UC 406
LBL-27088

Lawrence Berkeley Laboratory
UNIVERSITY OF CALIFORNIA

Engineering Division

Presented at the IEEE Nuclear Science Symposium,
San Francisco, CA, January 16–18, 1990, and
to be published in the Proceedings

The Feasibility of Images Reconstructed with the Method of Sieves

E. Veklerov and J. Llacer

April 1989

For Reference

Not to be taken from this room

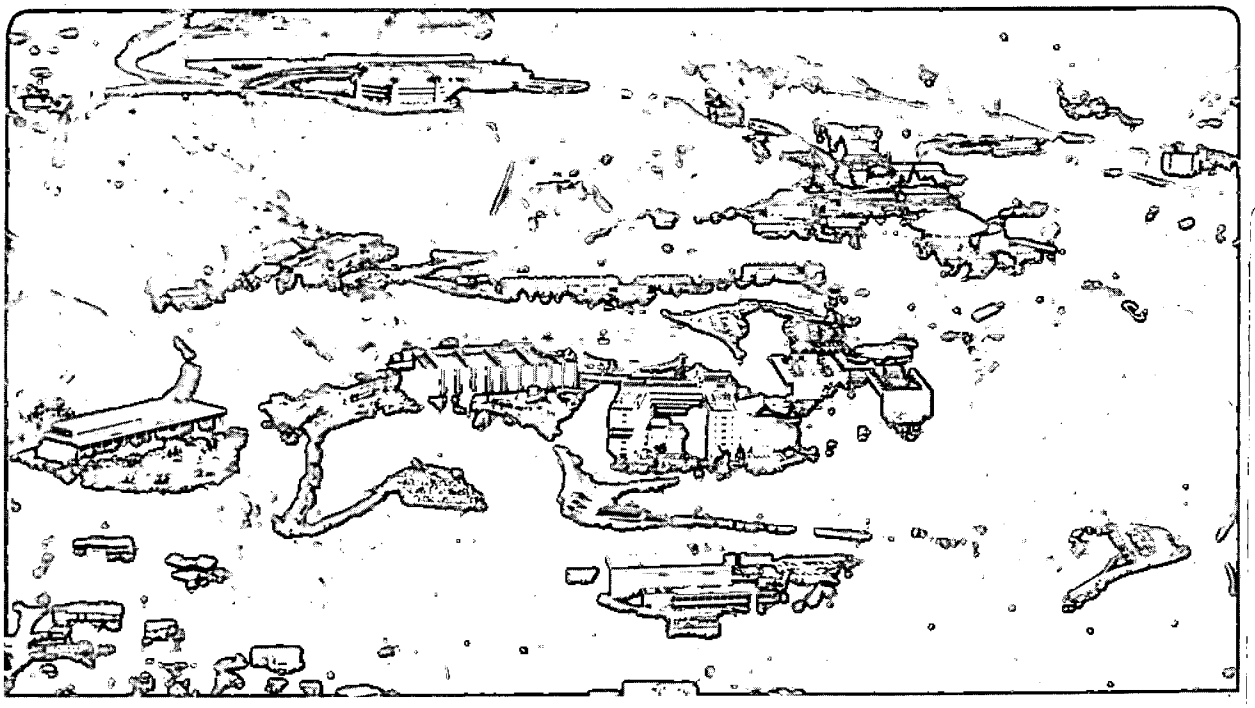

Prepared for the U.S. Department of Energy under Contract Number DE-AC03-76SF00098.

Bldg. 50 Library.
Copy 1
LBL-27088

DISCLAIMER

This document was prepared as an account of work sponsored by the United States Government. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor the Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal 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 its 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, or the Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or the Regents of the University of California.

LBL-27088

THE FEASIBILITY OF IMAGES RECONSTRUCTED WITH THE METHOD OF SIEVES

Eugene Veklerov and Jorge Llacer
Engineering Division
Lawrence Berkeley Laboratory
1 Cyclotron Road, Berkeley, CA 94720

Abstract

The concept of sieves has been applied with the Maximum Likelihood Estimator (MLE) to image reconstruction. While it makes it possible to recover smooth images consistent with the data, the degree of smoothness provided by it is arbitrary.

It is shown that the concept of feasibility is able to resolve this arbitrariness. By varying the values of parameters determining the degree of smoothness, one can generate images on both sides of the feasibilty region, as well as within the region. Feasible images recovered by using different sieve parameters are compared with feasible results of other procedures. One– and two–dimensional examples using both simulated and real data sets are considered.

