Title

ON THE APPLICABILITY OF THE MAXIMUM LIKELIHOOD ESTIMATOR ALGORITHM FOR IMAGE RECOVERY IN ACCELERATED POSITRON EMITTER BEAM INJECTION

Authors

Llacer, J.
Andreae, S.
Chatterjee, A.

Publication Date

1986-03-01

Lawrence Berkeley Laboratory

UNIVERSITY OF CALIFORNIA

Engineering Division

Presented at the International Workshop on Physics and Engineering of Computerized Multidimensional Imaging and Processing, Newport Beach, CA, April 2-4, 1986; and to be published in SPIE Proceedings, Vol. 671

ON THE APPLICABILITY OF THE MAXIMUM LIKELIHOOD ESTIMATOR ALGORITHM FOR IMAGE RECOVERY IN ACCELERATED POSITRON EMITTER BEAM INJECTION

J. Llacer, S. Andreae, and A. Chatterjee

March 1986

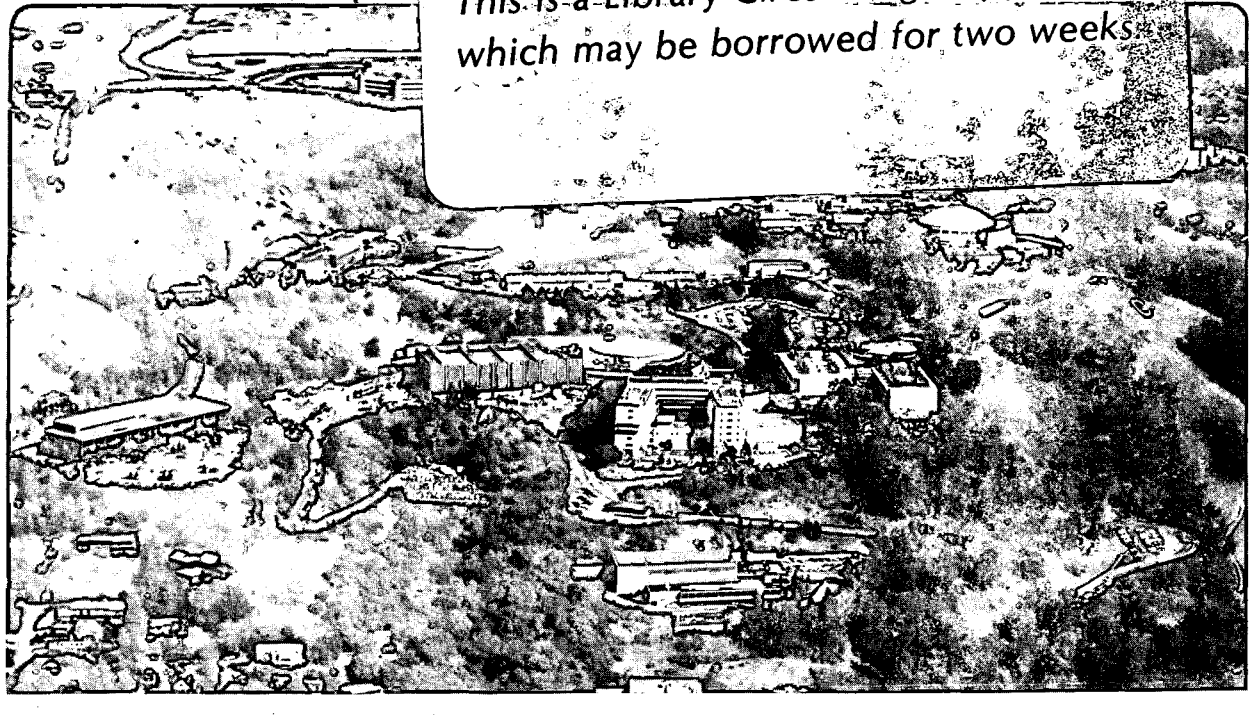

Prepared for the U.S. Department of Energy under Contract DE-AC03-76SF00098

DISCLAIMER

On the applicability of the Maximum Likelihood Estimator algorithm for image recovery in accelerated positron emitter beam injection

Jorge Llacer, Sypko Andreae and Aloke Chatterjee

Lawrence Berkeley Laboratory, University of California, Mailstop 29-100, Berkeley, California, 94720

Abstract

The Maximum Likelihood Estimator (MLE) algorithm for tomographic image reconstruction is being investigated in substantial detail by a number of research groups, as it appears to promise images with very low noise and increased sharpness when compared with filtered backprojection techniques. Recently, however, it has been found that the reconstruction of data from uniform activity distributions exhibits strong peaks and valleys when the number of iterations increases toward a maximum in the likelihood function. This problem has now been investigated with our Positron Emitter Beam Analyzer (PEBA) camera, which, because of its small size and favorable geometry, has allowed an analysis with enough detail to find the origin of that apparent instability.

The findings can be summarized as follows:

1) The very low noise of the MLE reconstructions comes about by the ability of the Poisson-based MLE algorithm to generate an image which favors the matching of experimental data (detector pairs) containing few counts.
2) The image instability at a high number of iterations is a direct consequence of the above characteristic.
3) The matrix of probability elements needed for the MLE reconstruction provides the link between the two above observed phenomena. It appears that, by proper system design, it is possible to obtain the favorable low noise characteristic without the instability.

