

Methods of Detecting Objects in Photon-Limited Images

Ahmad Abu-Naser

Department of Electrical and Computer Engineering

Illinois Institute of Technology

Chicago, IL 60616, USA

Nikolas P. Galatsanos

Department of Computer Science

University of Ioannina

Ioannina, Greece 45110

Miles N. Wernick

Department of Electrical and Computer Engineering, and Medical Imaging Research Center,

Illinois Institute of Technology

Chicago, IL 60616, USA

ABSTRACT

Herein we investigate the problem of detecting and localizing a known signal in a photon-limited image, where Poisson noise is the dominant source of image degradation. For this purpose we developed and evaluated three new algorithms. The first two are based on the impulse restoration (IR) principle and the third is based on the generalized likelihood ratio test (GLRT). In the IR approach, the problem is formulated as one of restoring a delta function at the location of the desired object. In the GLRT approach, which is a well-known variation on the optimal likelihood ratio test, the problem is formulated as a hypothesis testing problem, in which the unknown background intensity of the image and the intensity scale of the object are obtained by maximum-likelihood estimation. We used Monte Carlo simulations and localization receiver operating characteristic curves (LROC) to evaluate the proposed algorithms quantitatively. LROC curves demonstrate the ability of an algorithm to detect and locate objects in a scene correctly. Our simulations demonstrate that the GLRT approach is superior to all other tested algorithms.

1. Introduction

Template matching, also known as signal-known-exactly (SKE) detection is the problem of detecting and localizing a known template (signal) image within a cluttered background image in the presence of noise [1]. This is one of the most basic and widely studied tasks in image processing, with many applications in fields such as object recognition, novelty detection, motion estimation, industrial inspection, medical imaging and sensor fusion [2].

Most research on SKE detection has focused on the Gaussian noise model. In this paper, we propose algorithms for signal detection in photon-limited images, where the dominant source of noise is caused by lack of light, and is correctly modeled by a Poisson law [3]-[5]. Important photon-limited

applications are night vision and medical imaging (such as nuclear medicine [6]) where the number of collected photons is low. The same theory applies to low-dose electron microscopy, where the limiting factor is the low number of electrons used to form the image.

Object recognition in photon-limited imagery can be very challenging for a human observer, but a numerical detector, informed by the Poisson noise model, can be surprisingly effective [3],[4]. Template matching in photon-limited imaging was first described in [3], in which a simple correlation detector was used. The correlation detector lends itself to real-time computation; however, it is by no means an optimal detector for the Poisson noise case. An approximate likelihood-ratio test (LRT) for the photon-limited image classification problem was derived in Ref. [4] for the simplified case in which the template resides in a black background. In Ref. [5] the detection problem was investigated for photon-limited images passed through a linear system.

In this paper, we develop, evaluate, and compare three algorithms for SKE detection in Poisson noise: two based on the concept of *impulse restoration* (IR), and one consisting of a generalized likelihood ratio test (GLRT). We evaluate these methods using localization receiver-operating-characteristic (LROC) analysis, which captures both the detection and localization performance of the algorithms.

The IR approach is a less-traditional method, in which the detection problem is framed as one of restoring a delta function that indicates the spatial location of the detected signal. This can be viewed as an improvement on classical template-matching methods, such as the matched filter, which aim to produce a peak in their output at the detected signal location. The potential benefit of IR methods is that they produce a sharp delta function at the desired position. The first explicit application of the IR principle was reported in Refs. [7] and [8]. More recently, the relationship between linear-minimum-

mean-square-error (LMMSE) IR and object recognition was recognized and IR object-recognition filters were proposed [9]-[11]. IR-based methods that are robust to orientation and scale uncertainties were presented in Refs. [12] and [13]. A potential advantage of the IR viewpoint is that it permits the substantial base of knowledge gained in the image-restoration field to be brought to bear on the template-matching problem. This observation was exploited for the Gaussian noise case in [14] and [15] and for the Poisson noise case in [16]. The work presented in this paper expands on the discussion of the IR methods described in Ref. [16], and introduces the GLRT, which we find outperforms the IR methods.

The GLRT is a standard and widely used variation on the likelihood ratio test (LRT), which is a pillar of basic decision theory [1]. Many optimal decision strategies—e.g., Bayes risk, Neyman-Pearson, maximum-likelihood—have the form of an LRT. In a GLRT, the unknown parameters of the required likelihood functions are replaced by maximum-likelihood (ML) estimates of these parameters. Many well-known hypothesis-testing techniques, including the classical Student t -test, are examples of GLRTs.

The rest of this paper is organized as follows. In Section 2 we present the IR formulation for the photon-limited noise case, and we solve the problem in two ways: 1) maximum-likelihood (ML) estimation and 2) maximum *a posteriori* (MAP) estimation. In Section 3 a GLRT is developed, which accounts in a simple way for local variations in background intensity and intensity scaling of the template, and thus significantly boosts performance over a previous LRT implementation [4]. In Section 4 we present quantitative performance evaluations based on localization ROC (LROC) curves, which plot the probability of detection and correct localization of an object versus the false-alarm probability. Finally, in Section 5 we present our conclusions.