I. INTRODUCTION

The use of the maximum likelihood criterion in emission tomography has been recognized as a promising approach to image reconstruction. Its numerous merits have been widely discussed in literature. An "unfortunate fact of life" for this criterion is also well known: the reconstructions maximizing the criterion suffer from the image break–up effect. This effect takes place regardless of any specific algorithm maximizing this criterion. Furthermore, it holds for other criteria, such as the weighted likelihood criterion proposed by Llacer and Veklerov in [1], that are based on matching the image projections and the data as closely as possible in a sense different from the maximum likelihood. The essence of the image breakup effect has been adequately explained by several authors and can be summarized as follows. Due to the finiteness of the number of counts, the imperfection of data sampling and of the transition matrix, the data contain noise and the maximum likelihood image tries to follow the noise in the data too closely. For the badly posed inverse problem of tomographic reconstruction, the result is an amplification of the noise in the image.

*This work has been supported by a grant from the National Cancer Institute (CA-39501) and the U.S. Department of Energy under Contract No. DE-AC03-76SF00098.

To overcome this effect, two remedies have been proposed by adding a constraint to the solution, implemented in the *image space* and in the *projection space*, respectively. The first remedy is due to Snyder, Miller and others, see [2] and [3]. It utilizes the method of sieves which restricts the class of permitted images to those satisfying a smoothness constraint. Alternatively, Snyder et al point out that, instead of estimating an image, we can estimate another quantity which is its filtered version. This approach uses a "resolution" kernel. The former approach (sieves) stems from the belief that, regardless of the data, the underlying image that generated the data must have had some degree of smoothness. The latter (resolution kernel) addresses the issue of the finiteness of the resolution of the instrument which does not allow us to see very small details. Therefore, the authors postulate, if they cannot be seen anyway, let us see their blurred version. In other words, both approaches make tacit use of *a priori* information about the image to be recovered by discriminating in favor of smooth images. In fact, Snyder et al, in [3], propose to use both methods simultaneously to control both the image breakup and the overshoot resulting from the use of sieves alone.

The main difficulty of the method of sieves, so far as its implementation is concerned, lies in the arbitrariness of the degree of smoothness. The method offers a family of reconstructions, rather than a unique reconstruction, the two extreme members of the family being the regular broken–up maximum likelihood image and the uniformly grey image. The resolution kernel has a direct physical interpretation and, therefore, its parameters can be quantified. However, as its authors point out, the resolution kernel alone (without sieves) cannot remedy the problem.

The second remedy, in projection space, was proposed by the authors of this paper in [4] and in subsequent papers. They defined the concept of feasibility as follows: an image is said to be feasible if, taken as a radiation field, it could have generated the initial projection data by the Poisson process that governs radioactive decay. Historically, the concept of feasibility appeared in the works of Trussell [11,12], who applied it to image restoration. As applied to tomography, it appeared first as a "stopping" rule for the EM algorithm in the MLE method. The following observation was reported in [4]. All feasible images

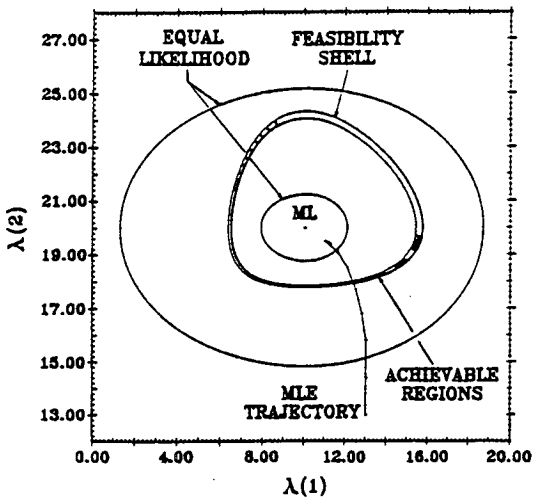

Figure 1: The feasibility shell and a trajectory of the MLE. Ordinate and abcissa are two out of many projection data variables.

make up a closed shell surrounding the maximum likelihood point, which is shown schematically in Fig. 1. All points outside the shell correspond to images for which the data and the forward projections are, on the average, too far apart to be consistent with the Poisson distribution. All points inside the shell correspond to images for which the data and the forward projections are too close. Finally, all points belonging to the shell correspond to images for which the differences between the data and the forward projections are consistent with the Poisson assumption. As the MLE iterative process progresses, it traverses the area outside the shell, passes through the shell inward and finally converges to a point inside the shell.

It was realized later [7,10] that feasibility is a fundamental property of a reconstruction and a necessary condition for an image to be acceptable. Moreover, it was found in [7] that the MLE is but one way of recovering feasible images. Thus [7] describes a Bayesian algorithm recovering feasible images which are distinct from those recovered with the MLE and discuses the finding that it is also possible to obtain feasible images by iterating with the MLE past the feasibility point and post-filtering with a Gaussian kernel to return to the feasibility shell.

