

THE USE OF A STOPPING RULE IN ITERATIVE IMAGE RECONSTRUCTION

Jorge Llacer and Eugene Veklerov
Lawrence Berkeley Laboratory
University of California
Berkeley, CA 94720

ABSTRACT

During a study of the characteristics of the Maximum Likelihood Estimator (MLE) method of image reconstruction from Positron Emission Tomography (PET) data, we have found that the requirement that the reconstructed image, if it were a source, could have generated the original data imposes the stopping of the iterative procedure at some point. This requirement appears to be in contradiction with the concept of a maximum likelihood estimator until one realizes that maximizing the likelihood results in reaching for the top of the measurement probability distribution $P(g|f)$ in Bayes' formula:

$$P(f|g) = P(g|f)P(f)/P(g)$$

where the results measurement vector g can be described by $g = Hf + n$, with f being a source, H the response matrix of the measurement system and n the noise vector in the measurement.

From a Bayesian point of view, we should maximize the Maximum a Posteriori Probability (MAP) $p(f|g)$, i.e., the probability that the source be f given a measurement g , and for that purpose one needs to have the *a priori* source distribution $P(f)$, which is not usually available with any degree of reliability.

The use of the stopping rule does not reach for the MAP solution but it uses some important physical prior information: it requires that the iterations stop as soon as the image obtained has a characteristic that we call "feasibility". We define a feasible image as one that could have given the original data by the physical process that governs the measurement.

We show that MLE reconstructions started from a uniform image field pass through a region of feasibility in which the images represent a good compromise between sharpness and smooth regions of high activity, without the "noise artifact" having yet set in. The shape and characteristics of the region of feasibility are described and implications for future work are described.

1. Introduction

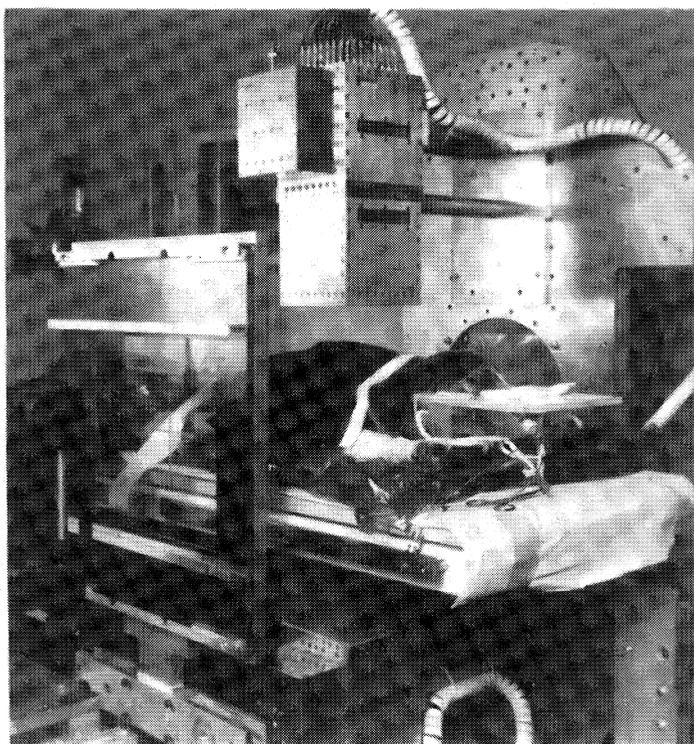
As part of the research program on cancer therapy with heavy ion radiation, it has become important to ascertain the penetration distance of accelerated heavy ion beams in the human body with accuracies on the order of 1 mm. It is one important characteristic of those accelerated ions that their cell destructive power is concentrated in the last few mm of their track. A well placed radiation dose in a small, inoperable tumor can result in its eradication with a tolerable perturbation to surrounding tissues. An error in the estimation of the energy absorbing power of the tissues that the beam must traverse in order to arrive at the target can, however, result in serious damage to a patient if a vital organ or structure is in the vicinity of the target volume. The calculation of absorbing power is normally done by converting the approximate tissue density information obtained from *X-ray CT* scans to electron density and it is known that the process can result in combined errors of 3 to 5% over pathlengths of 20 cm, for example.

Fortunately, it is possible to inject into patients accelerated ions that decay by positron emission, like Ne-19, C-10 and others, so that the penetration distance of the ions in the patient can be measured by detecting the coincidence gamma-rays resulting from the positron annihilations, much like in the case of Positron Emission Tomography (PET). Proper registration by X-radiographs in three dimensions between patient position and a measurement gamma-ray camera can result in the determination of that penetration distance with the desired accuracy.

The process of image restoration from positron emitting ions injected into the human body by a high energy accelerator has been under study at Lawrence Berkeley Laboratory for a number of years [1]. Figure 1 shows an anesthetized dog placed between the two banks of detectors of a Positron Emitter Beam Analyzer (PEBA) camera (solid cubes above and below the animal) during early verification experiments. Each detector bank consists of 64 Bismuth Germanate scintillation crystals with individual photomultiplier tubes that transform the energy of the gamma rays detected to electrical signals. The arrangement of the crystals is shown in Figure 2. The skewing between rows and 90° rotation between banks results in a sampling distance at the central image plane four times smaller than the center-to-center detector distance.

Since the plane of gamma ray emission is defined by the beam injection geometry, the image restoration problem is equivalent to time-of-flight PET with near perfect knowledge of the point of emission in the direction perpendicular to the detector planes. The function of the image restoration algorithm is then limited to deblurring and correcting for the different sensitivities of the camera at different points in the image plane. Because of the high spatial frequency sampling and the non-tomographic nature of the problem, the matrix that characterizes the linear system has a very low condition number and a pseudo-inverse restoration method can be used.

In spite of the success of the pseudo-inverse restoration method, we were motivated to attempt a Maximum Likelihood Estimator (MLE) reconstruction on the PEBA data using the EM algorithm first developed for PET by Shepp and Vardi [2] for two reasons: (1) we deal with a very small number of counts, so that the data exhibit strong Poisson characteristics, and (2) early tests showed that MLE reconstructions were very robust to even violations of the sampling theorem, so that we thought it possible to obtain improved resolution by using pixels smaller than those tolerated by the linear reconstruction method.