2. Impulse Restoration (IR) Methods

We begin by reviewing the formulation of object recognition as an impulse-restoration problem. Let the observed $M_0 \times N_0$ photon-limited image be denoted by $g(\mathbf{m})$, where $\mathbf{m} = (m, n)$ denotes the discrete spatial coordinates of a pixel ($m = 0, \dots, M_0 - 1$, $n = 0, \dots, N_0 - 1$), and let \mathbf{m}_p , $p = 1, \dots, P$, denote the set of unknown positions within the image at which the signal is located. In photon-limited imaging, the image obeys the Poisson law, i.e., $g(\mathbf{m}) \sim \text{Poisson}(E[g(\mathbf{m})])$. Here, we model the image as a scene containing one or more instances of a signal object $f(\mathbf{m})$ so that the expected value of the image is:

$$E[g(\mathbf{m})] = \sum_{p=1}^P f(\mathbf{m} - \mathbf{m}_p), \quad (1)$$

or, equivalently,

$$E[g(\mathbf{m})] = f(\mathbf{m}) * \tilde{\delta}(\mathbf{m}), \quad (2)$$

where

$$\tilde{\delta}(\mathbf{m}) = \sum_{p=1}^P \delta(\mathbf{m} - \mathbf{m}_p). \quad (3)$$

Here, $\delta(\mathbf{m})$ represents a discrete impulse function and $*$ denotes convolution.

For notational convenience, we will use the following matrix-vector notation to describe the imaging model in Eq. (2):

$$E[\mathbf{g}] = \mathbf{F}\boldsymbol{\delta}, \quad (4)$$

where \mathbf{F} is a doubly block-circulant matrix composed from the elements of $f(\mathbf{m})$ in such a way that $\mathbf{F}\boldsymbol{\delta}$ represents circular convolution of the zero-padded versions of f and $\tilde{\delta}$. This representation is equivalent to Eq. (2) if the images are zero-padded to dimension $M_1 \times N_1 (= N)$ and arranged as $N \times 1$

vectors by using lexicographic ordering. Thus, in photon-limited imaging, the observed image can be modeled as $\mathbf{g} \sim \text{Poisson}(\mathbf{F}\boldsymbol{\delta})$. Such a statistical relationship is commonly known as a *Poisson linear model*.

In the IR formulation, the detection problem reduces to one of deconvolving the observed image \mathbf{g} to obtain the indicator image $\boldsymbol{\delta}$, which contains an impulse at each location where the known signal is present. In the following sections, we describe a maximum-likelihood (ML) and a maximum *a posteriori* (MAP) method of estimating $\boldsymbol{\delta}$.

2.1 ML Solution

We model the observed photon-limited image by the following Poisson likelihood function [6]:

$$p(\mathbf{g} | \boldsymbol{\delta}) = \prod_{i=0}^{N-1} \frac{[\mathbf{F}\boldsymbol{\delta}]_i^{g_i} e^{-[\mathbf{F}\boldsymbol{\delta}]_i}}{g_i!}, \quad (5)$$

where $[\mathbf{F}\boldsymbol{\delta}]_i$ is the i^{th} element of the vector $\mathbf{F}\boldsymbol{\delta}$. The aim of impulse restoration is to estimate the indicator image $\boldsymbol{\delta}$ from the observed image \mathbf{g} . Here we pursue a maximum-likelihood (ML) estimation strategy [1] to obtain the desired estimate $\hat{\boldsymbol{\delta}}$, i.e.,

$$\hat{\boldsymbol{\delta}} = \arg \max_{\boldsymbol{\delta}} \{\ln p(\mathbf{g} | \boldsymbol{\delta})\}. \quad (6)$$

Taking the natural logarithm of (5) we obtain the required log-likelihood function

$$\ln [p(\mathbf{g} | \boldsymbol{\delta})] = \sum_{i=0}^{N-1} [\mathbf{F}\boldsymbol{\delta}]_i + g_i \ln ([\mathbf{F}\boldsymbol{\delta}]_i) - \ln (g_i!). \quad (7)$$

From Eq. (7) it is clear that a closed-form solution for the ML estimate cannot be found; therefore, we employ the widely used the expectation-maximization (EM) algorithm [17] to obtain the solution iteratively.

In the present problem we define the complete data z_{ij} as the number of photoevents in g_i which were emitted, transmitted, or reflected by a region of the object that maps onto image location j . Under this definition, the functional relationship between the complete and incomplete data is

$$g_i = \sum_{j=0}^{N-1} z_{ij}, \quad i = 0, 1, \dots, N-1. \quad (8)$$

where

$$z_{ij} = F_{ij} \delta_j. \quad (9)$$

A similar relationship between complete- and incomplete-data sets has been used to solve the reconstruction problem in emission tomography, which leads to the following well-known iterative formula for the EM algorithm [18]-[20]:

$$\hat{\delta}_j^{(l+1)} = \frac{\hat{\delta}_j^{(l)}}{\sum_{i=0}^{N-1} F_{ij}} \sum_{i=0}^{N-1} \left\{ \frac{F_{ij} g_i}{\sum_{k=0}^{N-1} F_{ik} \hat{\delta}_k^{(l)}} \right\}, \quad j = 0, \dots, N-1. \quad (10)$$

2.2 MAP Solution

Because the ML principle has no intrinsic form of regularization, it can produce undesirable results if used without modification. Thus, it is common instead to use a maximum *a posteriori* (MAP) estimation procedure, which regularizes the solution through the use of a prior $p(\boldsymbol{\delta})$ on the signal.