The rest of this paper examines the feasibility of reconstructions recovered with the method of sieves, including the resolution kernel. It is demonstrated that it is possible to specify sieve and resolution parameters in such a way that the generated trajectory converges to a point lying in the feasibility shell. This resolves the uncertainty of the method of sieves mentioned above. Finally, we compare different images generated with the method of sieves with feasible images generated by the MLE.

II. ONE–DIMENSIONAL CASE

The one–dimensional case serves as a convenient test–bed model in this paper. The idea of feasibility is applied to a one–dimensional experiment described in [2]. The experiment was as follows.

A smooth Gaussian shaped profile was defined in the interval [0.25, 0.75]. The interval was further subdivided into 512 pixels (bins) and a Poisson process was generated in each pixel with the mean proportional to the integral of the Gaussian profile over the pixel. The total number of counts was, on the average, 1000.

We take the results of the Poisson process as measured data which are to be used to reconstruct the original Gaussian. An unconstrained MLE process returns the exact noisy data as a solution, since that is the solution with maximum likelihood. In order to introduce a smoothness constraint, two schemes were proposed in [2]. One is based on using a penalty function, the other on convolution–kernel sieve. We have implemented a slightly modified version of the first scheme and the second scheme in its original form. For the details, see Appendix A.

Both schemes include the parameter BW (bandwidth) that determines the degree of smoothness and is defined in [2]. Figure 2 shows the estimates derived using the first scheme with 4 values of that parameter (for aesthetic reasons we averaged the values in every 3 adjacent bins in the figure and showed them as one point).

BW	H
0.0001	172
0.001	87
0.005	25
0.01	16
0.05	24
0.1	14
0.3	39
0.5	118

Table 1: The feasibility of one–dimensional reconstructions, as measured with the H parameter, as a function of BW.

The computed "reconstructions" were then tested for feasibility, as described in [4]. Note that the Gaussian profile itself is feasible by definition. Roughly speaking, we might expect that those reconstructions that on the average follow the data much more closely, or much less closely, than the profile are not feasible. This guess was confirmed by the results of the feasibility tests. A small sample of these results is shown in Table 1. The parameter H (hypothesis testing function) reflects the correctness of the hypothesis that the reconstruction could have generated the measured data by a Poisson process. For values H below approximately 30 to 40, we cannot reject that hypothesis, i.e., we accept it.

When the values of BW are too small (less than 0.003 in our experiment), the reconstructions are much farther away from the data than the Gaussian profile is from the

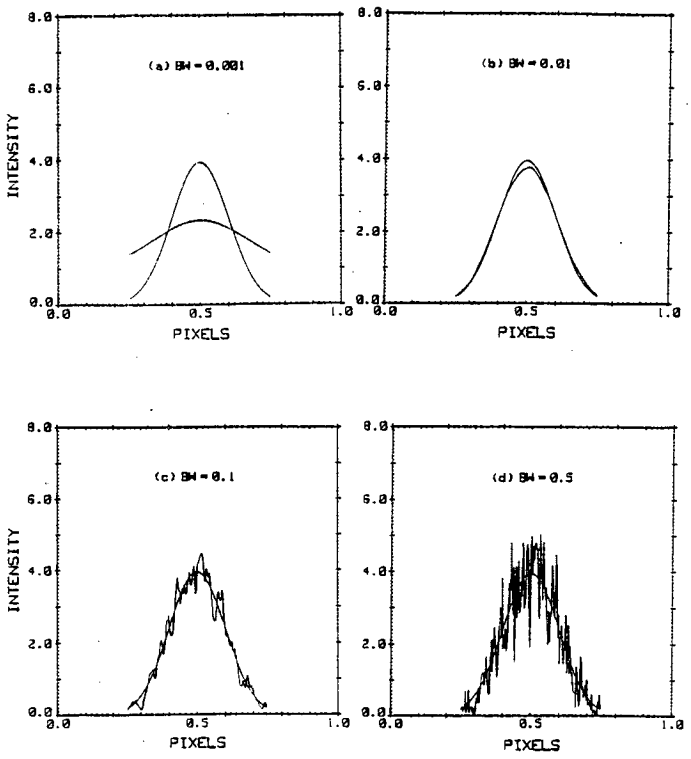

Figure 2: Original and reconstructed values of the one–dimensional Poisson process generated with a Gaussian mean rate. The parameter BW is 0.001 (a), 0.01 (b), 0.1 (c), and 0.5 (d).

data and the points of convergence lie outside the feasibility shell. When BW is too large (more than 0.2), the reconstructions follow the data much closer than the profile follows the data and the points of convergence lie inside the shell. Finally, the remaining range of BW (0.003 to 0.2) is such that the corresponding points of convergence belong to the shell.

