

Assessing the accuracy of estimates of the likelihood ratio

Eric Clarkson^{a,b}, Matthew A. Kupinski^{a,b} and John W. Hoppin^c

^aOptical Sciences Center

^bDepartment of Radiology

^cProgram in Applied Mathematics

The University of Arizona, Tucson, AZ

ABSTRACT

There are many methods to estimate, from ensembles of signal-present and signal-absent images, the area under the receiver operating characteristic curve for an observer in a detection task. For the ideal observer on realistic detection tasks, all of these methods are time consuming due to the difficulty in calculating the ideal-observer test statistic. There are relations, in the form of equations and inequalities, that can be used to check these estimates by comparing them to other quantities that can also be estimated from the ensembles. This is especially useful for evaluating these estimates for any possible bias due to small sample sizes or errors in the calculation of the likelihood ratio. This idea is demonstrated with a simulation of an idealized single photon emission detector array viewing a possible signal in a two-dimensional lumpy activity distribution.

1. INTRODUCTION

The receiver operating characteristic (ROC) curve is often used to quantify the performance of an observer on a signal-detection task. The area under the ROC curve (AUC) is a number that can be estimated from ROC data and used as a figure of merit for observer performance on the task in question. The ideal observer is a mathematical observer whose decisions can be arrived at by computations, if we have full knowledge of the statistics of the data under both the signal-absent and signal-present hypotheses. Given a task and the statistical model for the data, the ideal observer computes a statistic called the likelihood ratio and compares it to a threshold in order to decide whether the signal is present or absent. The AUC of the ideal observer is greater than or equal to the AUC for any other observer for a given detection task and can therefore be used as a measure of the quality of the data for the task. This in turn implies that the ideal-observer AUC can be used as a figure of merit for the imaging system that generates the data. Unfortunately, the AUC for the ideal observer can be difficult to compute for realistic imaging system models and detection tasks.

We have recently developed a method for computing the likelihood ratio for simulations that accounts for randomness in the objects being imaged as well as measurement noise in the imaging system. We do not use any Gaussian assumptions for either source of randomness, and this fact results in the need to compute the value of a large-dimensional integral in order to estimate the likelihood ratio. The nominal dimension for this integral is equal to the number of detectors in the imaging system. By making use of parameterized object models, such as lumpy backgrounds and clustered lumpy backgrounds, to describe object randomness, we can effectively reduce this dimension to the number of free parameters in the model, that is, the number of parameters it takes to specify an individual object. The object model itself has adjustable parameters, which are chosen to optimize the agreement between the statistical properties of the model with the statistics of real objects[1,2]. We then employ Markov-chain Monte Carlo techniques in this relatively low-dimensional parameter space to estimate the value of the integral [3].

This estimate of the likelihood ratio is necessarily a random variable itself. One source of randomness in this number is the data vector. However, even when we fix the data vector, the random choices that are made when we run the Markov chain are another source of variation in the end result. Our estimate of the AUC for the ideal observer is based estimating the likelihood ratio for each member of a large sample of data vectors.

Corresponding author: E.C., E-mail: clarkson@radiology.arizona.edu, Address: Department of Radiology, The University of Arizona, Tucson, AZ 85724

This sample introduces an additional source of randomness in our estimate of the AUC. The bias and variance of this AUC-estimate will converge to zero as the sample size, and the length of the Markov chains for each data vector in the sample, are all increased. We are, of course, restricted to finite samples and finite-length Markov chains, and we therefore would like to have information about the bias and variance of our estimate of the AUC. The variance of can be estimated using standard resampling techniques, repeating the calculation with different image ensembles to account for the variation from the random data vectors, and repeating the Markov-chain calculations for fixed data vectors, to account for the random variation due to the Markov-chain procedure.

It is more difficult to get information about the bias of our estimate. To this end we have developed tests which are based on equalities and inequalities for the AUC of the likelihood ratio that have been developed by ourselves and others. These relations have been derived from properties of the likelihood ratio and will not, in general, hold for other test statistics. Furthermore, we can estimate the quantities on both sides of each relation from the samples of data vectors that we use to estimate the AUC. If the estimates fail to satisfy the relations, at least to within the degree of approximation determined by variance estimates, then we know that we must increase either the sample size or the length of the Markov chains. If the estimates satisfy these relations, then this is no guarantee of their reliability, but in combination with the variance estimates it does give us some confidence in our estimates of the AUC of the ideal observer.

In what follows we will first briefly summarize our method for computing the likelihood ratio, which has been presented in more detail elsewhere [3]. Then we will present the equations and inequalities that relate the AUC for the likelihood ratio to other quantities that can also be estimated from our simulations. We will also give some details on how to estimate these other quantities. Finally, we show some results from a simulation study of an idealized parallel-hole-collimator SPECT system viewing a two-dimensional lumpy object which may or may not have a known signal superimposed upon it.