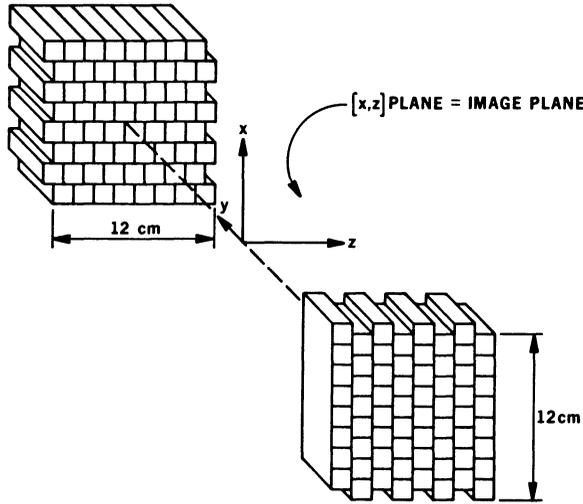
XBB 839-7929

Figure 1. Photograph of PEBA camera in position at the BEVALAC treatment room for an experiment with a live dog. A beam collimator appears at left and an X-ray plate holder with cross-hairs for accurate positioning is also shown.

Our findings [3] were disappointing. Figure 3a shows a pseudo-inverse PEBA image for a flat beam of Ne-19 ions that stopped inside a plastic cylinder within the field of view of the PEBA camera. One portion of the beam was obstructed by a certain thickness of calcium compound, generating the visible step in the image. Figures 3b through f show the MLE reconstructions for 3, 9, 21, 100 and 200 iterations. As the number of iterations grows we observe that we have called “image breakup”, also described as “noise artifact” by Snyder and Miller [4]. It became evident that, if we were to obtain some benefit from the MLE reconstructions, we should investigate the nature of the image breakup and avoid it by a well understood procedure.

2. A physicist’s critical look at maximum likelihood

It is evident that, in the absence of any other information, the best estimate of the mean of a random variable that can be obtained from a single measurement of that variable is the result of the measurement itself. Maximum Likelihood image reconstruction for PET uses that idea for the purpose of making an estimate of an image



BiGeO DETECTOR ARRAYS FOR PEBA II

XBL 8210-2750

Figure 2. 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 arrangement results in a sampling distance of $s/4$ for the image plane, where s is the detector center-to-center spacing.

$\lambda(1), \lambda(2), \dots, \lambda(B)$ whose projection into detector-pair (tube) data space is given by:

$$\lambda^*(d) = \sum_{b=1}^B p(b, d)\lambda(b). \tag{1}$$

The transition matrix values $p(b, d)$ are the probabilities that a gamma-ray pair emitted from pixel b will be detected in tube d . The notation introduced by Shepp and Vardi [2] is used throughout this paper. Maximum Likelihood obtains an estimate of that image by maximizing the likelihood function:

$$L(\lambda) = P(n^*|\lambda) = \prod_{d=1}^D e^{-\lambda^*(d)} \lambda^*(d)^{n^*(d)} / n^*(d)! \tag{2}$$

in which $n^*(1), n^*(2), \dots, n^*(D)$ are the results of a single measurement of each of the independent tube counts. It is evident from looking at Eq.2 that the maximum of that function will occur when the values of $\lambda^*(d)$ are as close as possible to $n^*(d)$. Because of the nature of the radioactive disintegration process, the data $n^*(d)$ are not going to be an absolutely accurate representation of the probabilities $p(b, d)$; i.e., there is Poisson noise in the data elements. The maximum of L will result in a varying degree of “reconciliation” between the measurement data n^* and resulting image projection λ^* .

Early in our study of the MLE, first in the PEBA geometry and later with a true PET configuration (ECAT-III 512-detector ring at U.C.L.A. [5]), we realized that the onset of the image breakup phenomenon occurred when the estimate λ^* and the data

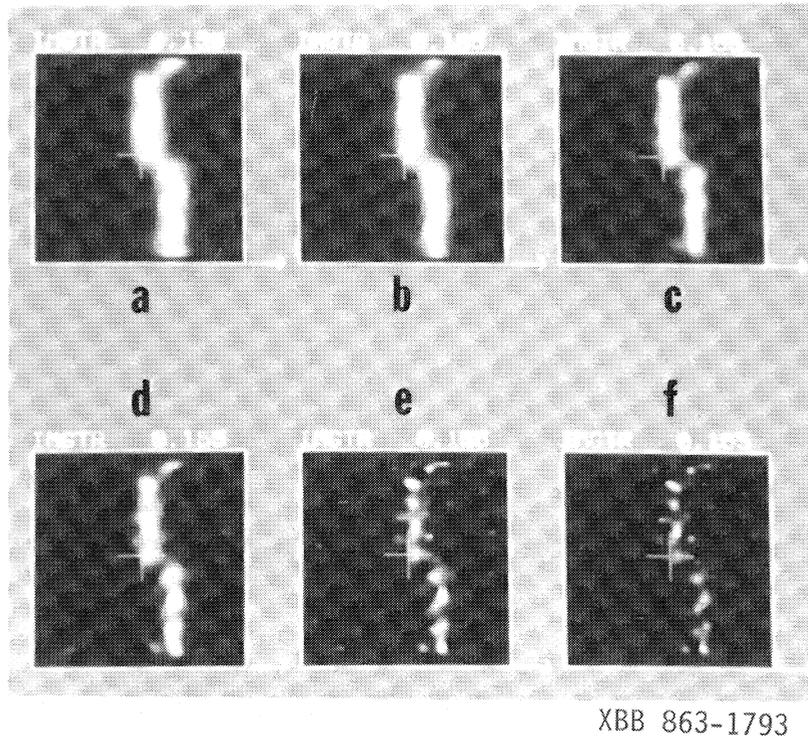


Figure 3. Reconstructed images from a line source obtained by the injection of a beam of Ne-19 into a plastic cylinder, with a partial obstruction of a Ca compound, causing a step. (a) image obtained by pseudo-inverse method; (b) MLE, 3 iterations; (c) MLE, 9 iterations; (d) MLE, 21 iterations; (e) MLE, 100 iterations; (f) MLE, 200 iterations.

n^* for tubes with high number of counts were being reconciled rather closely [3, 6, 7]. The nature of Eq.2 is such that tubes with low counts are reconciled very early in the iterative process (notice that $|dL/d\lambda^*|$ is higher at low values of n^* than at high values for equal errors $|\lambda^* - n^*|$), leaving the detailed reconciliation of tubes with high counts for the later iterations.