The range of acceptable BW computed above, 0.003 — 0.2, is rather wide but it is consistent with the results of other tests. For example, the weak feasibility test defined in [7] yielded the "best" value of BW between 0.08 and 0.09.

Finally, similar results were obtained when we applied the second reconstruction scheme of Snyder et al to the same problem of deriving maximum–likelihood estimates of one–dimensional Poisson data.

III. TWO–DIMENSIONAL CASE

The results of the previous section suggest that the method of sieves may produce feasible images if the corresponding "smoothness" parameters are appropriately adjusted. These results have been extended on PET tomography. To accomplish this, we chose to reconstruct Hoffman's brain phantom, as we did in our previous papers, such as [10], using real PET data supplied by UCLA.

Specifically, we applied the discrete version of the iterative scheme proposed in [3] that is capable of suppressing both the noise and the edge artifact. Let $\lambda^n(b)$ be the estimated b-th pixel value at the n-th iteration and $n^*(d)$ — the number of counts detected in the d-th tube. Then, the iterative scheme is as follows:

$$\lambda^{n+1}(b) = \lambda^n(b) \sum_{d=1}^{D} \frac{n^*(d) k(b,d)}{\sum_{c=1}^{B} \lambda^n(c) k(c,d)} \qquad (1)$$

where $k(b,d)$ is the convolution of the original transition matrix with a Gaussian density of standard deviation $\sigma_k^2 = \sigma_s^2 - \sigma_r^2$. The original matrix was computed for the geometry of the tomograph by the prescription of Shepp and Vardi in [9]. B and D are the numbers of pixels and tubes, respectively.

The parameters σ_s and σ_r here are the standard deviations of the sieve and resolution kernels, respectively. The former has the effect of suppressing noise, the latter — the edge artifact. This iterative scheme was carried out to infinity (300 iterations sufficed for a data set with 1M counts) and the final image was obtained by convolving the results of the last iteration with the sieve kernel.

As we had done in our earlier work, the algorithm took advantage of a pre–computed matrix. However, since the matrix $k(b,d)$ is the result of a convolution operation, it may not be as sparse as it was in the "pure" MLE case which could preclude storing it efficiently. Fortunately, for the range of the parameters that are of interest to our experiment, we could truncate the Gaussian density beyond a square of 7 by 7 pixels in a 128 by 128 grid. This resulted in a somewhat larger number of non–zero elements in the matrix but the increase was still manageable. Just as it was in the one–dimensional case, the purpose of this two–dimensional experiment was to try to recover feasible images by varying the "free" parameters σ_s and σ_r.

As was shown in [10], the original feasibility test introduced in [4] for simulated data is not applicable to real life situations where characteristics of the system are not known exactly. That is why we developed a modified feasibility test in [10] which was employed in the experiments presented in this section. The resulting hypothesis testing function H has values that depend on a parameter ϵ which quantifies the degree of imprecision in the knowledge of the matrix elements. Throughout the present experiments we have used the same value of ϵ that led to reliable feasibility indications in [10] for the same data set. Note that the modified feasibility test can distinguish only the points outside the feasibility shell from the points belonging to the shell but not the latter from the points inside the shell. Hence, an image is said to be feasible in the generalized sense if it lies at or near the outer boundary of the shell.

Table 2 shows a few pairs of the parameters σ_s and σ_r, the unit of length being a pixel side (2.033 mm) on a 128 by 128 grid, and the resulting values of the parameter H

that quantifies feasibility. The reader is referred to [10] for the details of computing H. Taking into account the preceding discussion, feasible images are those for which H is below 40 and which lie near the outside of the feasibility shell.

σ_s	σ_r	$\sqrt{\sigma_s^2 - \sigma_r^2}$	H
0.5	0.50	0.0	20.9
0.6	0.60	0.0	33.1
0.7	0.70	0.0	59.3
0.4	0.26	0.3	4.1
0.5	0.40	0.3	16.9
0.6	0.52	0.3	31.5
0.7	0.63	0.3	55.2
0.5	0.00	0.5	6.6
0.6	0.33	0.5	14.8
0.7	0.49	0.5	25.2
0.8	0.62	0.5	65.1
1.0	0.00	1.0	17.3
1.1	0.46	1.0	32.5
1.2	0.66	1.0	94.1
1.5	0.00	1.5	20.1
1.6	0.56	1.5	31.5
1.7	0.80	1.5	56.6

Table 2: Pairs of σ_s and σ_r and the resulting H values.