The applicability of the above findings to true tomography (PEBA does not carry out a true tomographic reconstruction) seems direct, but confirmation should be obtained by further research on the question.

Introduction

The reconstruction of tomographic images by Maximum Likelihood Estimator (MLE) methods is being investigated in substantial detail by a number of research groups,[1-7] as it appears to promise images with very low noise and increased sharpness when compared with filtered backprojection techniques. Recently, however, Snyder and Miller at Washington University have shown examples of what they call the "noise artifact" in reconstructions of simulated data and propose the use of sieves to overcome it in part.[8] The "noise artifact" manifests itself in the appearance of strong peaks and valleys in the reconstruction of areas of uniform non-zero activity as the number of iterations increases. Furthermore, Snyder and Miller show that the effect is fundamental in the application of unconstrained maximum likelihood estimation of density functions based on point-process data, i.e., in processes in which the measurements are made at single points.

At Lawrence Berkeley Laboratory we have also observed the apparent breakup of images of continuous line sources obtained with our PEBA camera[9] when the reconstruction is carried out with a large number of iterations of the MLE algorithm described by Shepp and Vardi.[1,2] Since the PEBA camera (two 64-crystal arrays equidistant from a plane of activity injected by the BEVALAC accelerator into a cancer patient) does not correspond well to the point-process concept of Snyder and Miller (crystals are 1.25 x 1.25 cm spaced by 1.5 cm between centers and the measurements cannot be considered to be made at a point), we felt that we needed to analyze our results with substantial detail in order to ascertain why the image breakup occurs. Corrective steps can be considered subsequently.

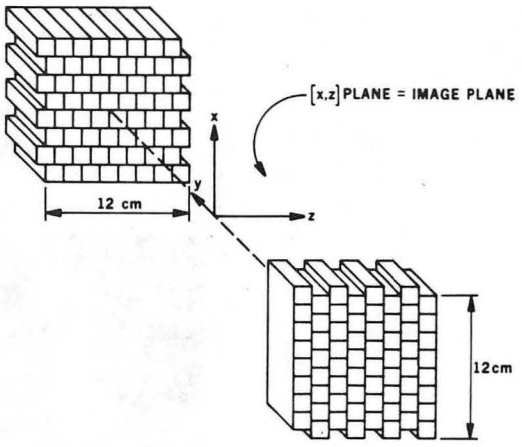

BiGeO DETECTOR ARRAYS FOR PEBA II

XBL 8210-2750

Fig. 1 Schematic description of the two 64-detector arrays of the PEBA camera. The image plane is equidistant from the two detector banks and parallel to them. The arrangements result in a sampling distance of s/4 for the image plane (where s is the detector center-to-center spacing).

The geometry of the PEBA camera, Fig. 1, is such that no tomography is being done during the imaging process. The plane of the Ne-19 atoms injected by the BEVATRON into a patient is well defined and is equidistant to the two camera detector banks. In that situation, the principal function of an image reconstruction algorithm is only to restore the sharpness of the source which has been degraded by the dimensions and crosstalk of the detector crystals. In those conditions, the reconstruction problem is mathematically well posed and the blurring functions of the system connect each pixel only to the more immediate neighbors, instead of having the typical 1/r function of tomography. In this simpler environment we have been able to observe some interesting properties of the MLE which explain the ability of the method to yield very low noise reconstructions. Furthermore, by reducing the imaging problem to two 3-crystal detector banks with 9 coincidences and 9 pixels, we have been able to follow the iterative process in substantial detail and show the direct connection between the favorable low noise character of the MLE and the unfavorable "noise artifact" and ringing of solutions.[8,10] In this paper we shall describe our findings and discuss our present understanding of the conditions under which the advantages of an unmodified MLE method can be obtained without the appearance of the image breakup or "noise artifact", at least in the case of our camera.

Statement of the problem

Figure 2a shows the image of a line source injected by a beam of Ne-19 into a plastic cylinder, as obtained by our standard pseudo-inverse method of reconstruction. The step in the line source resulted from the presence of a 9 mm thickness of a Ca compound in the beam path to simulate bone. The cross section of the reconstructed line is approx. 7.5 mm FWHM. The same data set has also been used to carry our reconstructions with the MLE algorithm described by Shepp and Vardi,[1,2] using a probability matrix calculated for the particular geometry and detector characteristics of the PEBA camera by a combination of Monte-Carlo and deterministic methods.[9] Results after 3, 9, 21, 100 and 200 iterations are shown in Figs. 2b - 2f. The 9 x 9 cm image plane was divided into 48 x 48 pixels of 0.1875 x 0.1875 cm each, with all the activity assumed at the center of each pixel. The sampling distance at the image plane is 0.375 cm in each dimension, for the camera in the configuration of Fig. 1.