It then appeared fair to ask the question: Is it physically correct that the estimates λ^* should become so close to the tube data n^* ? A simple hypothetical example will help clarify the nature of the question.

Consider Figure 4a to be a graph of one projection from a simple PET measurement. The true, unmeasurable average of counts in each tube is represented by the solid line with squares. One standard deviation is shown by the interval between the two solid lines. The crosses indicate one realization of the Poisson processes with the means indicated in each tube. They form a set of noisy data that can result from one single measurement by a PET instrument.

An MLE reconstruction, when carried to a large number of iterations, will try to find an image whose projection is as close as possible to the initial data, as shown hypothetically by the solid line in Figure 4b. We then ask the question: Could an image whose projection is shown in Figure 4b have generated the original noisy data by

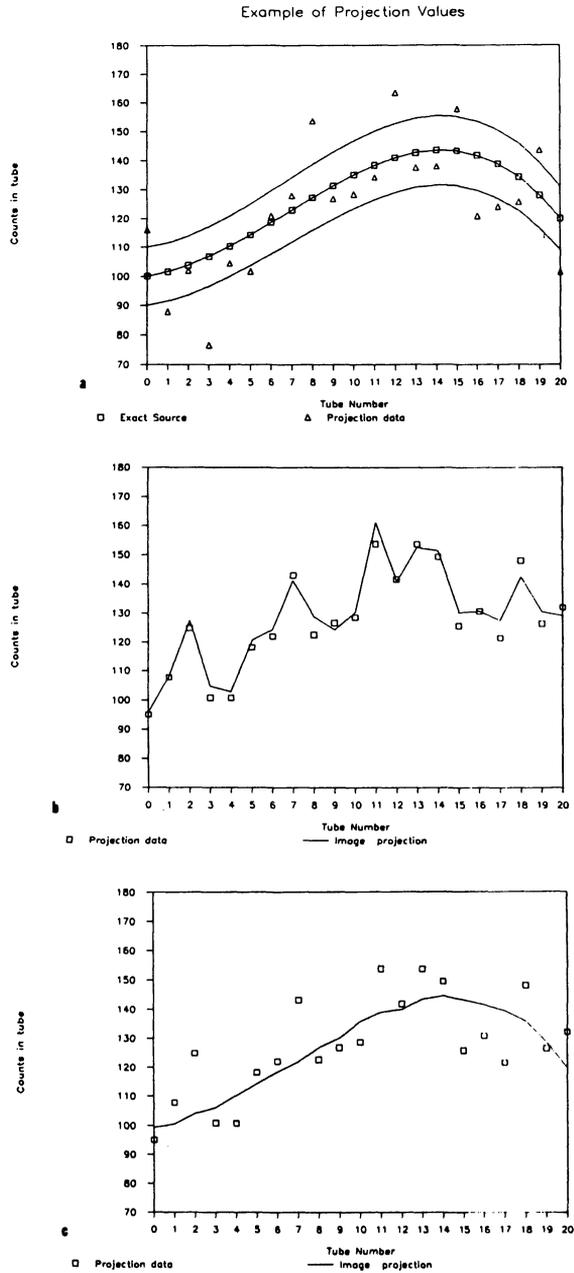


Figure 4. Hypothetical example of the process of data generation and image reconstruction in projection data space, illustrating the question asked of whether a reconstruction is physically consistent with the measured data.

a Poisson process?

If we calculate the mean standard error between the solid line of Figure 4b and the corresponding measurement points, we find that number to be ~ 1 count for this

hypothetical experiment. For means in the order of 100 counts and a Poisson process, that mean standard error should be in the vicinity of 10 counts. We would conclude that the image that gave rise to the projection shown in Figure 4b cannot have caused the experimental data, although it had a high likelihood.

On the other hand, an image whose projection is shown by the solid line of Figure 4c will be of lower likelihood, but would be closer to the true image and have a mean standard error of ~ 10 counts, in agreement with the physical process that generated the data. It is the function of a reconstruction algorithm to obtain that line from the available experimental points.

3. The concept of image feasibility

After having asked a physically meaningful question, we have proceeded to define a feasible image as an image that, if it were a true radiation field in the image plane, could have caused the experimental data obtained by the physical process that governs the data acquisition mechanism.

Definition. The image $\lambda(1), \lambda(2), \dots, \lambda(B)$ is said to be a feasible image with respect to data $n^*(1), n^*(2), \dots, n^*(D)$ if and only if the statistical hypothesis that $n^*(1), n^*(2), \dots, n^*(D)$ are a Poisson sample with the means $\lambda^*(1), \lambda^*(2), \dots, \lambda^*(D)$, respectively, can be accepted (not rejected).

We can then ask: How large is the set of feasible images and where are they with respect to the maximum likelihood image?

Let us consider the measurement data (projection) space and imagine that any set of projections $\lambda^*(d), d = 1, \dots, D$ is achievable; that is, it can be obtained from a distribution of intensities in the pixel space (we will reexamine this assumption below). If any point is achievable, then the maximum likelihood solution is such that $\lambda^*(d) = n^*(d)$ for all $d = 1, \dots, D$.

We assert that all feasible images belong to a volume resembling an ellipsoidal shell that surrounds the point with $\lambda^*(d) = n^*(d)$. Indeed, if an image is feasible, the mean value of the quantity:

$$\frac{[n^*(d) - \lambda^*(d)]^2}{\lambda^*(d)}$$

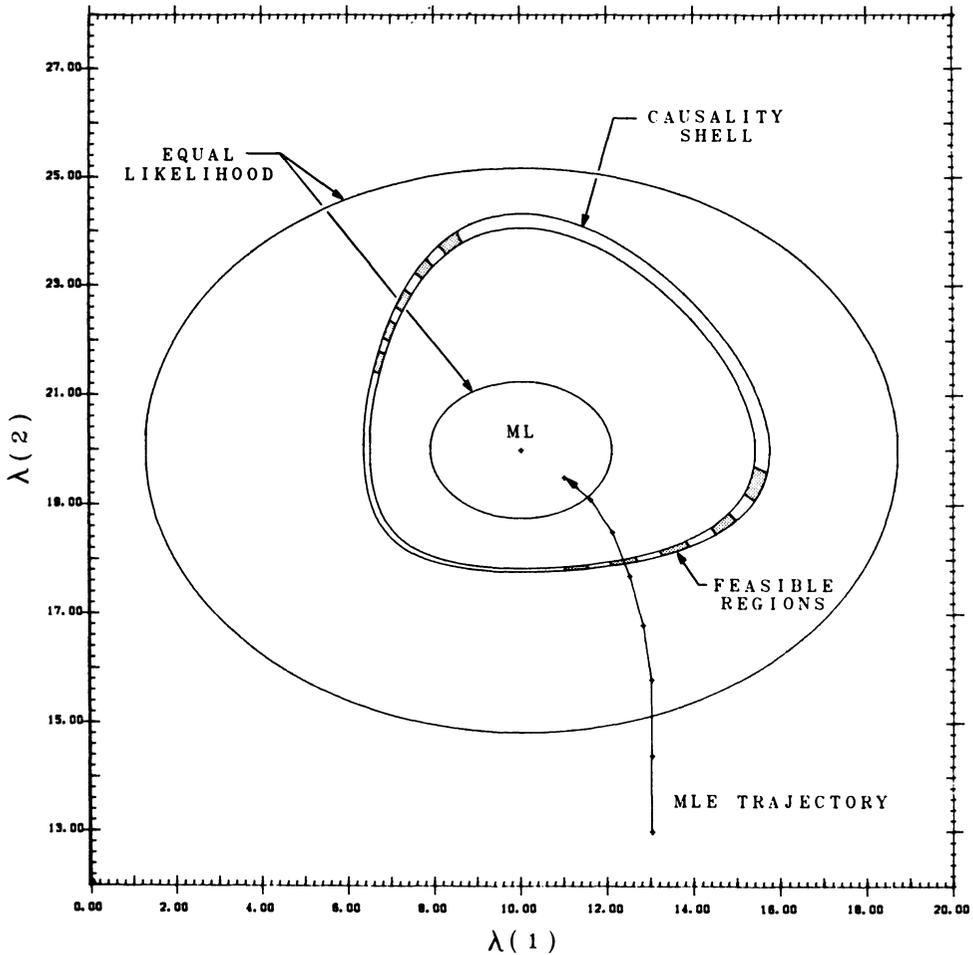
equals 1 for the Poisson distribution, because the numerator is the variance, while the denominator is the mean. Therefore, for a feasible image:

$$\sum_{d=1}^D \frac{[n^*(d) - \lambda^*(d)]^2}{\lambda^*(d)} = D \pm \varepsilon \quad (3)$$

All points satisfying (3) form the space between two ellipsoid-like surfaces. They would be exact ellipsoids if the denominators were constants. The parameter ε reflects the expectation that the region of feasibility has a certain width.

For the case of any set of projections being achievable, Figure 5 shows ellipses of equal likelihood centered at the point of maximum likelihood, which for that case is the same as the measurement vector $n^*(1) = 10, n^*(2) = 20$. It also shows the region of feasibility for the simple case of $D = 2$. A value $\varepsilon = 0.1$ was used in the example. The

measurement point n^* is enclosed by the region of feasibility (3) and, in general, the maximum likelihood point can be expected to be also enclosed. It is possible to show, however, that in very pathological cases the latter may not be true.



XBL 887-2476

Figure 5. Equal likelihood ellipses and feasibility shell with achievable regions for a case with two-dimensional projection data space.

A closer look at the reasoning above shows that while all feasible images must belong to the region (3), not all points of the region correspond to a feasible image. This is because (3) was derived using only the property of the Poisson distribution that its variance equals its mean. That condition is not sufficient for a distribution to be Poisson. Hence, the true feasible images in the sense of the definition above form subset of the region (3) which contains some gaps. The gaps are large, as can be deduced from the following reasoning: A point may satisfy (3) even if $n^*(d) > \lambda^*(d)$ for all $d = 1, \dots, D$. But $n^*(d)$ are measurements and $\lambda^*(d)$ are their means; therefore some $n^*(d)$ should be larger than $\lambda^*(d)$, while others should be smaller. The simple sign test [8] shows that the hypothesis that $\lambda^*(d)$ are the means of $n^*(d)$ should be rejected if for $D = 1000$, for example, there are less than 450 measurements below their averages or if

there are less than 450 measurements above their averages. This test eliminates most of the points in the region (3).

Now let us return to the assumption that all points in the projection space are achievable. If we consider that the matrix $p(b, d)$ maps a B -dimensional space of pixel intensities into a D -dimensional space of measurements, D being typically greater than B , it is clear that not all sets of projections $\lambda^*(d)$ are achievable from a distribution of intensities in the pixels. This fact further limits the region of feasibility (3) by placing holes into the remaining parts of the region. It should be pointed out, however, that the trajectory of likelihood in projection space obtained during the iterative reconstruction process, if it crosses the feasibility region, will escape those holes, since the result of the iterations is an image whose projections have to be in the solid part of the feasibility region.

Figure 5 shows hypothetical achievable regions and a trajectory of likelihood points for a reconstruction starting from the arbitrary vector $\lambda^*(1) = \lambda^*(2) = 13$. The trajectory passes through the feasible region and continues towards the maximum likelihood.

On a historical note, the idea of feasibility has been known in a different context in radio astronomy prior to our present work. Thus, it was used by Skilling and Bryan [9] in 1984. However, Skilling and Bryan considered Gaussian rather than Poisson data. As a result, their feasibility region is a regular ellipsoid rather than a shell, as it is in our case. More importantly, Skilling and Bryan's definition of feasibility is based on the Gaussian version of Eq. 3 which, as we know, provides only a necessary but not sufficient condition for causality. Several variations on the same theme in case of Gaussian data have been reported by other authors working in radio astronomy; see Ables [10], Gull and Daniels [11] and Narayan and Nityananda [12]. The fact that not any point of the Gaussian version of Eq. 3 statistically fits the data was noted by Reiter and Pfeleiderer [13].