The MAP estimate is obtained by the following maximization:

$$\begin{aligned} \hat{\boldsymbol{\delta}} &= \arg \max_{\boldsymbol{\delta}} \{ \ln p(\boldsymbol{\delta} | \mathbf{g}) \} \\ &= \arg \max_{\boldsymbol{\delta}} \{ \ln p(\mathbf{g} | \boldsymbol{\delta}) + \ln p(\boldsymbol{\delta}) \}. \end{aligned} \quad (11)$$

We assume a Gibbs prior on the indicator image δ , an approach that is now widely used in image processing and pattern analysis [21]. The Gibbs prior is defined as

$$p(\delta) = \frac{1}{Z} \exp[-\beta U(\delta)] \quad (12)$$

where

$$U(\delta) = \sum_{j,k \in C_j} V(\delta_j, \delta_k). \quad (13)$$

The function U is called the energy function and is the weighted sum of potential functions V associated with individual *cliques* C_j . The parameter β determines the relative weight of the prior term in the posterior function, and the parameter Z is a normalization factor which is unnecessary to calculate.

For the Poisson linear model, the MAP solution can be obtained by the one-step-late algorithm [22],[20]:

$$\hat{\delta}_j^{(l+1)} = \frac{\hat{\delta}_j^{(l)}}{\left\{ \sum_{i=0}^{N-1} F_{ij} \right\} + \beta \frac{\partial U(\delta)}{\partial \delta_j^{(l)}} \sum_{i=0}^{N-1} \left\{ \frac{F_{ij} g_i}{\sum_{k=0}^{N-1} F_{ik} \delta_k^{(l)}} \right\}}, \quad j = 0, \dots, N-1. \quad (14)$$

Note that the MAP algorithm differs from the ML EM algorithm only in the extra term in the denominator. As β approaches zero, the MAP algorithm places diminishing emphasis on the prior information, and approaches the ML solution.

In traditional applications of MAP estimation to image reconstruction and restoration, the purpose of the prior is to encourage smoothness of the solution, under the assumption that the true image is correlated whereas the noise is not. In our application, the true signal (the delta function) is not low-

pass; therefore, standard image priors are not appropriate. Therefore, we propose a modification of the traditional approach to make it suitable for our problem.

Specifically, we propose the following potential function, which is an adaptation of a form that has been used widely in other image-recovery problems [22]:

$$V(r) = -\lambda \ln \left[\cosh\left(\frac{r}{\lambda}\right) \right] \quad (15)$$

which yields

$$\frac{\partial V(r)}{\partial r} = -\frac{e^{r/\lambda} - e^{-r/\lambda}}{e^{r/\lambda} + e^{-r/\lambda}}. \quad (16)$$

where

$$r = \frac{\hat{\delta}_j}{\hat{\delta}_i}, \quad (17)$$

with j and k denoting the pixel indices of neighboring pixels. Examples of the derivative of the potential function, $dV(r)/dr$ are shown in Fig. 1 for various values of the parameter λ .

The practical effect of this prior on the iterative formula can be understood by considering a location j where the value δ_j is much larger than at the surrounding locations, so that $r \gg 1$. This leads to a negative value for $dV(r)/dr$, which decreases the denominator in Eq. (14), thus increasing the estimate of δ_j . If the value δ_j is not significantly larger than its surrounding locations (so that r is nearly one), the derivative term will be nearly zero and the prior will have little effect on the estimate. The parameter λ controls the sensitivity of the potential function to differences among neighboring pixels.

Note that the aim of this prior is contrary to the usual purpose of priors in image recovery. Whereas priors usually are used to encourage a smooth image solution (based on the prior assumption that natural images are smooth), our potential function seeks to emphasize the impulsive nature of the output (based on our knowledge that δ is a collection of delta functions). Those familiar with Gibbs priors will recognize that, whereas r is usually defined as the difference between neighboring pixel values, we have instead defined it as a *ratio* of these values. We have also negated the function $V_1(r)$ to make it have the desired sense of optimality.

3. Generalized Likelihood Ratio Test (GLRT)

In this section, we describe a second approach to the problem detecting a signal in Poisson noise, namely a GLRT [1]. The GLRT is a standard hypothesis test for choosing between competing hypotheses based on noisy observations. In general, the GLRT solves the binary detection problem when the likelihood function $p(\mathbf{g}|H_j, \boldsymbol{\theta}_j)$ describing hypothesis H_j , $j=1,2$, is specified by an unknown parameter vector $\boldsymbol{\theta}_j$. In a GLRT, the log-likelihood ratio is evaluated by replacing the unknown parameter vectors with their maximum-likelihood (ML) estimates, i.e.,

$$\Lambda(\mathbf{g}) = \frac{p(\mathbf{g}|H_1; \hat{\boldsymbol{\theta}}_1)}{p(\mathbf{g}|H_0; \hat{\boldsymbol{\theta}}_0)}, \quad (18)$$

where $\hat{\boldsymbol{\theta}}_j$ is ML estimate of $\boldsymbol{\theta}_j$. The hypothesis test is performed by comparing the log-likelihood ratio

$\ln \Lambda(\mathbf{g})$ to a pre-determined threshold T , and making a decision accordingly, i.e.,

$$\ln \Lambda(\mathbf{g}) \underset{d_0}{\overset{d_1}{\gtrless}} T, \quad (19)$$

where d_j indicates decision in favor of hypothesis H_j . Methods for selecting the threshold T are well known [1]. In this paper, we evaluate the results using localization receiver-operating-characteristic (LROC) curves [23], which examine the entire range of possible choices for T .