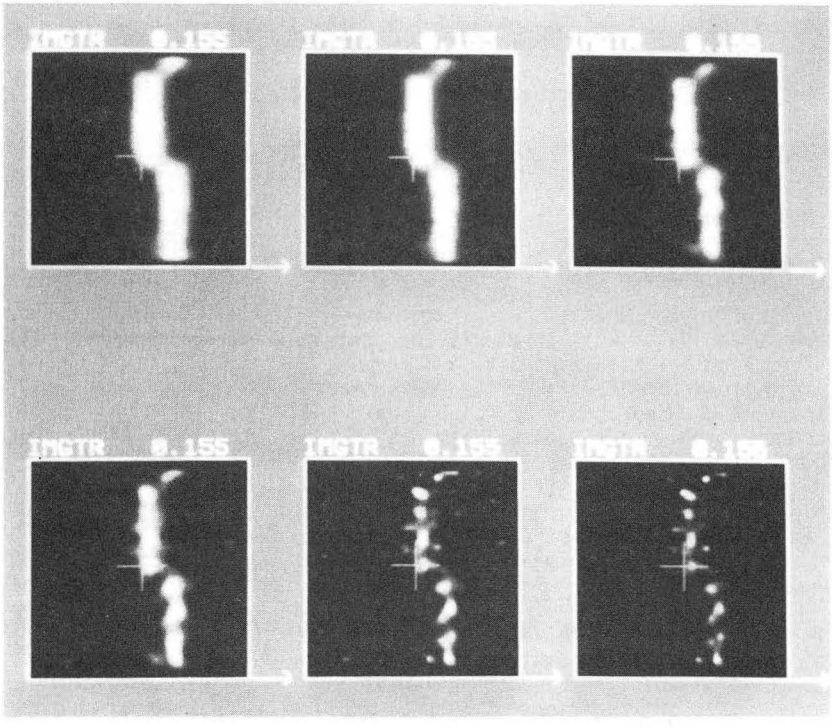

XBB 863-1793

Fig. 2 Reconstructed images from a line source obtained by the injection of a beam of Ne-19 into a plastic cylinder, with a partial obstruction of a Ca compound, causing the step. From left to right, top to bottom: a) image obtained by pseudo-inverse method; b) MLE, 3 iterations; c) MLE, 9 iterations; d) MLE, 21 iterations; e) MLE, 100 iterations, and f) MLE, 200 iterations.

The MLE image after 3 iterations is very similar to the one obtained by our standard pseudo-inverse method and the results after 9 iterations show a substantial improvement in sharpness. Beyond that, image breakup occurs with a periodicity which is clearly related to the structure of the detector banks. MLE reconstructions with pixels of 0.375 x 0.375 cm behaved in the same manner, although going to 0.75 x 0.75 cm pixels (activity at center in all cases) resulted in smooth images even after 100 iterations, when further calculations were useless due to exhaustion of the single precision capabilities of our array processor.

In order to eliminate the effect of inaccurate values of the probability matrix in relation to the true detector efficiencies, the same program that generated the matrix of probabilities was used to simulate a string of point sources forming a line at one edge of the field of view for reconstruction experiments. The data obtained could be modified by the introduction of Poisson or Gaussian statistical fluctuations (depending of the number of counts in a particular detector pair) to simulate any desired activity level in the line source. The choice of location of the line at an edge of the field of view was made so that relatively few coincidence pairs would have to be followed in the subsequent analysis.

Figures 3a and b show the images obtained after 9 and 100 iterations of the MLE algorithm for the simulated line source before statistical fluctuations were introduced. Minor truncation errors existed in the data by the change of the calculated detection probabilities to integer number of counts. Figures 3c and d show the results obtained with the same number of iterations for the simulated line when the statistical fluctuations corresponding to 0.35 μCi/cm counted during 10 s were introduced. The total number of coincidence pairs collected by the detectors was 1065, with a distance between detector banks of 51 cm. This level of counts is realistic for our application to radioactive beam injection. While Figs. 3a and b show a continuous narrow line image with only minor imperfections, and Fig. 3c is still reasonably continuous, we find in Fig. 3d that, after 100 iterations, image breakup similar to the one observed with real camera data has also occurred. We have to conclude, therefore, that the observed image breakup is not necessarily caused by discrepancies between the calculated detection probability functions and the true probabilities of a detector system.