4. Limited test for image feasibility

In order to answer the question of whether the image resulting from a particular iteration of the MLE procedure, or any other reconstruction method for that matter, can be physically meaningful in the sense indicated above, we have developed a limited test for feasibility. We test collectively whether each of the projection data points $n^*(d)$ can be an instance of a Poisson variable with the corresponding $\lambda^*(d)$ as mean. We assume independence between each and all the $\lambda^*(d)$ and $n^*(d)$; i.e., we disregard the fact that the tested $\lambda^*(d)$ have been obtained from the ensemble of $n^*(d)$ by a specific iterative reconstruction method. With this simplifying assumption it is possible to define a simple hypothesis test with a well defined number of degrees of freedom and, consequently, with known confidence levels for our results, regardless of the method of reconstruction. The test also contains a vastly reduced number of degrees of freedom when compared to the fitting problem and is, therefore, much less strict than a likelihood ratio test, for example. Full details of the development of the test are given by Veklerov and Llacer [14]. We only summarize here the main ideas.

The first step of the test consists of scaling the differences $\lambda^*(d) - n^*(d)$ for each data pair to a new variable x which is uniformly distributed between 0 and 1 if $n^*(d)$ is a realization of a Poisson process with mean $\lambda^*(d)$. Next, a histogram with N bins is generated with the values x for all the data pairs. We conclude by testing the hypothesis that x is uniformly distributed between 0 and 1 by Pearson's procedure with $N - 1$

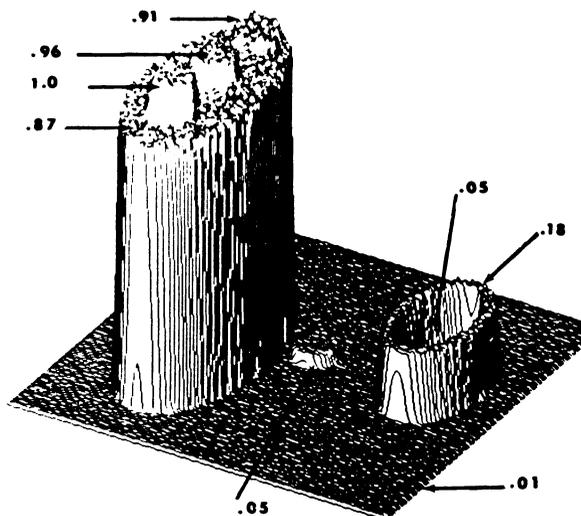
degrees of freedom. The histogram testing function H is defined as:

$$H = \sum_{i=1}^N \frac{(h_i - D/N)^2}{D/N} \quad (4)$$

where h_i is the frequency of x falling in bin i and D/N is the expected frequency if x is uniformly distributed.

5. Experimental results

We have carried out several groups of MLE reconstruction experiments with computer generated data and with real data from the ECAT-III tomograph of U.C.L.A. [5] for the experiments with computer generated data we first used the phantom shown in Figure 6, with the relative amounts of radioactivity indicated. The generation of projection data was carried out in two steps: (1) generation of the "source" image by a random process based on the probabilities indicated in Figure 6 with a finite number of counts, and (2) generation of the projection data by taking each count in the source image and using a second random process based on the transition probabilities $p(b, d)$ to determine which detector tube would receive the count. The transition matrix used both in the data generation and in the reconstruction was calculated with the parameters of the ECAT-III tomograph by the prescription of Shepp and Vardi [2]. All the reconstruction experiments begin with a uniform image as the first guess.



XBB 860-10391-A

Figure 6. Source image of a "liver and heart" phantom with 32 million counts obtained by a random process with probabilities corresponding to the relative activities indicated. Image discretized in 128×128 pixels in all sources and reconstructions.

Figure 7 shows cuts through the source image, a filtered back-projection reconstruction (Shepp-Logan filter) and MLE reconstructions at 9, 32 and 200 iterations for the case with 2 million (2M) counts in the source image. Figure 8 shows the values of the hypothesis testing function H of Eq. 3 obtained during the iterative MLE process for the 2M case, as well as those for 8 million (8M) and 32 million (32M) counts. The horizontal lines correspond to critical values of 0.1 and 0.01 for the test. During the early iterations the histograms of x are high at the extreme bins and low at the center, indicating excessive distance between values of H the histograms are rather flat, while at high number of iterations the histograms of x are bunched in the center, indicating that the projection values λ^* are too close to the data n^* for the image being tested to be causal.

The images with acceptable values of H (below ~ 40) invariably exhibit a good compromise between sharpness and noise in the regions of high activity. The image breakup phenomenon or noise artifact has not yet set in. It is evident from Figure 8 that images with more counts take more iterations to get to their acceptable range. The final results are correspondingly better.

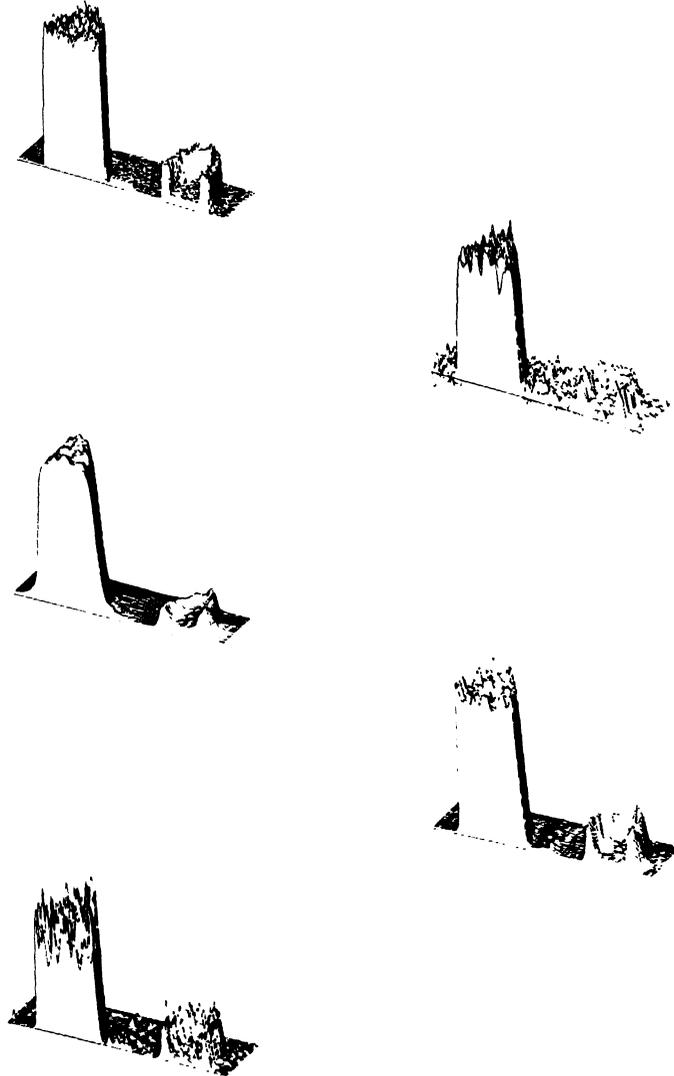
We have tested a second computer generated phantom with considerable more image detail than the first one, mimicking in some way complex brain structures. Figure 9 shows a 2M source image, a filtered back-projection (Shepp-Logan filter) reconstruction and MLE reconstruction results at 21, 32, 45 and 100 iterations. The results of the hypothesis testing are shown in Figure 10, indicating the first acceptable images to be in the vicinity of iteration 40. Indeed, the reconstructions of iterations 32 and 45 appear adequately sharp and clean.