We frame the object detection problem as a binary hypothesis test at every image location, in which one chooses between hypothesis H_0 (that the object is not present at a given location) and H_1 (that the object is present at a given location). Let us define \mathbf{g}_i as a vector containing the photon counts in a small $K_1 \times K_2 (= N_w)$ image window W_i centered at the current test pixel i . Using this definition, we model the competing hypotheses as:

$$\begin{aligned} H_0 : \mathbf{g}_i &\sim \text{Poisson}(b_i \mathbf{1}) \\ H_1 : \mathbf{g}_i &\sim \text{Poisson}(a_i \mathbf{f}_i) \end{aligned} \quad (20)$$

where $\mathbf{1}$ is a vector representation of a $K_1 \times K_2$ image window filled with ones, b_i is the unknown scalar value of the background in the vicinity of test location i , and a_i is an unknown intensity scaling factor for the template object \mathbf{f}_i .

In Eq. (20), we are posing the task as one of deciding whether the image mean within the test window W_i is the known signal \mathbf{f}_i (within an intensity scaling factor a_i) or a locally uniform background with an unknown amplitude b_i . Clearly, we do not expect the image to be truly uniform, even locally, when the object is not present. However, given our lack of information about the null hypothesis, the local-uniformity assumption is a suitably neutral statement about the background intensity. While local uniformity would seem to be a very simplistic assumption, our results show that it is surprisingly effective.

Now we derive the specifics of the GLRT algorithm. The likelihood functions for the hypotheses in Eq. (20) are:

$$p(\mathbf{g}_i | H_0; b_i) = \prod_{j \in W_i} \frac{b_i^{g_{i,j}} e^{-b_i}}{g_{i,j}!} \quad (21)$$

and

$$p(\mathbf{g}_i | H_1; a_i) = \prod_{j \in W_i} \frac{(a_i f_{i,j})^{g_{i,j}} e^{-a_i f_{i,j}}}{g_{i,j}!}, \quad (22)$$

where $f_{i,j}$ and $g_{i,j}$ denote the value of pixel j in test window W_i for the template \mathbf{f}_i and observations \mathbf{g}_i , respectively. Substituting (21) and (22) into (18), it is easy to show that

$$\Lambda(\mathbf{g}) = \prod_{j \in W_i} \left(\frac{\hat{a}_i f_{i,j}}{\hat{b}_i} \right)^{g_{i,j}} e^{-\hat{a}_i f_{i,j} + \hat{b}_i}. \quad (23)$$

where \hat{a}_i and \hat{b}_i are ML estimates of a_i and b_i , respectively. Taking the natural logarithm of (23) yields:

$$\ln \Lambda(\mathbf{g}) = \sum_{j \in W_i} \left[g_{i,j} \ln \left(\frac{\hat{a}_i f_{i,j}}{\hat{b}_i} \right) - \hat{a}_i f_{i,j} + \hat{b}_i \right]. \quad (24)$$

Now the ML estimates \hat{a}_i and \hat{b}_i are obtained simply by solving the following equations:

$$\frac{\partial}{\partial a_i} p(\mathbf{g}_i | H_1; a_i) = 0 \quad (25)$$

$$\frac{\partial}{\partial b_i} p(\mathbf{g}_i | H_0; b_i) = 0. \quad (26)$$

It is easy to show that these equations yield the following solutions:

$$\hat{a}_i = \frac{\sum_{j \in W_i} g_{i,j}}{\sum_{j \in W_i} f_{i,j}}. \quad (27)$$

$$\hat{b}_i = \frac{1}{N_W} \sum_{j \in W_i} g_{i,j}. \quad (28)$$

Substituting Eqs. (24), (27), and (28) into Eq. (19), we obtain the following decision rule for the GLRT:

$$\sum_{j \in W_i} g_{i,j} \ln(N_W \tilde{f}_{i,j}) \underset{d_0}{\overset{d_1}{\geq}} T, \quad (29)$$

where $\tilde{f}_{i,j} = f_{i,j} / \sum_{j \in W_i} f_{i,j}$ is a normalized version of the template.

Note that the GLRT decision rule in Eq. (29) can be computed simply by a cross correlation of the observed image $g_{i,j}$ with the kernel $\ln(N_W \tilde{f}_{i,j})$. This can be computed rapidly as a convolution using fast Fourier transforms, so it is a much faster technique than the IR methods described earlier.

In the following section, we compare the performance of our GLRT and IR methods with a more-traditional method, which we call the “exact LRT”, which has been proposed before [4]. This “exact LRT” makes the simplistic assumption that the template intensity is known exactly (there is no scaling factor a_i as in our GLRT) and that the background is zero. These assumptions lead to the following decision rule:

$$\sum_{j \in W} g_{i,j} \ln(f_{i,j}) \underset{d_0}{\overset{d_1}{\geq}} T. \quad (30)$$

Note that this method cannot be implemented in most practical applications, because it requires exact knowledge of the template. In our experiments, we provided this “exact” method with the true template

intensity to determine its best-case performance. In spite of this, it substantially underperformed the GLRT, which had no such information available to it. Note also that the “exact LRT” in Eq. (30) is almost identical in form to the GLRT in Eq. (29). The GLRT owes its performance advantage to the correct template normalization. Specifically, the argument of the logarithm must sum to N_w .

4. Experimental results

To evaluate our results we use a variation of the receiver operating characteristic (ROC) curve, a well known and comprehensive way to describe the detection performance of a human or machine observer [1]. The ROC curve is a plot of the probability of correct detection versus the probability of false alarm for the continuum of possible decision thresholds. Thus, it summarizes the range of trade-offs between missed detections and false alarms. The limitation of the ROC curve in characterizing a template-matching algorithm is that it only captures the ability of the algorithm to decide the presence of the object, but does not test whether the perceived location of the object is correct.