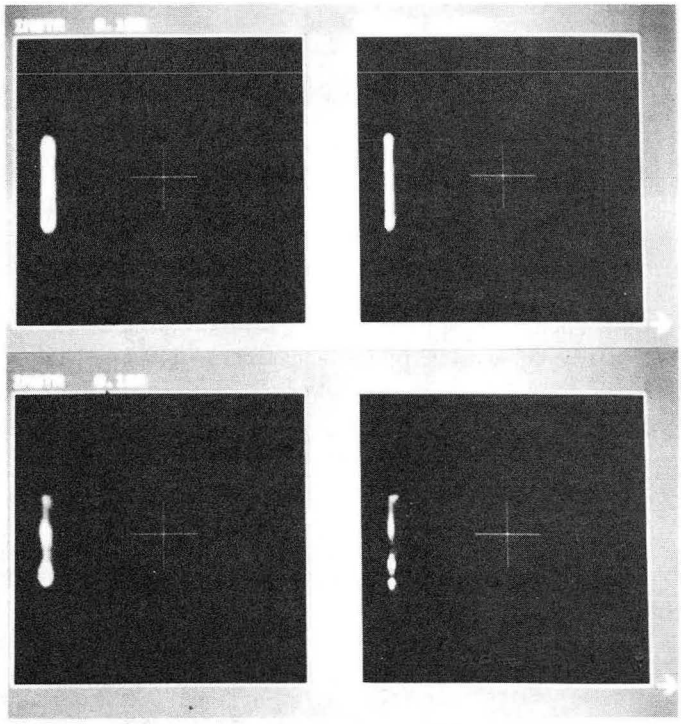

XBB 863-1915

Fig. 3 Reconstructed images from simulated data for a line source at an edge of the image plane. All reconstructions by the MLE algorithm. a) and b), no statistical fluctuations, 9 and 100 iterations, respectively; c) and d), with fluctuations corresponding to 0.35 μCi/cm in the line source, counted during 10 s, 9 and 100 iterations, respectively.

Analysis of the iterative process

We begin the detailed analysis of the MLE process by considering what the MLE algorithm is expected to do - find an image activity distribution $\lambda(b)$ that maximizes the probability of having obtained from it a set of experimental coincidence results $n^*(d)$. The notation of Shepp and Vardi[1,2] will be used throughout this paper. The activity distribution is divided into boxes b and the pairs of detector coincidences are indicated by index d. We presume the existence of a matrix of probabilities $p(b,d)$ which correspond to the probability that a gamma ray pair emitted from box b is detected in detector pair d. The sum of the $p(b,d)$ elements corresponding to a box b is set to unity by pre-normalization of the true calculated detection probabilities. This pre-normalization can be undone after convergence of the MLE solution and it is necessary when the detector system does not have the same total efficiency for all boxes b.

After a number of iterations i, we consider that we have found an image $\lambda^i(b)$. The projection of that image into the detector-pair space is:

$$n^{*i}(d) = \sum_b p(b,d)\lambda^i(b). \qquad (1)$$

If the MLE algorithm is working properly, the total probability:

$$P^i{}_{tot} = \sum_d P[n^*(d)|n^{*i}(d)] \qquad (2)$$

should be increasing as the number of iterations increases. P[j|m] is the probability of obtaining j coincidences when the mean of the distribution is m and it corresponds to the Poisson distribution in our case.

We begin the analysis by checking that, indeed, the total probability of Eq. 2 increases as the number of iterations increases. Table 1 shows the values of $P^i{}_{tot}$ for the simulated line source of Fig. 3, without and with statistical fluctuations, as a function of iteration number.

Table 1. Total Probability that a Reconstructed Activity Distribution has Resulted in a Set of Experimental Detector Coincidences

	$P^i{}_{tot}$	
No. of Iterations	Exact data	0.35 μCi/10 s
3	3780.96	3870.54
9	3831.81	3907.17
21	3856.06	3921.31
100	3879.09	3931.45

The value of the total probability for the images of Fig. 3 is already high after the first iteration because there is a large number of detector pairs without counts and P[0|0] = 1.0. Each correct estimation of zero counts results in the addition of 1.0 to the total probability.

Evidently, the image breakup process is a correct result of the MLE algorithm given the conditions of the problem and it should be interpreted as indicating that the broken source distribution of Fig. 3d has a higher probability of resulting in the particular set of coincidences that were obtained in the simulation (with statistical fluctuations) than a smooth source distribution.

Why this breakup occurs cannot be seen with the application of a global formula like that of Eq. 2, but considerable insight can be gained by looking in detail at the behavior of some of the terms in the sum of Eq. 2 as the number of iterations increases. In particular, we have chosen a 5 x 5 pixel region at the lower end of the line of Fig. 3 for detailed analysis. There are 153 possible detector pairs that can have counts from activity in that region.

In Figs. 4a and b, we have plotted the absolute value of the differences between the number of counts in each of the 153 detector pairs obtained in the simulation and the projected number of counts obtained by Eq. 1 from an MLE reconstruction after 3 and 100 iterations. This difference represents a count error in each coincidence pair. The abcissa has 153 bins ordered by an increasing number of counts in each bin in the experiment simulation data. In the case of Fig. 4a, for exact data, we observe that in going from 3 to 100 iterations most of the bins with 0 experiment counts end up with very small errors. Bins with one experimental count have somewhat larger errors and bins with counts above 29 counts, for example, end up with errors of 1 or 2 counts, which represents roughly a 5% error. When we carry out a similar analysis for the source with statistical fluctuations, Fig. 4b, we again observe small errors for the bins with 0 counts, but much larger errors for the bins with the larger numbers of counts. For bins above 10 counts, for example, bins with 20–40% errors are common.

If we observe at the same time the contribution of the 153 detector pairs to the total probability of Eq. 2, we see in Figs. 5a and b that a much larger increase in probability is contributed by bins with zero or low counts in going from 3 to 100 interations than is contributed by bins with high counts. This can be interpreted as indicating that a maximum likelihood estimator with a Poisson-based likelihood function will favor matching experimental data with low number of counts over data with high counts. This result is not surprising in view of the fact that the derivative of the Poisson function with respect to the mean decreases as the mean increases, for a number of experimental counts in the vicinity of the mean. The ability of MLE reconstructions in tomography to yield images with small noise magnification ratios is then explained by this particular characteristic of the Poisson function. The reconstruction algorithm "focuses" on accurate reconstruction in the regions of detector pairs with low counts.