Note that the concept of feasibility, like any other concept based on testing statistical hypotheses, is not crisp in the sense that there is a grey area separating feasible and unfeasible images. This lack of crispness follows from a measure of arbitrariness in choosing the allowable testing error. Hence, there is a small amount of slack in pinpointing the threshold values.

The images corresponding to several pairs of the parameters in Table 2 are shown in Fig. 3. The data set that generated the images contained 1 million (1M) counts and was obtained from an ECAT-III tomograph at UCLA. The approximately 6 per cent delayed coincidences corresponding to random events were subtracted previous to the reconstruction, setting negative differences to zero. A feasible image obtained by the standard MLE by computing to iteration 50 (past the feasibility point) and post-filtering with a Gaussian 2-dimensional kernel of σ=0.6 pixels is shown in Fig. 3a). Figure 3b) shows a standard MLE reconstruction after 500 iterations, without any filtering, showing the typical image breakup phenomenon. This image is not feasible. Figure 3c) shows another non-feasible image which has converged in the region between the feasibility shell and the ML point. It corresponds to a value of σ_s=0.3 and σ_r=0.0. Figure 3d) is a feasible image that has been obtained with σ_s=0.7 and σ_r=0.49. It still exhibits the image breakup effect. Figure 3e), also feasible, corresponds to σ_s=1.6 and σ_r=0.55. It does not exhibit image breakup. Finally, Fig. 6f) is a reconstruction that is towards the outer layers of the feasibility shell (H=53.19), which appears to be excessively smooth. It has σ_s=1.8 and σ_r=0.995.

The characteristics of the different feasible images obtainable from the sieve and resolution kernel methods have been studied further by reconstructing one data set with 55 million (55M) counts from the same ECAT-III tomograph by several methods. In particular, attention has been focused on bias in narrow valleys and ridges, which would indicate overshoot of the solution, or "edge artifact". With the very large number of counts, statistical errors are kept to a minimum. Convergence has been assured by carrying the iterative process to 500 iterations, where an additional 100 iterations did not change the measured parameters significantly. The resulting reconstructions have been compared to filtered-backprojection results from the same data set, obtained with the "ramp" and the Shepp-Logan filters, which were almost identical to each other and considered unbiased.

The results of the analysis of the 1M and 55M reconstructions indicate that:

a) Feasible sieve reconstructions from the 1M count data set which are acceptably smooth, but not too smooth, having σ_r significantly smaller than σ_s suffer some degree of edge artifact, resulting in biased images in narrow ridges and valleys. The magnitude of peak-to-valley ratios may be in error by as much as 8 per cent in the phantom studied, as determined from the 55M reconstructions.

b) As we make σ_r similar in value to σ_s to avoid edge artifact, the resulting 1M images are, at first ($\sigma_s = \sigma_r = 0.6$), still quite noisy. As we increase the value of the parameters they become suitably smooth (at approximately $\sigma_s = \sigma_r = 0.9$) but then, they are leaving the feasibility region.

Thus, although there are some reconstructions with excellent appearance, it is not easy to obtain images that fulfill the three requirements of being adequately smooth, little or no edge artifact and feasibility at the same time if one's demands are strict. There is a region, however, in which the three conditions are close to being fulfilled.

It is interesting to note that the "optimum" image described above (feasible, smooth and without overshoot), can be obtained by letting the standard MLE iterate past the stopping point and filtering back to feasibility with a Gaussian kernel of σ=0.6 pixels, as in Fig. 3a). Letting the MLE process iterate to "infinity" and then post-filtering, which is equivalent to a sieve reconstruction with $\sigma_s = \sigma_r = 0.6$ pixels, results in a feasible but noisy image.

IV. CONCLUSIONS

This work continues our study of feasible images and various ways of recovering them from real data sets in cases in which the values of the transition matrix (representing the instrument that generated the data) are not known with great accuracy. It has been shown that the method of sieves can be used as such a vehicle and, in