The localization ROC (LROC) curve [23] remedies this shortcoming by taking into account localization performance as well as detection performance. An LROC curve is a plot of the probability of detection *and* correct localization P_{DL} versus the probability of false alarm P_{FA} . Thus, a decision is said to be correct only if the object is both detected and located correctly by the object-recognition algorithm. Whereas the ROC curve has the property that $P_D \rightarrow 1$ as $P_{FA} \rightarrow 1$, the LROC has $P_{DL} \rightarrow K$ where $K \leq 1$ as $P_{FA} \rightarrow 1$.

In this paper LROC curves were computed using Monte Carlo simulations. In each experimental trial, an image of a target object (either a tank, truck, or car (Figure 2)) was artificially embedded at a

random location within one of five different 256×256 background scenes. An example of a noise-free scene is shown in Figure 3. The other background images that were used are shown in Figure 4.

Poisson noise was added to each test scene to simulate photon-limited imaging. In this type of imaging, the signal-to-noise ratio (SNR) increases in proportion to the mean image intensity (because the variance of a Poisson random variable is equal to its mean) [6],[24]. Therefore, images with lower average numbers of photons have poorer SNR. In Fig. 5 we show noisy versions of the test scene in Fig. 3 at total mean photon counts of 25,000, 50,000 and 100,000.

For each image, all the algorithms were applied three times, once for each of the candidate signals (tank, truck, and car). The peak of the output was taken to be the algorithm's decision as to the location of the object, and the output value of the algorithm (the likelihood ratio in the GLRT or the estimate delta signal in the IR methods) was used as a decision variable t . If t exceeded a decision threshold T , then the object was said to be present at the location of the peak; otherwise, the object was said to be absent.

In order to obtain the LROC, P_{DL} and P_{FA} must be computed as functions of the decision threshold T . These probabilities were obtained by numerical evaluation of the following integrals:

$$\begin{aligned} P_{DL} &= \int_T^{\infty} p(t|H_1)dt \\ P_{FA} &= \int_T^{\infty} p(t|H_0)dt \end{aligned} \tag{31}$$

where $p(t|H_j)$ is the conditional probability density function (PDF) of t given hypothesis H_j . The conditional PDFs $p(t|H_1)$ and $p(t|H_0)$ were obtained by generating sets of 50 images, each with a different noise realization and with different random object location. In one set of images, the signal was

present in each scene; in the other set, the signal was absent. In the case where the object was present, each algorithm was applied to every image, and the magnitude t of the peak value of the decision function was recorded, provided that the peak correctly indicated the template location. The normalized histogram of these values was then used as $p(t|H_1)$. To find $p(t|H_0)$ the set of images in which the object is absent was used. Again, the algorithms were applied to each image, but the magnitude of the peak value of the decision function was used regardless of its location. The normalized histogram of the recorded peak values of the decision function obtained was used as $p(t|H_0)$. The LROC curves were obtained by varying the decision threshold T and computing the integrals in Eq. (31).

Examples of the outputs of the algorithms tested are shown in Fig. 6. The overall LROC curves are shown in Fig. 7. The GLRT performed best overall in this experiment, in all cases producing by far the highest detection/correct-localization probability at any given false-alarm probability. As expected, the performance of the GLRT decreased with decreasing photon counts (decreasing SNR), but maintained excellent performance even at only 25,000 counts. Even though the “exact LRT” was provided with exact information as to the intensity of the template, it was unable to come close to matching the performance of the more-realistic GLRT method (which did not have access to this information).

The IR methods also require exact knowledge of the template intensity. To make these methods realistic, each method estimated the template intensity by assuming it to be equal to the average intensity of the scene, as measured from the observed data. The IR algorithms could be modified to incorporate local intensity estimation at every pixel (like the GLRT does), but unlike the GLRT, the result would be an iterative space-variant calculation, which would be far too computationally intensive for most practical applications.

From the example output images in Fig. 6, one gets the impression that the MAP method produces the sharpest output; however, this impression not meaningful, because the LROC curves show that the MAP method does not perform as well as the GLRT.

5. Conclusions and future work

We have studied three algorithms for object detection in Poisson noise: an EM IR method, a MAP IR method, and a GLRT. The GLRT is based on a seemingly simplistic model that the image is locally uniform when the signal is absent. In spite of this, the GLRT performed extremely well in our experiments, even at very low photon counts.

In our prior work in Gaussian noise applications (e.g., [14]), we have found the impulse restoration (IR) approach to work well, but it did not perform well in the Poisson-noise setting studied in this paper.

In this experiment, we did not consider geometrical distortions of the object (e.g., rotation and scale); however, these can be incorporated in the GLRT framework by considering them as additional parameters to be estimated within the parameter vector θ_j (in a manner similar to our inclusion of intensity parameters a_i and b_i). We will study this extension of the GLRT method in future work.

Acknowledgments

This work was supported in part by National Institutes of Health NIH/NCI Grant CA89668. We thank the anonymous reviewers for several helpful suggestions that contributed to improving the paper. The corresponding author is Miles N. Wernick (wernick@iit.edu).