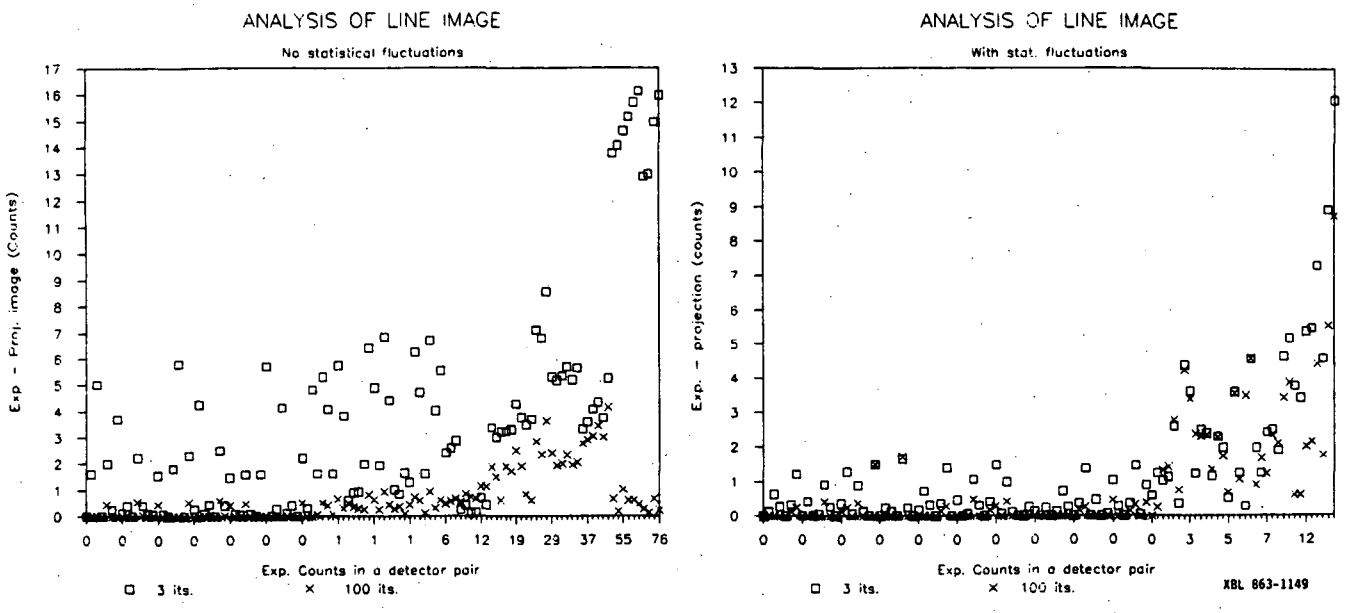

Fig. 4 Errors between "experimental" counts in each detector pair analyzed and the projection of reconstructed images after 3 and 100 iterations for the line source of Fig. 3. a) without statistical fluctuations; b) with fluctuations.

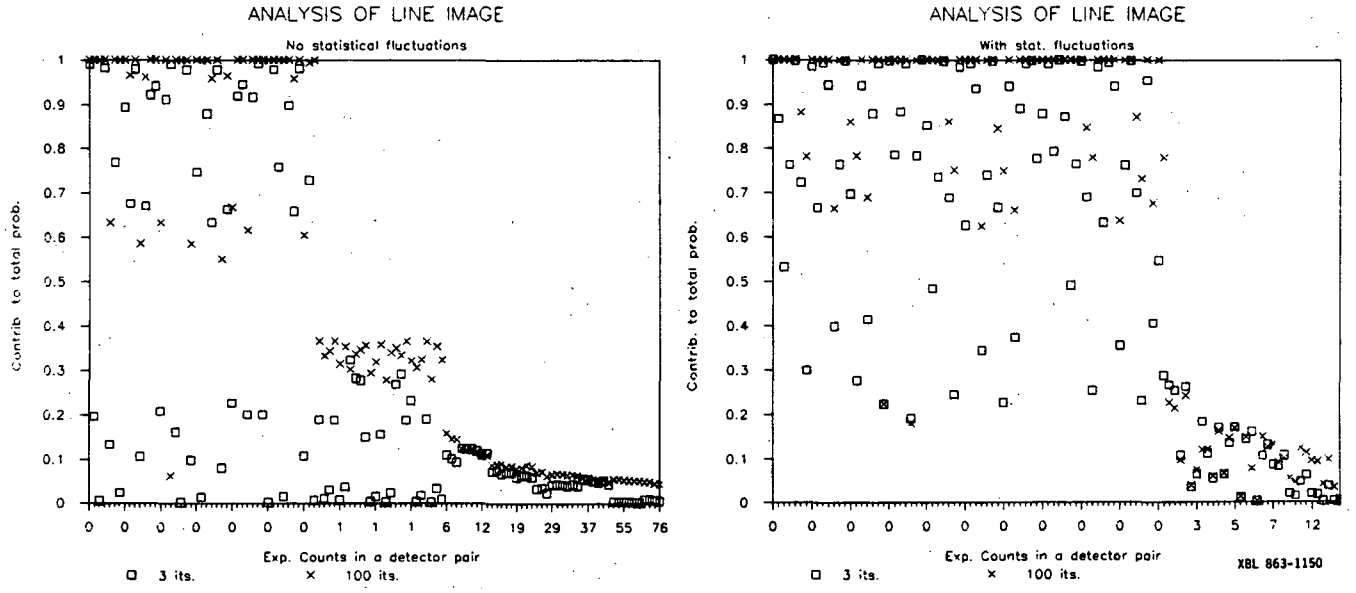

Fig. 5 Contribution to total probability by algorithm's attempt to match "experimental" counts with projection values, after 3 and 100 iteration. a) without statistical fluctuations; b) with fluctuations.