The reconstruction experiments with real PET data have presented some difficulties: although the resulting images exhibited the expected behavior as the number of iteration increased, the values of H never became low enough to lead to accepting the images as feasible. A detailed investigation has shown us that the success of the feasibility test, as presently implemented, is highly dependent on knowing the transition matrix $p(b, d)$ with high accuracy. With real tomographs, this knowledge cannot be expected and we have had to relax the hypothesis testing to include a parameter that is proportional to the estimated error with which we know the values of $p(b, d)$. We are now preparing to report our new results [15].

6. The stopping rule and Bayesian reconstruction

At this point in the development of our ideas it is important to pause and attempt to understand the meaning of using a rule that stops the estimation of an image before the maximum of the likelihood estimator for that image is reached. We will be helped in that understanding by considering Bayesian reconstruction.

First, we would like to point out that it appears self evident that an image presented to a radiologist for a clinical evaluation should be feasible as defined above; i.e., it should be possible for the image, if it were a true radiation source, to have generated the projection data. We feel that a radiologist could not have confidence in the results of a reconstruction if the contrary were true. Second, we have shown that an MLE reconstruction, starting from a uniform field as a first image guess, passes through a region of feasibility as the iterative process progresses towards maximum likelihood and we have devised a limited test to detect the passage through that region of feasibility. If the iterative process is allowed to continue past feasibility toward maximum likelihood,



XBB 871-535

Figure 7. Cuts through source and reconstructed images. (a) source image with 2 million counts; (b) reconstruction by filtered back-projection (Shepp-Logan filter); (c) MLE reconstruction, 9 iterations; (d) MLE, 32 iterations; (e) MLE, 200 iterations.

the image becomes non-feasible and should be rejected.