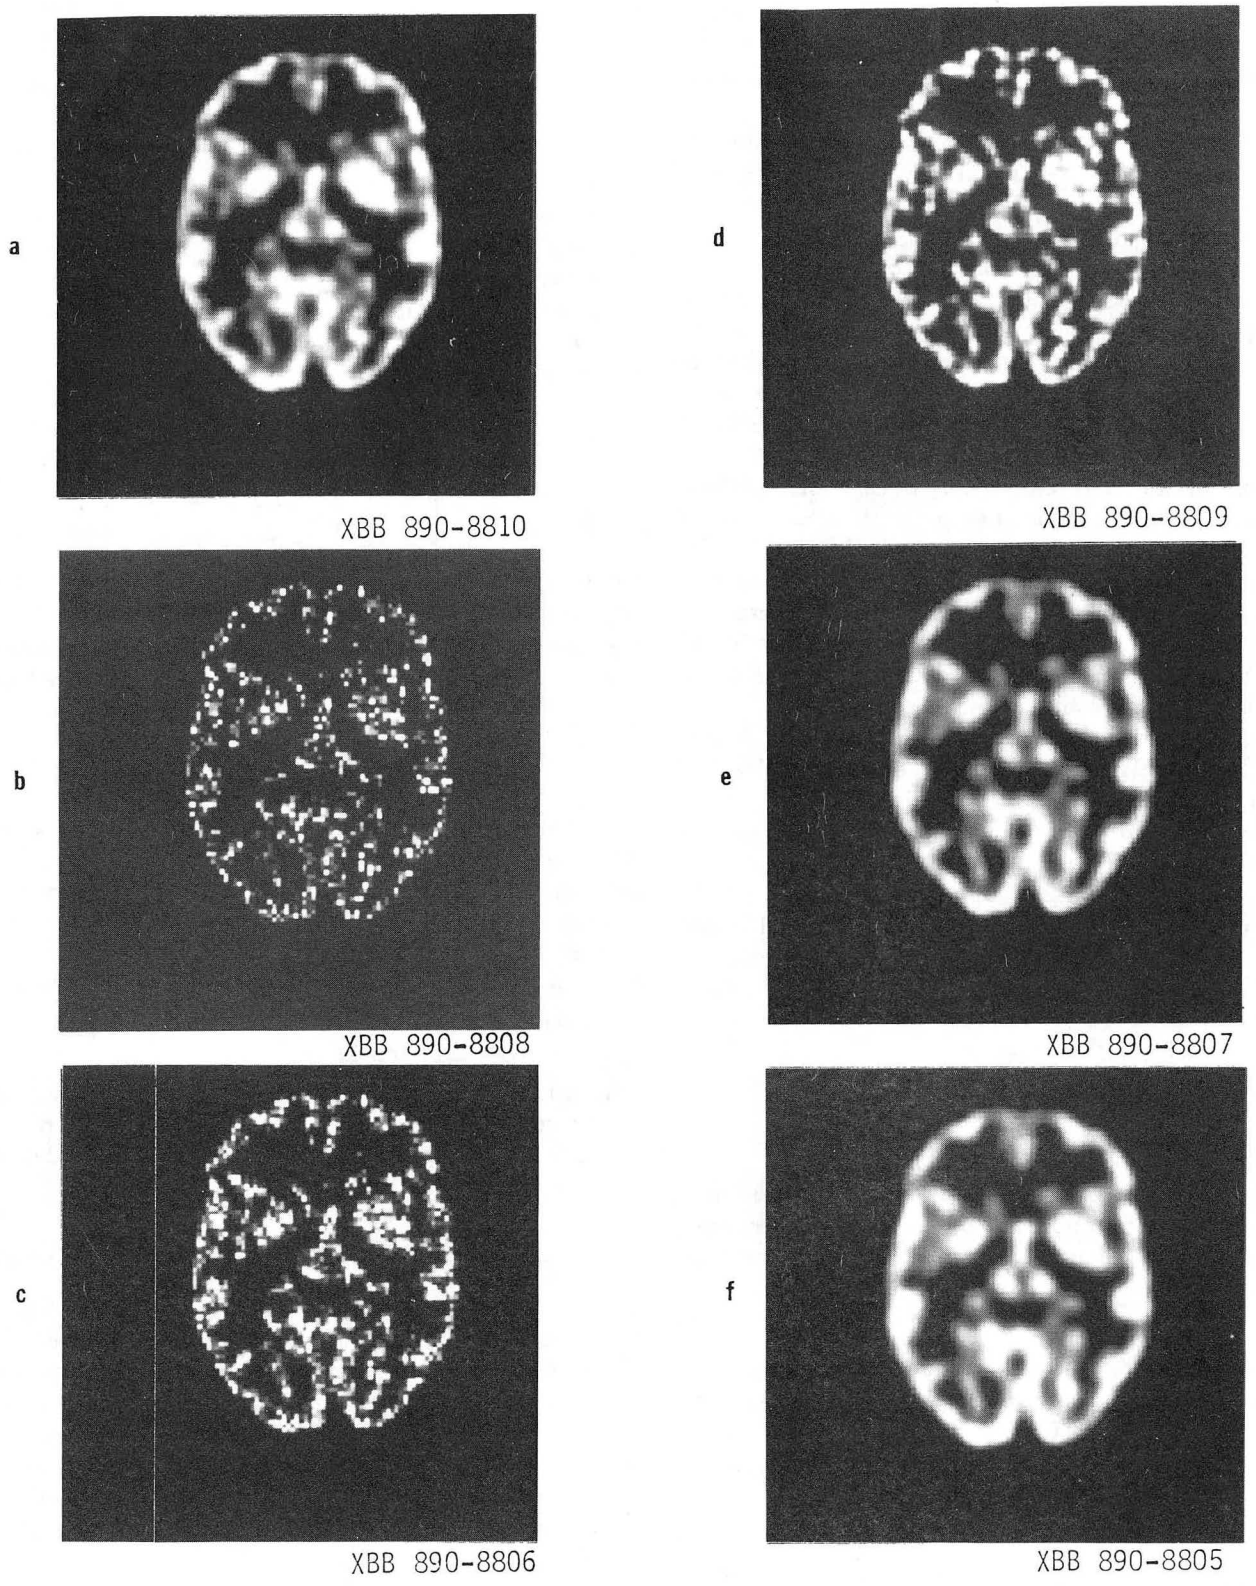

Figure 3: Reconstructions of the 1 million count data set generated by the ECAT-III tomograph. Feasible image by MLE and post-filtering (a), non-feasible MLE results at iteration 500 showing image breakup (b), non-feasible sieve reconstruction with σ_s=0.3, σ_r=0.0 (c), feasible sieve results with parameters 0.7 and 0.49 (d), feasible sieve results with 1.6, 0.49 (e) and marginally feasible results with parameters 1.8 and 0.995 (f).

fact, it allows the user to recover a broad range of feasible images, although only a small band of parameter values yield acceptable images according to criteria of feasibility, smoothness and lack of overshoot.

An interesting question to consider is why are the images recovered by the method of sieves in the region of approximately $\sigma_s = \sigma_r = 0.9$ not feasible and why should they, therefore, not be accepted.

The tomographic image reconstruction problem can be described as an inverse problem that aims are removing the "blurring" effects caused by two actions. The first action is a convolution with a 1/r function due to taking line integrals over a plane and the second action is a convolution with an approximately Gaussian kernel due to the finite size and detection efficiency of the detectors. If the aim of a reconstruction is to remove only the 1/r effects, then, the unsharp reconstructions obtained with $\sigma_s = \sigma_r = 0.9$ pixels and above may be quite acceptable. The test for feasibility, however, will not be passed by such reconstructions because feasibility implies that the reconstruction, if it were a radiation field, could have generated the initial data by a Poisson process. Those sieve images are too unsharp to have given the projection data.