2. COMPUTING THE LIKELIHOOD RATIO

The ideal observer takes the data vector \mathbf{g} and computes the likelihood ratio

$$\Lambda(\mathbf{g}) = \frac{pr(\mathbf{g}|H_1)}{pr(\mathbf{g}|H_0)}. \quad (1)$$

The numerator is the probability density for the data under the signal present hypothesis, while the denominator is the probability density under the signal absent hypothesis. This test statistic is compared to a threshold Λ_0 and the signal is declared to be present if $\Lambda(\mathbf{g}) > \Lambda_0$, otherwise it is declared to be absent.

For our purposes, i.e. SPECT imaging, the data vector is a Poisson random vector whose mean $\bar{\mathbf{g}}$ is given by

$$\begin{aligned} H_0 : \bar{\mathbf{g}} &= \mathcal{H}f_\theta \\ H_1 : \bar{\mathbf{g}} &= \mathcal{H}(f_\theta + s) \end{aligned} \quad (2)$$

under the signal-absent and signal-present hypotheses respectively. The K -dimensional vector θ is the vector of parameters that specifies the background function f_θ . This is a random vector with probability density $pr(\theta)$. The background model, which is the combination of the density $pr(\theta)$ and the procedure for determining f_θ from θ , is chosen to mimic certain statistical properties of real world backgrounds. The process for choosing this model is detailed elsewhere [1,2]. The operator \mathcal{H} maps functions to M -dimensional vectors in data space and is intended to model the imaging system under consideration. For simplicity, the signal s will be fixed throughout this discussion, but we may also account for randomness in this function by including additional integrals over signal parameters in the expressions below

By defining $\mathbf{b}(\theta) = \mathcal{H}f_\theta$ and $\mathbf{s} = \mathcal{H}s$, we may write the two hypotheses as,

$$\begin{aligned} H_0 : \bar{\mathbf{g}} &= \mathbf{b}(\theta) \\ H_1 : \bar{\mathbf{g}} &= \mathbf{b}(\theta) + \mathbf{s}. \end{aligned} \quad (3)$$

We now have conditional probability densities on the data for fixed θ under each hypothesis given by

$$\begin{aligned} pr(\mathbf{g}|\theta) &= \prod_{m=1}^M \frac{[b_m(\theta)]^{g_m}}{g_m!} \exp(-b_m(\theta)) \\ pr(\mathbf{g}|\theta, \mathbf{s}) &= \prod_{m=1}^M \frac{[b_m(\theta) + s_m]^{g_m}}{g_m!} \exp(-b_m(\theta) - s_m). \end{aligned} \quad (4)$$

The numerator and denominator of the likelihood ratio are given by the following K -dimensional integrals over a region R in parameter space

$$\begin{aligned} pr(\mathbf{g}|H_1) &= \int_R pr(\mathbf{g}|\theta, \mathbf{s}) pr(\theta) d^K\theta \\ pr(\mathbf{g}|H_0) &= \int_R pr(\mathbf{g}|\theta) pr(\theta) d^K\theta. \end{aligned} \quad (5)$$

One approach to estimating the likelihood ratio would be to try to estimate each of these integrals directly via Monte Carlo techniques, but we have had no success with this approach due to the sharply peaked nature of the conditional probabilities in the integrands.

Another approach, which we have had more success with, involves defining the background-known exactly (BKE) likelihood ratio via

$$\Lambda_{BKE}(\mathbf{g}|\theta) = \frac{pr(\mathbf{g}|\theta, \mathbf{s})}{pr(\mathbf{g}|\theta)} = \prod_{m=1}^M \left[\frac{b_m(\theta) + s_m}{b_m(\theta)} \right]^{g_m} \exp(-s_m). \quad (6)$$

We also define a posterior probability density on the background parameter vector θ using Bayes' rule

$$pr(\theta|\mathbf{g}, H_0) = \frac{pr(\mathbf{g}|\theta) pr(\theta)}{pr(\mathbf{g}|H_0)}. \quad (7)$$

In terms of these two functions the likelihood ratio may now be written as a single integral in parameter space

$$\Lambda(\mathbf{g}) = \int_R \Lambda_{BKE}(\mathbf{g}|\theta) pr(\theta|\mathbf{g}, H_0) d^K\theta. \quad (8)$$