There seems to be something wrong about using a reconstruction algorithm based

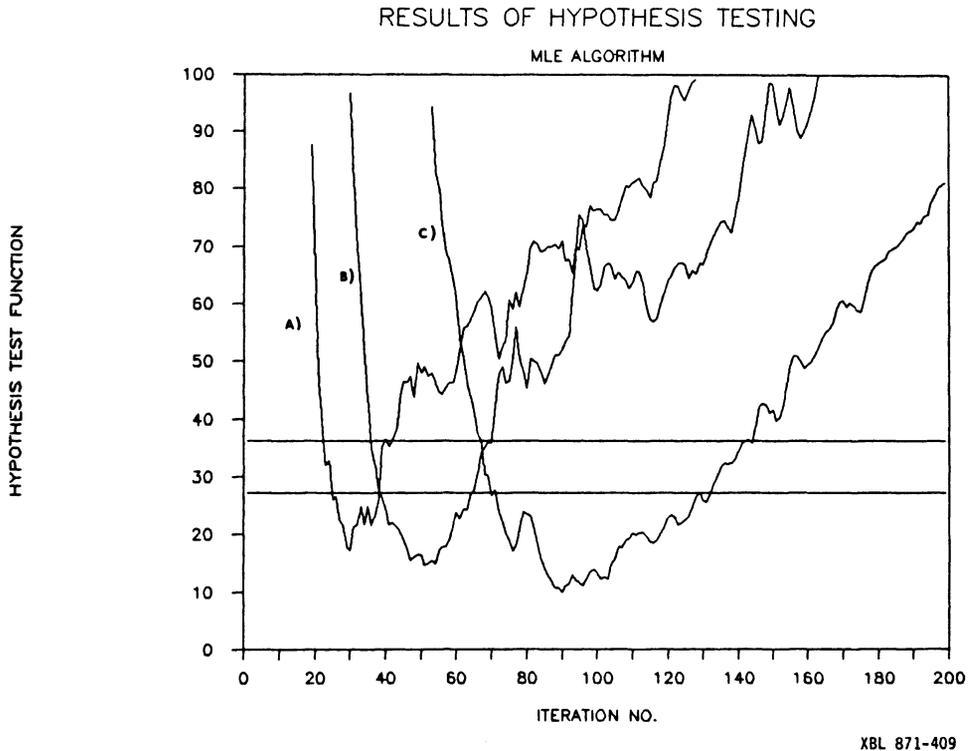


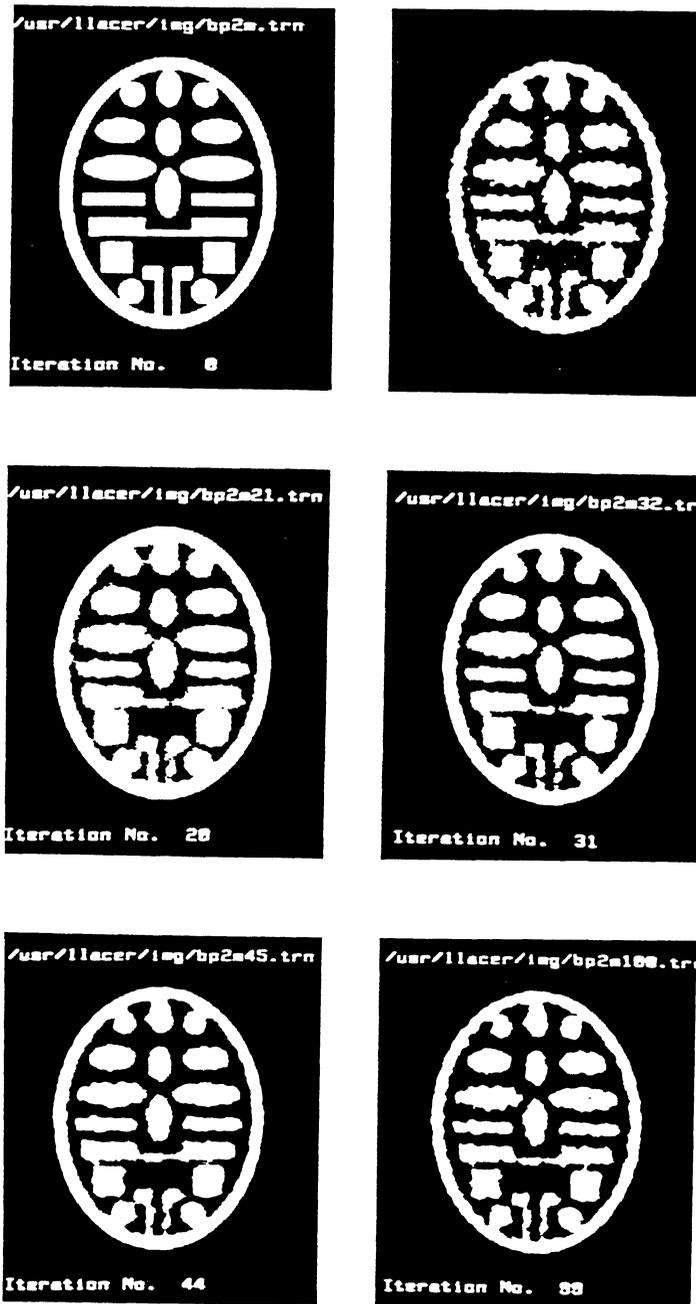
Figure 8. Value of the hypothesis testing function H vs. iteration number for MLE reconstructions of the source distribution of Figure 6. Curve (a) is for 2 million counts in the source image; (b) 8 million; (c) 32 million counts. The line at a value of 36.2 is the limit below which the probability of accepting an image when it should be rejected is 0.1. The line at 27.2 is the corresponding limit for a probability of 0.01. The confidence levels are based on the assumption that projection data and reconstructed images are independent of each other.

on maximizing the likelihood and stopping it before that likelihood is maximized. The main objection is that the MLE is known to lead to a unique solution and stopping before the maximum could result in images that are different depending on the starting point.

The drawing of Figure 5, describing qualitatively the shape of the feasibility region, appears closely related to that of Figure 11 (taken from Hanson [16]) describing a Maximum a Posteriori (MAP) reconstruction, in which a solution is looked for that maximizes the Bayes criterion:

$$p(f|g) = p(g|f)p(f)/p(g) \tag{5}$$

for $g = Hf + n$. For the imaging problem, f is an image, H is a mapping into the measurement data space, n is the noise in the image and g is the result of a measurement. In a Bayesian reconstruction we would look for the image f which has the maximum probability given the data g ; i.e., we would try to maximize $p(f|g)$. In order to do that we need to maximize the product $p(g|f) p(f)$. The MLE solution addresses the



XBB 878-6851

Figure 9. Left to right, top to bottom: source image with 2 million counts, filtered back-projection (Shepp-Logan filter) and MLE reconstructions for 21, 32, 45 and 100 iterations, presented with 16 levels of gray. The “best” image is between iterations 32 and 45, determined by the value of the hypothesis testing function H in Figure 10.

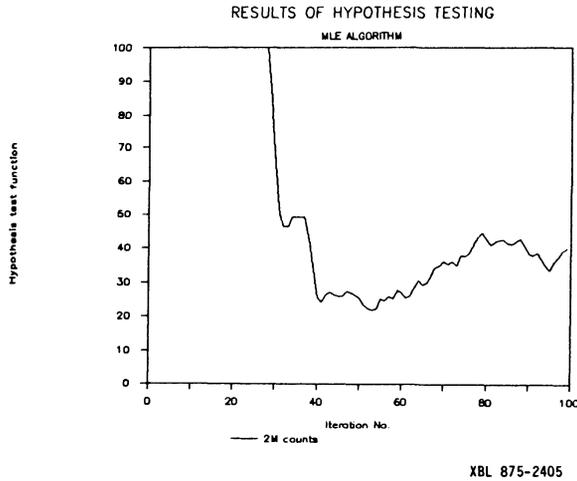


Figure 10. Hypothesis testing function for the reconstruction of the phantom of Figure 1, with 2 million counts in the source.

maximization of $p(g|f)$ only; i.e., it tries to maximize the probability of g being the data that would result from a given image and invariably leads to an image in which Hf is too close to g for acceptability in terms of the physical data generation process.

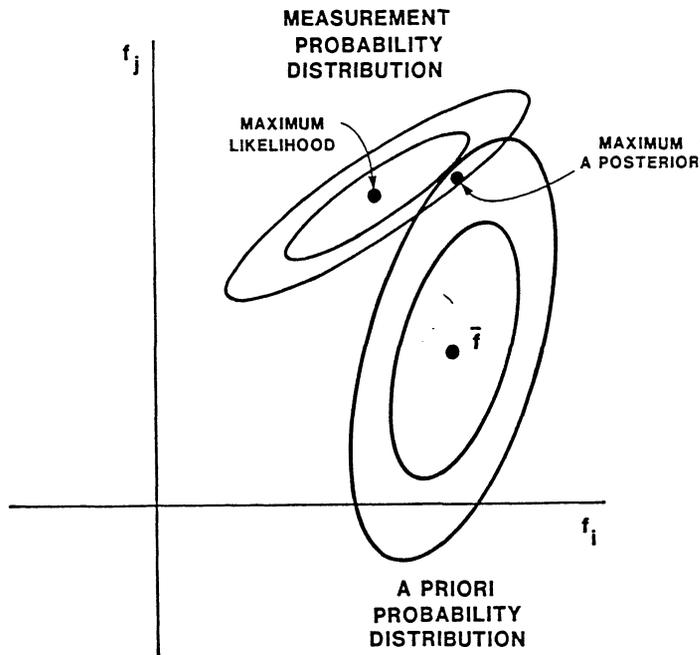


Figure 11. Contour plots of the probability distributions associated with prior information and the available measurements displayed as a function of two components of the image projection vector, for Gaussian distributions. (From Hanson, Ref. 16).