References

- [1] S.M. Kay, *Fundamentals of Statistical Signal Processing, Detection Theory*. Upper Saddle River, New Jersey: Prentice Hall, 1998.
- [2] A.K. Jain, *Fundamentals of Digital Image Processing*. Englewood Cliffs, New Jersey: Prentice Hall 1989.
- [3] G.M. Morris, "Scene matching using photon-limited images," *J. Opt. Soc. Am. A*, 1(5), 482-488, (1984).
- [4] M.N. Wernick, and G.M. Morris, "Image classification at low light levels," *J. Opt. Soc. Am. A*, 3(12), 2179-2187, (1986).
- [5] R. E. Sequira, , J.Gubner, , and B.E.A. Saleh, "Image detection under low-level illumination," *IEEE Trans. Image Processing*, 2(1), 18-26, (1993).
- [6] M.N. Wernick, and J.N. Aarsvold, eds., *Emission Tomography : The Fundamentals of PET and SPECT*, Academic Press 2004.
- [7] J. Ben-Arie, and K. Rao, "A novel approach to template matching by non-orthogonal image expansion," *IEEE Trans. Circuits and Systems for Video Technology*, 3(1), 71-84, (1993).
- [8] J. Ben-Arie and K. Rao, "Optimal template matching by non-orthogonal image expansion using restoration," *International Journal of Machine Vision and Applications*, 7(2), 69-81, (1994).
- [9] C. Chatwin, R. Wang, and R. Young, "Assessment of a Wiener filter synthetic discriminant function for optical correlation," *Journal of Optics and Lasers in Engineering*, 22(1),33-51, (1995).
- [10] E.Marom and H. Inbar , "New interpretations of wiener filters for image recognition," *J. Opt. Soc. Am. A*, 13(7), 1325-1330, (1996).
- [11] B. Javidi, , F. Parchekani, and G. Zhang, , "Minimum-mean-square error filters for detecting a noisy target in background noise," *Appl. Opt.*, 35, 6964-6975, (1996).

- [12] R. Dufour, E. Miller and N.P. Galatsanos, "Geometric parameter estimation with a multiscale template library," in *Automatic Target Recognition X*, F.A. Sadjadi, ed., *Proc. SPIE*, **4050**, 2000.
- [13] R. Dufour, E. Miller, and N.P. Galatsanos, "Template matching based object recognition with unknown geometry," *IEEE Trans. Image Processing*, **11**(12), 1385-1396, (2002).
- [14] A. Abu-Naser, "Object Recognition Based on Impulse Restoration for Images in Gaussian Noise," M.S. Thesis, Illinois Institute of Technology, Chicago, 1996.
- [15] A. Abu-Naser, N.P. Galatsanos, M.N. Wernick, D. Schonfeld, "Object recognition based on impulse restoration with use of the expectation maximization algorithm," *J. Opt. Soc. Am. A*, **15**(9), 2327-2339, (1998).
- [16] A. Abu-naser, N. P. Galatsanos and M. N. Wernick, "ML and Bayesian impulse restoration based object recognition in photon limited noise," *Proceedings of the 2002 IEEE International Conference on Image Processing*, **2**, 837-840, 2002
- [17] A. Dempster, N. Laird, and D. Rubin, "Maximum-likelihood from incomplete data," *J. Roy. Stat. Soc.-B*, **39**,1-38, (1977).
- [18] L. Shepp and Y. Vardi, "Maximum likelihood reconstruction for emission tomography," *IEEE Trans. Med. Imaging*, **1**(2), 113-122, (1982).
- [19] K. Lange and R. Carson, "EM reconstruction algorithms for emission and transmission tomography," *J. Comput. Assist. Tomog.*, **8**(2), 306-316, (1984).
- [20] D.S. Lalush and M.N.Wernick, "Iterative Image Reconstruction," in *Emission Tomography: The Fundamentals of PET and SPECT*, M.N. Wernick and J.N. Aarsvold, eds., Academic Press: New York, 2004.
- [21] S. Geman and D. Geman, "Stochastic distributions, and the Bayesian restoration of images," *IEEE Trans. Patt. Anal. Machine Intell.*, **6**, 721-741, (1984).
- [22] P.J. Green, "Bayesian reconstructions from emission tomography data using a modified EM algorithm," *IEEE Trans. Med. Imaging*, **9**(1), 84-93, 1990.
- [23] R. Swensson, "Unified measurement of observer performance and localizing target objects on images," *Med. Phys.*, **23**, 1709-1725, (1996).
- [24] B. E. A. Saleh and M. C. Teich, *Fundamentals of Photonics*, Wiley, New York, 1991.

FIGURE CAPTIONS

Figure 1. Plot of the derivative of the clique energy V as a function of r for several values of λ .

Figure 2. Target objects that were embedded within the simulated test scenes.