If by some appropriate method we obtain an image that not only "deconvolves" the 1/r effect of tomography, but also some part of the Gaussian blurring function caused by the detectors, we obtain sharper images that pass the feasibility test.

Thus, whether one accepts an unfeasible, or marginally feasible, image that is more blurred than a feasible one depends on the criterion with which one is satisfied. We consider it important to attempt to extract some of the resolution lost in the detection process.

V. APPENDIX

In order to derive our algorithm for the scheme based on using a penalty function, let us reproduce the necessary equations from [2]. The goal is to maximize $G(\lambda)$ given by

$$G(\lambda) = L(\lambda) + \Phi(\lambda) \tag{2}$$

which consists of two components: $L(\lambda)$ is the log–likelihood function and $\Phi(\lambda)$ is a function penalizing roughness. Since the estimate $\lambda(x)$ is non–negative, the function $\gamma(x)$ is introduced, such that $\gamma^2(x) = \lambda(x)$ and equation (2) becomes

$$G(\gamma^2) = L(\gamma^2) + \Phi(\gamma^2) \tag{3}$$

It was proved in [2] that for one particular choice of the penalty function, which is a soft bandwidth constraint on γ, the maximum of $G(\gamma^2)$ is achieved when $\gamma(x)$ is the solution of the equation:

$$\hat{\gamma}(x) = \sum_{i=1}^{N} \frac{h(x - x_i)}{\hat{\gamma}(x_i)} \tag{4}$$

where

$$h(x) = c_1 \cdot \exp(c_2 \cdot \mid x \mid) \tag{5}$$

and the parameters c_1 and c_2 are related to the smoothness constant.

To solve this equation, Snyder and Miller proposed an iterative procedure and observed that the procedure converges for all simulation experiments they attempted. In the present paper we propose a similar but computationally somewhat simpler procedure.