The prior probability $p(f)$ is difficult to obtain in a general way. We will mention here five attempts at defining $p(f)$, without claiming completeness. Miller and Snyder [4, 17], have used sieves as a method of regularizing the MLE iterative procedure, with images that do not exhibit the noise artifact and with adequate sharpness. Geman and McClure [18] use a prior that requires a certain degree of smoothness in image subregions. Levitan and Herman [19] have used a smoothed, non-negative filtered back-projection reconstructed image to require that the MLE reconstruction does not deviate excessively from it. Smith, Barrett and Paxman [20] use the method of simulated annealing to, in effect, define a certain $p(f)$ that allows for a number of pixel counts that do not maximize the likelihood function. Finally, Skilling and Gull [21] have shown that, if one accepts three fundamental axioms, the prior $p(f)$ has to be a measure of the entropy in the image. Evidently, more work needs to be done in this area both by evaluating the presently proposed prior probability functions in terms of image quality and feasibility and by developing better priors, if needed.

The use of the stopping rule makes use of prior knowledge about the statistics that govern the physical process of a tomographic measurement but does not incorporate that knowledge into a prior probability $p(f)$.

7. Conclusion

In this paper we have shown the motivation that we have had to study the behavior of the MLE method of image reconstruction in emission tomography and described our observations on its behavior. From a physicist's point of view, we have shown that the MLE alone does not lead to acceptable images because its reconstructions violate a physical principle, which we have called "feasibility". This principle is in the form of a priori knowledge that does not lend itself to casting into a prior probability to regularize the MLE reconstruction. Instead, we have used that knowledge to stop the MLE iterative procedure as soon as the feasibility condition is successfully met. We feel that, at this time, the use of the stopping rule may be the most economical (in terms of computing time) and simplest way to use a statistically based reconstruction method in a practical medical environment. Further study of the characteristics of images obtained by the stopping rule starting from different initial estimates in comparison with the presently proposed regularization methods will help clarify its role in practical emission tomography.

Acknowledgements

The authors would like to thank L. A. Shepp for his encouragement of our work and David J. Strauss and Noel A.C. Cressie for their constructive criticism of our stopping rule. This work was supported in part by the National Cancer Institute (CA-39501) and by the U.S. Department of Energy under contract No. DE-AC03-76SF00098.

References

- [1] Llacer, J., Chatterjee, A., Alpen, E. L., Saunder, W., Andreae, S., & Jackson, H. C. (1984). Imaging by injection of accelerated radioactive particle beams. *IEEE Transactions on Medical Imaging* **MI-3** (2), 80–90.
- [2] Shepp, L. A., & Vardi, Y. (1982). Maximum likelihood reconstruction for emission tomography. *IEEE Transactions on Medical Imaging* **MI-1**, 113–122.
- [3] Llacer, J., Andreae, S., & Chatterjee, A. (1986). *On the applicability of the maximum likelihood estimator algorithm for image recovery in accelerated positron emitter beam injection*. SPIE Proc. International Workshop on Phys. and Eng. of Comp. Multidimensional Imaging and Proc., Newport Beach, CA, 671.
- [4] Snyder, D. L., & Miller, M. I. (1985). The use of sieves to stabilize images produced with the EM algorithm for emission tomography. *IEEE Transactions on Nuclear Science* **NS-32** (5), 3864–3872.
- [5] Hoffman, E. J., Ricci, A. R., van der Stee, L. M. A. M., & Phelps, M. E. (1983). ECAT-III basic design considerations. *IEEE Transactions on Nuclear Science* **NS-30** (1), 729–733.
- [6] Llacer, J., Veklerov, E., & Hoffman, E. J. (1987). On the convergence of the maximum likelihood estimator method of tomographic image reconstruction. *SPIE Proceedings of the Conference on Medical Imaging* **767**, 70–76.
- [7] Llacer, L., & Veklerov, E. (1987). *The high sensitivity of the maximum likelihood estimator method of tomographic image reconstruction*. Int'l. Symp. of Computer Assisted Radiology (CAR '87), Springer, 108–112.
- [8] Freund, J. (1971). *Mathematical Statistics*. Prentice-Hall, 343–344.
- [9] Skilling, J., & Bryan, R. K. (1984). Maximum entropy image reconstruction: general algorithm. *Mon. Not. R. Astr. Soc.* **211**, 111–124.
- [10] Ables, J. G. (1974). Maximum entropy spectral analysis. *Astron. Astrophys. Suppl.* **15**, 383–393.
- [11] Gull, S. F., & Daniell, G. J. (1978). Image reconstruction from incomplete and noisy data. *Nature* **272**, 686–690.
- [12] Narayan, R., & Nityananda, R. (1986). Maximum entropy image restoration in astronomy. *Ann. Rev. Astron. Astrophys.* **24**, 127–170.
- [13] Reiter, J., & Pfeiderer, J. (1986). Improvement of MEM-deconvolution by an additional constraint. *Astron. Astrophys.* **166**, 381–392.
- [14] Veklerov, E., & Llacer, J. (). Stopping rule for the MLE algorithm based on statistical hypothesis testing. *IEEE Transactions on Medical Imaging* **MI-6**, 313–319.
- [15] Llacer, J., & Veklerov, E. (). *Causal images and practical stopping rules in iterative image reconstruction*, to be published.
- [16] Hanson, K. M. (1987). *Image Recovery, Theory and Application*. Ch. 3, Bayesian and related methods in image reconstruction from incomplete data, Academic Press, 79–125.

- [17] Miller, M. I., & Snyder, D. L. (1987). The role of likelihood and entropy in incomplete-data problems: Applications to estimating point-process intensities and toeplitz constrained covariances. *Proceedings of the IEEE* **75** (7), 892–907.
- [18] Geman, S., & McClure, D. E. (1987). *Statistical methods for tomographic image reconstruction*. Proceedings of the 46th Session of ISI, Bull. of ISI, 52.
- [19] Levitan, E., & Herman, G. T. (1986). *A maximum a posteriori probability expectation maximization algorithm for image reconstruction in emission tomography*. Medical Image Processing Group Technical Report No. MIPG115, Dept. of Radiology, Univ. of Pennsylvania.
- [20] Smith, W. E., Barrett, H. H., & Paxman, R. G. (1983). Reconstruction of objects from coded images by simulated annealing. *Optical Letters* **8**, 199–201.
- [21] Skilling, & Gull (1988). personal communication, publication in progress.