Relationship with the "noise artifact"

Although one may suspect that the ability of the MLE method to reconstruct accurately regions of low counts may result in misbehavior in other regions, any attempt at finding confirmation by study of the example used above has failed because of excessive complexity. We have been able to find such a relationship, however, by studying a much smaller "camera" formed by two groups of 3 detectors facing each other, as shown in Fig. 6. A complete simulation by the program MATRIX which was used to design the PEBA camera[9] has been carried out. The detectors are 1.25 x 1.25 x 3 cm BiGeO crystals, with a center separation of 1.5 cm. The two detector groups have been assumed to be separated by 50 cm. Nine pixels of 0.375 x 0.375 cm, at midline between the detector groups, have been simulated, with activity concentrated at their centers. The choice of activity distribution in a pixel is not considered to be of importance for the present purpose.

A line source of 1.69 cm was simulated by 19 points of equal activity and was placed as indicated in Fig. 6. The reconstruction of the resulting exact data is shown in Fig. 7 after approx. 50 iterations (at convergence with single precision arithmetic). A number of different reconstructions incorporating the statistical fluctuations corresponding to 0.2 μCi counted for 1 s have been generated and a few have been selected for study. We will present here only the detailed study of one reconstruction which exhibited strong overshoot at the edge of the line source near the center of the detector system as the number of iterations increased. Figure 7 shows the reconstructed line of activity at 3, 21 and 200 iterations, in addition to the exact data results.

Table 2 shows the calculated activity distributions for the nine pixels at iterations 3, 9, 15, 21, 100 and 200 before final renormalization.

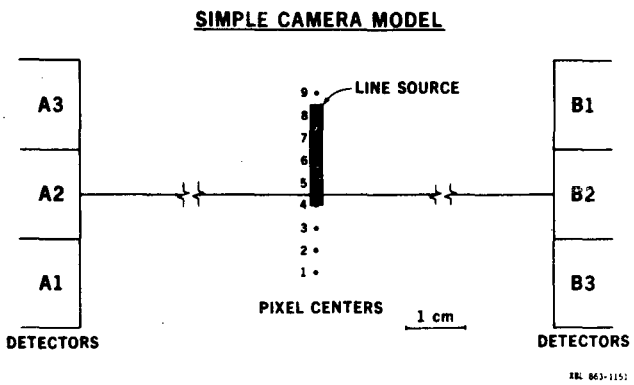

Fig. 6 Description of small coincidence "camera" used for the analysis of the "noise artifact".

Table 2. Calculated Activity Distributions vs. Iteration Number for Line Source and Small Camera

Pixel No.	Iteration Number 3	9	15	21	100	200
1	0.00	0.00	0.00	0.00	0.00	0.00
2	0.00	0.00	0.00	0.00	0.00	0.00
3	0.50	0.00	0.00	0.00	0.00	0.00
4	12.85	9.68	6.63	4.28	0.003	0.00
5	14.85	19.09	22.48	25.11	29.87	29.80
6	10.00	10.17	10.14	10.11	10.25	10.63
7	8.34	7.94	7.59	7.33	6.67	6.35
8	3.34	3.13	3.13	3.13	3.08	3.03
9	2.10	1.97	2.00	2.02	2.12	2.18

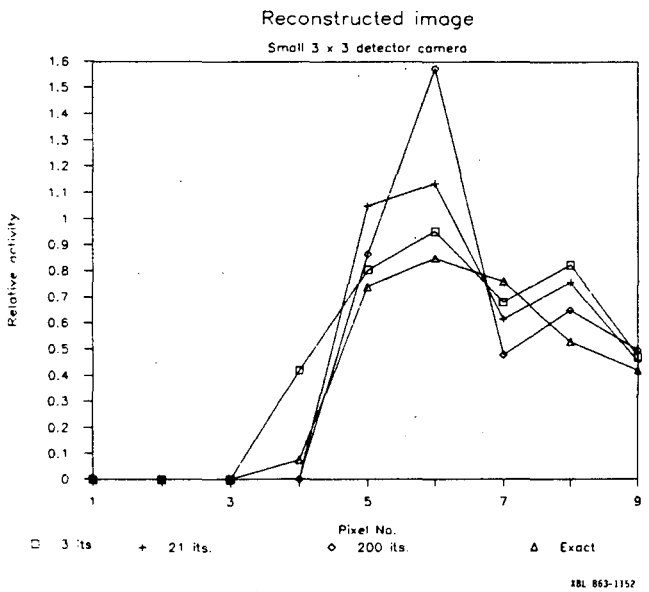

Fig. 7 Values of the image at the 9 pixels of the small "camera" after 3, 21 and 200 its., with statistical fluctuations corresponding to 0.200 μCi counted during 1 s, and exact (without fluctuations). Overshoot and/or "noise artifact" is clearly observable.