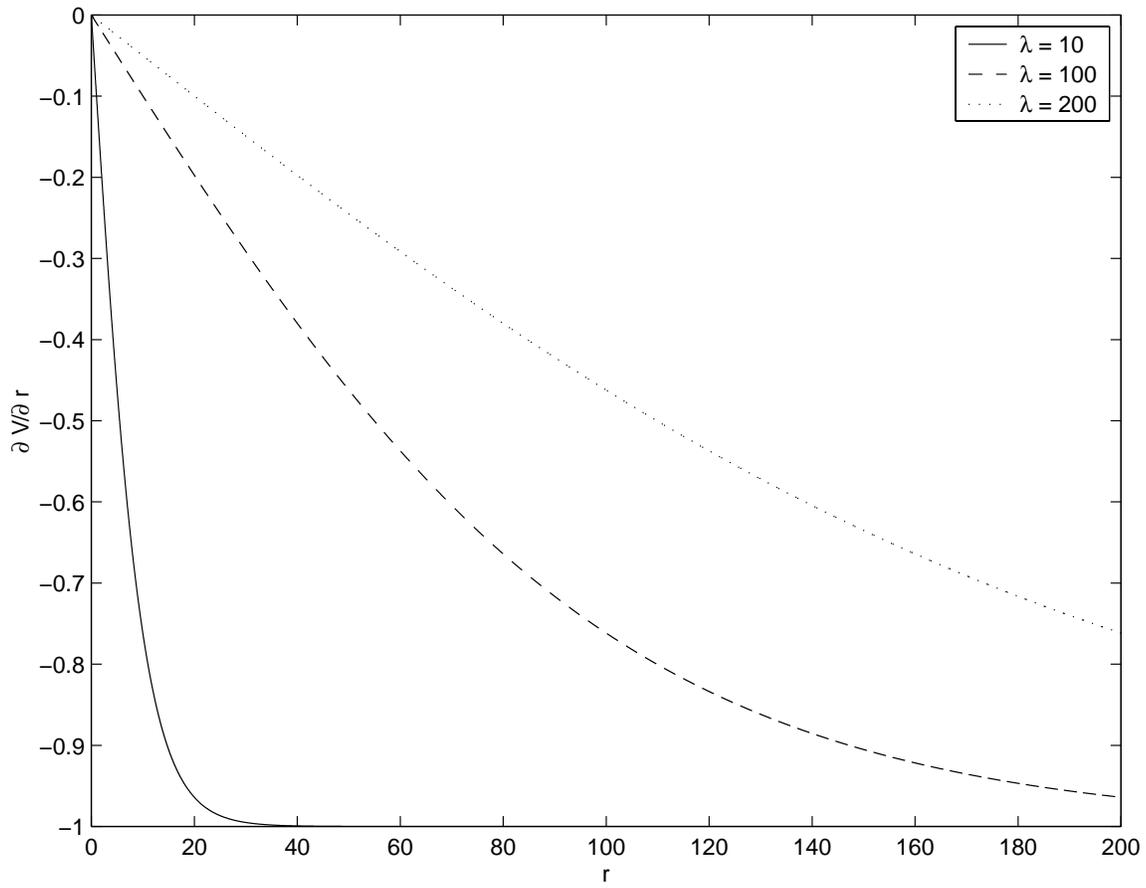
Figure 3. Example noise-free scene from which noisy test images were generated.

Figure 4. Additional background images used in the simulations. Test scenes were generated in equal numbers using these four backgrounds and the one shown in Fig. 3.

Figure 5. Examples of simulated photon-limited test scenes for mean total photon counts of 25,000 (left), 50,000 (center) and 100,000 (right). The object is virtually impossible to detect visually; however, the GLRT finds it easily.

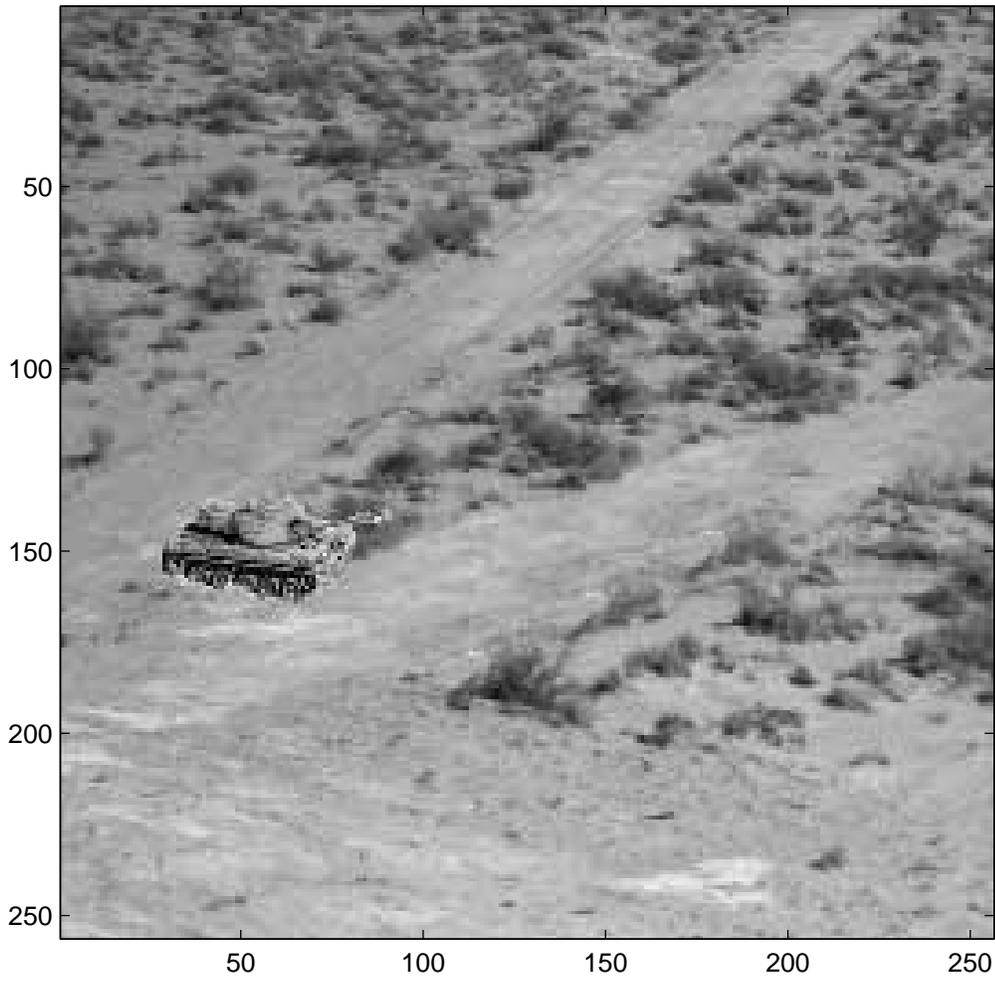
Figure 6. Examples of the output of tested algorithms for the 50,000 photon-count case: EM (upper left), MAP (upper right), exact LRT (lower left), GLRT (lower right). Arrows point to the target location.