Successively substituting each of $x_1, x_2, \ldots, x_N$ for x in Eq. (4) and letting y_k denote $\hat{\gamma}(x_k)$, we arrive at the following system of algebraic equations:

$$\begin{cases} y_1 = \sum_{i=1}^{N} h_{1,i}/y_i \\ y_2 = \sum_{i=1}^{N} h_{2,i}/y_i \\ \ldots\ldots\ldots\ldots\ldots\ldots \\ y_N = \sum_{i=1}^{N} h_{N,i}/y_i \end{cases} \tag{6}$$

where $h_{k,i} = h(x_k - x_i)$. A standard iterative method for solving System (6) is to start with an arbitrary approximation $y_1^{(0)}, y_2^{(0)}, \ldots, y_N^{(0)}$ and then compute the n-th $(n = 1, 2, \ldots)$ approximation using the iterative procedure:

$$y_k^{(n+1)} = \sum_{i=1}^{N} \frac{h_{k,i}}{y_i^{(n)}} \qquad (k = 1, \ldots, N) \tag{7}$$

It is known (see e.g. Wait [5]) that procedure (7) would converge to a unique solution if the mapping defined by the right–hand sides of Procedure (7) were contracting. However, it is not contracting in our case, which can easily be seen in the trivial case $N = 2$. Since by definition $h_{1,1} = h_{2,2}$ and $h_{1,2} = h_{2,1}$, if we begin with $y_1^{(0)} = y_2^{(0)}$, the successive iterates oscilate: $y_i^{(1)} = y_i^{(3)} = y_i^{(5)} = \ldots$ and $y_i^{(2)} = y_i^{(4)} = y_i^{(6)} = \ldots$.

In order to overcome this obstacle, we will employ the procedure described by Isaackson and Keller [6], pp. 120 – 122. In its simplest form, this procedure can be defined by

$$y_k^{(n+1)} = 0.5 \cdot \sum_{i=1}^{N} \frac{h_{k,i}}{y_i^{(n)}} + 0.5 \cdot y_k^{(n)} \qquad (k = 1, \ldots, N) \tag{8}$$

Note that procedure (8) is similar to the one suggested in [2], p. 3867, which used the geometric mean to compute the next iterate rather than the arithmetic mean in procedure (8). In all the simulation experiments we attempted, both schemes converged to the same solution, the rate of convergence in both cases was geometric but, insofar as computational complexity is concerned, procedure (8) is slightly preferable.

VI. REFERENCES

1. J. Llacer, E. Veklerov, "The high sensitivity of the maximum likelihood estimator method of tomographic image reconstruction". Proc. of International Symposium

CAR'87, Computer Assisted Radiology, Berlin, July, 1987, Springer-Verlag, 1987.
2. D. L. Snyder and M. I. Miller, "The use of sieves to stabilize images produced with the EM algorithm for emission tomography", *IEEE Trans. Nucl. Sci.*, NS-32, No. 5, 1985.
3. D. L. Snyder, M. I. Miller, L. J. Thomas, D. G. Politte, "Noise and edge artifacts in maximum–likelihood reconstructions for emission tomography," *IEEE Trans. Med. Imaging.*, vol. MI–6, pp. 228 – 238, 1987.
4. E. Veklerov, J. Llacer, "Stopping rule for the MLE algorithm based on statistical hypothesis testing," *IEEE Trans. on Med. Imaging*, vol. MI–6, pp. 313 – 319, 1987.
5. R. Wait, The numerical solution of algebraic equations, John Wiley & Sons, 1979.
6. E. Isaacson, H. B. Keller, Analysis of numerical methods, John Wiley & Sons, 1966.
7. J. Llacer, E. Veklerov, J. Nunez, "The concept of feasibility in image reconstruction", Proc. NATO Advanced Study Institute on Mathematics and Computer Science in Medical Imaging, Povoa, Portugal, 1988. Report OBO-25877.
8. J. Llacer, E. Veklerov, "The maximum likelihood estimator method of image reconstruction: its fundamental characteristics and their origin". Proc. of the Xth Information Processing in Medical Imaging (IMPI) International Conference, Utrecht, the Netherlands, 1987.
9. L. A. Shepp, Y. Vardi, "Maximum likelihood reconstruction for emission tomography," *IEEE Trans. Med. Imaging*, vol. MI-1, pp. 113 – 122, 1982.
10. J. Llacer, E. Veklerov, "Feasible images and practical stopping rules for iterative algorithms in emission tomography", *IEEE Transactions on Medical Imaging*, Vol. 8, No. 2,pp 186-193, 1989.
11. H. J. Trussell, "Convergence criteria for iterative restoration methods", *IEEE Trans. Acoust., Speech, Signal Processing*, vol. ASSP–31, Feb., 1983.
12. H. J. Trussell, M. R. Civanlar, "The feasible solution in signal restoration", *IEEE Trans. Acoust., Speech, Signal Processing*, vol. ASSP–32, Apr., 1984.

LAWRENCE BERKELEY LABORATORY
TECHNICAL INFORMATION DEPARTMENT
1 CYCLOTRON ROAD
BERKELEY, CALIFORNIA 94720