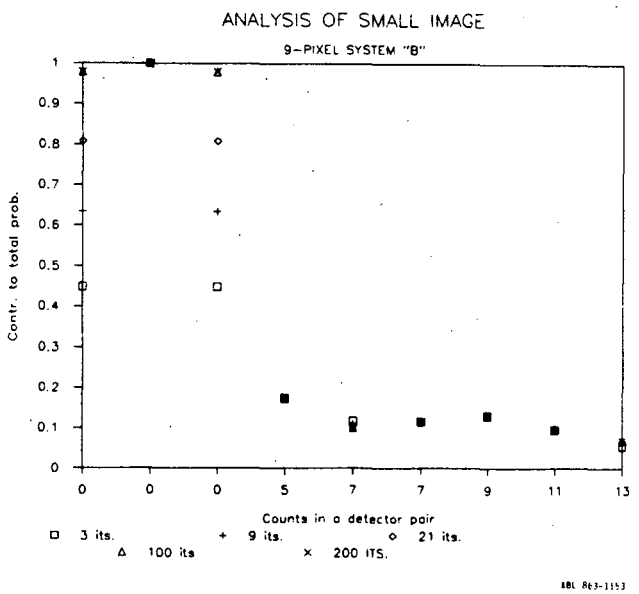

Fig. 8 Contribution to total probability for the 9 pixels of the small "camera" as the iterative process progresses.

The most prominent feature of the data in Table 2 is the decrease in activity in pixel No. 4 as iterations increase, with the concurrent increase in activity in pixel No. 5. This causes the "overshoot" observed in Fig. 7. It is interesting to note that the sum of activities in pixels 4 and 5 remains nearly constant. It appears that activity is being shifted between the two pixels during the iterative process.

Table 3 shows the number of counts in the simulation for each of the nine detector pairs and the projection of the distributions of Table 2, obtained using Eq. 1.

Table 3. Comparison of Simulation with Projection Data as a Function of Iteration Number

Coincidence	Detectors A	Detectors B	Counts, Simulation	Counts, Projection, Iteration 3	9	15	21	100	200
1	1	1	7	8.91	9.22	9.32	9.41	9.5737	9.5732
2	1	2	0	0.80	0.46	0.32	0.21	0.0214	0.0212
3	1	3	0	0.0008	0.0003	0.0002	0.0001	0.0000	0.0000
4	2	1	11	9.03	9.08	9.07	9.05	9.0070	9.0069
5	2	2	13	9.18	9.49	9.60	9.68	9.8373	9.8381
6	2	3	0	0.80	0.45	0.32	0.21	0.0218	0.0217
7	3	1	5	5.31	4.99	4.99	4.99	4.9982	4.9983
8	3	2	7	9.04	9.08	9.07	9.05	9.0133	9.0139
9	3	3	9	8.91	9.21	9.31	9.38	9.5270	9.5266

It would appear, at first, that no significant changes are occurring in the data of Table 3 while iterations are proceeding. It is important to note, however, that the projection values for coincidences 2 and 6 are steadily decreasing towards a final value near zero. Figure 8, showing the contribution to total probability by the different detector pairs as iteration progresses, clearly indicates a strong gain in P^i_{tot} (Eq. 2) in the first and third bins, corresponding to coincidences 2 and 6, respectively. In order to understand the significance of this detail, we show in Tables 4 and 5 the values of the probability elements p(b,d) for coincidences d=5 and d=2, respectively.

Table 4. Values of Probability Matrix for Coincidence No. 5

Pixel No. b	p(b,d)
1	0.0005
2	0.01
3	0.75
4	4.94
5	4.94
6	0.75
7	0.011
8	0.0005

Table 5. Values of Probability Matrix for Coincidence No. 2

Pixel No. b	p(b,d)
1	0.708
2	4.88
3	4.91
4	0.73
5	0.011
6	0.0005

As the iteration progresses, we have observed that pixel 5 (Table 2) increases strongly, while pixel 4 decreases with an almost constant sum. From Table 4 we see that the values of p(b,d) connecting pixels 4 and 5 through coincidence 5 are similar. The observed behavior of pixels 4 and 5 is, therefore, completely consistent with the near constant values of the projections for coincidence No. 5 (Table 3). We still need to know, however, what drives this exchange of activity from pixel 4 to 5.

We have observed that the projection for coincidence 2 decreases steadily because the MLE algorithm gains considerable likelihood if a coincidence with zero counts can be made to match. Since pixels 1, 2 and 3 do not have any activity, the most important remaining component of the projection for coincidence 2 is the activity in pixel 4 (see Table 5). Thus, while the MLE algorithm is pushing for a match of the 0 counts in coincidence 2 with the projection of the successive images, it pulls down the activity in pixel 4. In turn, the values in Table 4 insure that whatever activity is lost by that pixel is placed into pixel 5, leading to the overshoot observed. The role of coincidence 6 is found to be identical to that of coincidence 2.