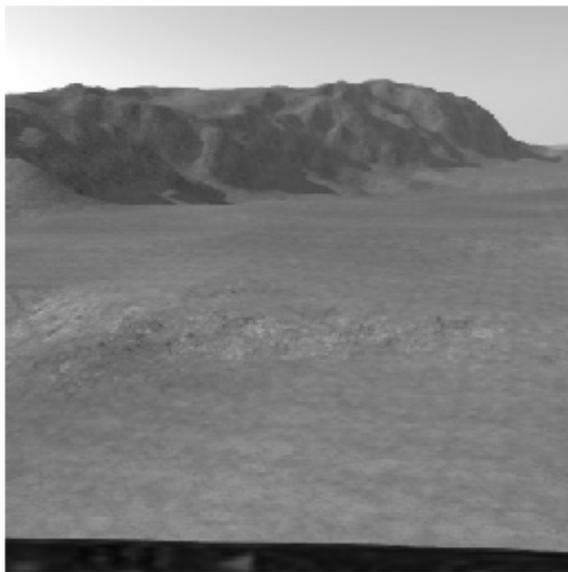
Figure 7. LROC curves for mean total photon counts of 25,000 (left), 50,000 (center) and 100,000 (right). The proposed GLRT performed best among the methods tested, and its performance improves consistently with increasing number of photons.



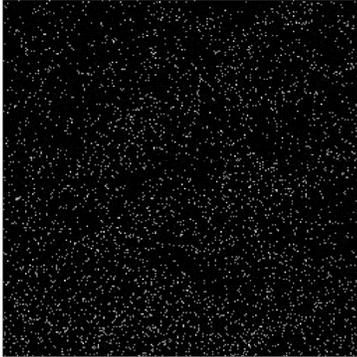


TANK AT LOCATION (133,28)

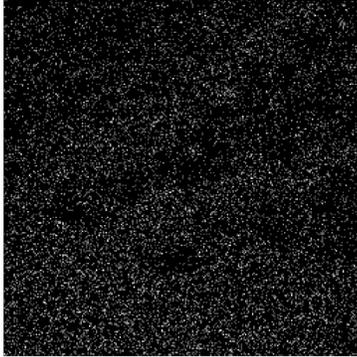




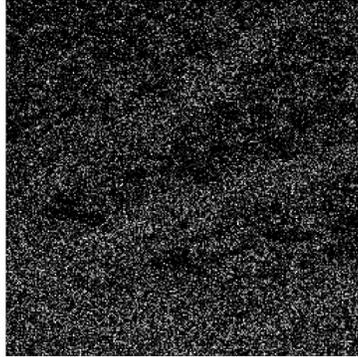
25k photons

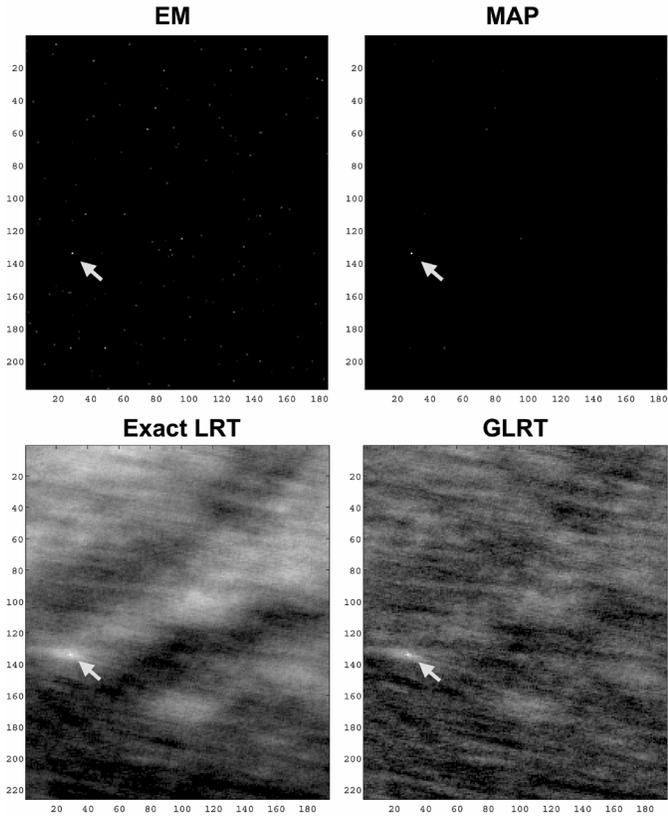


50k photons

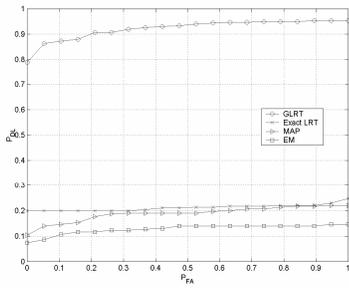


100k photons

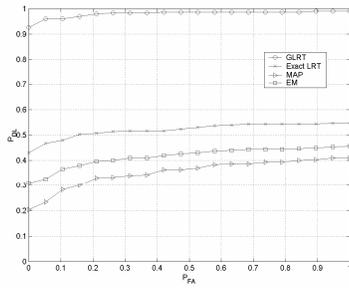




25k photons



50k photons



100k photons

