biol. 30, 489-518 (1985).
3. K.J. Myers, J.P. Rolland, H.H. Barrett, and R.F. Wagner, "Aperture optimization for emission imaging: effect of a spatially varying background," to be published in J. Opt. Soc. Am. A (June 1990).
4. H.H. Barrett, K.J. Myers, and R.F. Wagner, "Beyond signal-detection theory," Proc. SPIE 626, 231-239 (1986).
5. R.F. Wagner and H.H. Barrett, "Quadratic tasks and the ideal observer," Proc. SPIE 767, 306-309 (1987).
6. K.J. Myers and R.F. Wagner, "Detection and estimation: human vs. ideal performance as a function of information," Proc. SPIE 914, 291-297 (1988).
7. R.F. Wagner, K.J. Myers, M.J. Tapiovaara, D.G. Brown, and A.E. Burgess, "Maximum a posteriori detection: figures of merit for detection under uncertainty," Proc. SPIE 1231 (1990).
8. K.M. Hanson, "Method to evaluate image-recovery algorithms based on task performance," Proc. SPIE 914, 336-343 (1988).
9. K.M. Hanson, "POPART - Performance OPTimized Algebraic Reconstruction Technique," Proc. SPIE 1001, 318-325 (1988).
10. K.M. Hanson, "Optimization for object localization of the constrained algebraic reconstruction technique," Proc. SPIE 1090, 146-153 (1989).
11. K.M. Hanson, "Method to evaluate image-recovery algorithms based on task performance," to be published in J. Opt. Soc. Am. A (June 1990).
12. K.M. Hanson, "Object detection and amplitude estimation based on maximum a posteriori reconstructions," Proc. SPIE 1231 (1990).
13. A.E. Burgess, R.F. Wagner, R.J. Jennings, and H.B. Barlow, "Efficiency of human visual discrimination," Science 214, 93-94 (1981).
14. K.J. Myers, H.H. Barrett, M.C. Borgstrom, D.D. Patton, and G.W. Seeley, "Effect of noise correlation on detectability of disk signals in medical imaging," J. Opt. Soc. Am. A 2, 1752-1759 (1985).
15. A.J. Simpson and M.J. Fitter, "What is the best index of detectability?" Psych. Bull. 80, 481-488 (1973).
16. D.M. Green and J.A. Swets, *Signal detection theory and psychophysics*, John Wiley (New York, 1966).
17. R. Gordon, R. Bender, and G. Herman, "Algebraic reconstruction techniques for three-dimensional electron microscopy and x-ray photography," J. Theor. Biol. 29, 471-481 (1970).
18. S.F. Gull and J. Skilling, "Maximum entropy method in image processing," in *Indirect imaging*, J.A. Roberts, ed., Cambridge Univ. Press (1984).
19. MEMSYS 2^(c) maximum entropy data processing system, Maximum Entropy Data Consultants Ltd.

Table 1. Summary of ART results. The stated rms noise refers to the noise in the data before the pre-smoothing operation.

No. Views	No. Iter.	rms noise	λ_0, r_λ	rms residual		d'	
				without constraint	with constraint	without constraint	with constraint
100	10	8	nominal	3.55	3.85	1.99	1.82
100	10	8	optimiz.	3.62	3.99	2.01	1.91
100	100	8	optimiz.		3.80		1.84
8	10	0	nominal	0.06	0.76	0.46	0.65
8	10	0	optimiz.	0.96	0.18	0.48	5.14
8	100	0	optimiz.		0.02		11.69

Table 2. Summary of MaxEnt results for the 100-view case. The stated rms noise refers to the data before the pre-smoothing operation. "acc" is the expected rms noise in the data input to the MaxEnt algorithm.

No. Views	No. Iter.	χ^2_{aim}	rms. noise	acc.	rms. residual		d'
					aim	actual	
100	15	1600	8.0	8.0	2.83	3.79	1.82
100	15	2700	8.0	8.0	3.67	3.80	1.81
100	15	2900	8.0	8.0	3.81	3.81	1.86
100	15	3200	8.0	8.0	4.00	4.00	1.78
100	15	5000	8.0	8.0	5.00	5.00	1.01
100	15	8000	8.0	8.0	6.32	6.34	0.58

Fig. 1. Sample source object showing 10 high-contrast (amplitude=1.0) and 10 low-contrast (amplitude=0.1) discs.

Fig. 2a. Optimized reconstruction of object shown in Fig. 1 from 100 noisy (noise rms=8) projections after 100 iterations of constrained ART. ($d'=1.84$)

Fig. 2b. Reconstruction of object shown in Fig. 1 from 100 noisy (noise rms=8) projections after 15 iterations of optimized MaxEnt. ($d'=1.86$)

Fig. 3a. Optimized reconstruction of object shown in Fig. 1 from 8 noise-free views after 100 iterations of constrained ART. ($d'=11.69$)

Fig. 3b. Reconstruction of object shown in Fig. 1 from 8 noise-free views after 50 iterations of optimized MaxEnt. ($d'=14.53$)

Fig. 4. Effect of varying χ^2_{aim} on the expected rms residual and the average rms residual obtained from the MaxEnt reconstructions for the noise-free 8-view case.

Fig. 5. Plot of the MaxEnt detectability index as a function of the desired rms residual for the noise-free 8-view case. The desired rms residual was varied by varying χ^2_{aim} .

Fig. 6. Average value in the signal-present and signal-absent locations in the MaxEnt reconstructions as a function of χ^2_{aim} for the 8-view, noise-free case.