The above discussion, which in spite of the simple model, is still rather complex, can be summarized as follows:

The emphasis of the MLE algorithm for Poisson distributed processes to match data with low number of counts can result in overshooting and possibly other observed anomalies in regions of higher activity by the influence of p(b,d) terms that connect regions of low activity with those of high activity.

Discussion of results and conclusion

The above analysis indicates conclusively that, for the conditions of the problem studied, the MLE algorithm will result in images with a very low noise magnification ratio in the regions of low or no activity, while peaks and valleys may appear in regions with higher uniform activity when the statistical fluctuations caused by a limited number of counts in a measurement are substantial. The two effects are closely related and are caused by the bias of the MLE algorithm based on Poisson processes to favor matching data with a low number of counts at the expense of accurate matching at higher counts. The link between the two effects is found in the matrix of probabilities, p(b,d), which in turn depends on the design of the imaging device and the size and positioning of pixels for the reconstruction.

For the PEBA camera, we have a detector separation s = 1.50 cm. The normal sampling of the image plane, equidistant from the detector banks, would be at a sampling distance of s/2 for regularly placed detector arrays. For the geometry of Fig. 1, however, the sampling distance improves to s/4, or 0.375 cm. As we indicated in the statement of the problem, image breakup occurs with pixel sizes of 0.1875 and 0.375 cm, but not at 0.75 cm. We are now investigating the behavior for pixel sizes between the last two values. It appears, at this time, that in the same manner as pseudo-inverse solutions (which provide a minimum least squares error solution) are very sensitive to the violation or near violation of the sampling theorem, the MLE solutions will exhibit a similar behavior. Although the sampling theorem has no direct connection with the non-linear MLE algorithm, the attempt at obtaining information with data which are ambiguous to start with results in bad images. Least squares linear methods result in increased noise everywhere in the image field when the sampling theorem is violated, or nearly violated, MLE methods misbehave at areas of high activity.

An extrapolation to true tomographic processes from the PEBA analysis cannot be carried out directly with assurance, but it is not unreasonable to expect similar effects, explanations and possible solutions in terms of system design, and choice of pixel size and location. Research in these areas will continue.

Acknowledgments

The authors would like to thank L.A. Shepp and F.S. Goulding for valuable discussions. This work was supported in part by the National Cancer Institute (CA-27021) and the Office of Health and Environmental Research of the U.S. Department of Energy under Contract No. DE-AC03-76SF00098.

References

1. Shepp, L. A. and Vardi, Y., "Maximum Likelihood Reconstruction for Emission Tomography", IEEE Trans. Med. Imaging, Vol. MI-1, No. 2, pp. 113-121. 1982.

2. Shepp, L. A., Vardi, Y., Ra, J. B., Hilan, S. K. and Cho, Z. H., "Maximum Likelihood PET with Real Data", IEEE Trans. Nucl. Sci., Vol. MS-31, No. 2, pp. 910-913. 1984.

3. Snyder, D. L. and Politte, D. G., "Image Reconstruction from List-Mode Data in an Emission Tomography System Having Time-of-Flight Measurements", IEEE Trans. Nucl. Sci., Vol. NS-20, No. 3, pp. 1843-1849. 1983.

4. Floyd, C. E., Jaszczak, R. J. and Coleman, R. E., "Inverse Monte Carlo: a Unified Reconstruction Algorithm for SPECT", IEEE Trans. Nucl. Sci., Vol. NS-32, No. 1, pp. 779-785. 1985.

5. Lange, K. and Carson, R., "EM Reconstruction Algorithms for Emission and Transmission Tomography", J. Comp. Assisted Tomography, pp. 306-316. April 1984.

6. Llacer, J. and Meng, J. D., "Matrix-Based Image Reconstruction Methods for Tomography", IEEE Trans. Nucl. Sci., Vol. NS-32, No. 1, pp. 855-864. 1985.

7. Llacer, J., Andreae, S., Veklerov, E. and Hoffman, E. J., "Towards a Practical Implementation of the MLE Algorithm for Positron Emission Tomography", IEEE Trans. Nucl. Sci., Vol. NS-33, No. 1, pp. 468-477. 1986.

8. Snyder, D. L. and Miller, M. I., "The Use of Sieves to Stabilize Images Produced with the EM Algorithm for Emission Tomography", IEEE Trans. Nucl. Sci., Vol. NS-32, No. 5, pp. 3864-3872. 1985.

9. Llacer, J., Chatterjee, A., Batho, E. K. and Poskanzer, J. A., "Design Analysis and Performance Evaluation of a Two-Dimensional Camera for Accelerated Positron Emitter Beam Injections by Computer Simulation", IEEE Trans. Nucl. Sci., Vol. NS-30, No. 1, pp. 617-625. 1983.

10. Tanaka, E. and Yamamoto, M., "Utilization of Non-Negativity Constraints in Convolutional Reconstruction of Emission Tomograms", 9th Inter'l. Conf. on Information Processing in Medical Imaging, Washington, D.C. June 1985.

LAWRENCE BERKELEY LABORATORY
TECHNICAL INFORMATION DEPARTMENT
UNIVERSITY OF CALIFORNIA
BERKELEY, CALIFORNIA 94720