One might suspect that no progress has actually been made here since the denominator in the expression for the posterior density is also the denominator of the likelihood ratio, which, as noted above, is difficult to estimate. However, we use a Markov-chain Monte Carlo method to sample from the posterior density on θ in order to estimate the integral in (8). In order to do this we need only to be able to calculate ratios of the form

$$\frac{pr(\theta'|\mathbf{g}, H_0)}{pr(\theta|\mathbf{g}, H_0)} = \frac{pr(\mathbf{g}|\theta') pr(\theta')}{pr(\mathbf{g}|\theta) pr(\theta)}. \quad (9)$$

All of the terms in this fraction are known, and we can therefore estimate a value for the likelihood ratio by running a single Markov chain and averaging the values of $\Lambda_{BKE}(\mathbf{g}|\theta)$ that are generated as the chain proceeds. Details of this procedure have been presented elsewhere [3]. In particular the choice of a good proposal density and starting point for the chain are crucial for getting meaningful results, but we will not go into those problems here.

There is a technical point which we will note here and then ignore for the rest of this discussion. For Poisson noise the data vector takes a discrete set of values and the usual notion of a continuous probability density function does not make sense. However, the two functions in (4) or (5) are defined on the same discrete set of data vectors, and therefore the ratio of each pair makes sense as a function on this same set. The mathematical terminology for this situation is that both densities are absolutely continuous with respect to the counting measure on the (countable) set of data vectors. As a result of this fact, the likelihood ratio is a discrete random variable. Strictly speaking then, integrals involving this variable, many of which will appear below, should be replaced with sums. Similarly, any integrals involving \mathbf{g} as the integration variable should also be replaced with sums. We have not done this since the integral notation is more convenient and applies to situations where the data vector is not discrete. With this caveat in place, we will now proceed with our discussion.

3. EXPRESSIONS FOR THE AUC OF THE IDEAL OBSERVER

Since Λ is itself a random variable, it has probability densities $pr(\Lambda|H_1)$ and $pr(\Lambda|H_0)$ under the two hypotheses. The true positive fraction (TPF) and false positive fraction (FPF) for a given threshold are given by the standard expressions

$$\begin{aligned} TPF(\Lambda_0) &= \int_{\Lambda_0}^{\infty} pr(\Lambda|H_1) d\Lambda \\ FPF(\Lambda_0) &= \int_{\Lambda_0}^{\infty} pr(\Lambda|H_0) d\Lambda. \end{aligned} \quad (10)$$

The ROC curve is generated by plotting $TPF(\Lambda_0)$ versus $FPF(\Lambda_0)$ as the threshold is varied. Finally, the AUC for the ideal observer is defined by

$$AUC_{\Lambda} = \int_0^1 TPF d(FPF). \quad (11)$$

It has been shown that [4], for a likelihood ratio only, the AUC can be also computed from the equation

$$2(1 - AUC_{\Lambda}) = \int_0^{\infty} [FPF(\Lambda)]^2 d\Lambda. \quad (12)$$

It must be emphasized that this equation does not hold for arbitrary test statistics since it follows from the relation

$$pr(\Lambda|H_1) = \Lambda pr(\Lambda|H_0), \quad (13)$$

which is true if and only if Λ is a likelihood ratio. Note that the expression on the right in (12) can be estimated from an ensemble of signal-absent images. On the other hand, the quantity AUC_{Λ} can be estimated from two ensembles of images, one with the signal present and one with the signal absent. A comparison of the two sides of this equation can be used as a check on the AUC estimate.

The moment-generating function for Λ under the hypothesis H_0 is defined by

$$M_0(\beta) = \int_0^{\infty} \Lambda^{\beta} pr(\Lambda|H_0) d\Lambda = \langle \Lambda^{\beta} \rangle_0. \quad (14)$$

It follows from this definition and (13) that $M_0(0) = M_0(1) = 1$. The converse of this statement is also true. That is, if we have a positive test statistic $\tilde{\Lambda}$ which satisfies

$$\langle \tilde{\Lambda} \rangle_0 = \int_0^{\infty} \tilde{\Lambda} pr(\tilde{\Lambda}|H_0) d\tilde{\Lambda} = 1 \quad (15)$$

then $\tilde{\Lambda}$ is a likelihood ratio. This follows from the fact that (15), together with the positivity of $\tilde{\Lambda}$ implies that

$$pr_1(\tilde{\Lambda}) = \tilde{\Lambda} pr(\tilde{\Lambda}|H_0) \quad (16)$$

is a probability density. We then have

$$\tilde{\Lambda} = \frac{pr_1(\tilde{\Lambda})}{pr(\tilde{\Lambda}|H_0)}. \quad (17)$$

From this discussion it follows that, if we estimate $M_0(1)$ from an ensemble of signal absent images and get a number close to 1, then there are two possible conclusions that we could draw from this result. The first is that we are succeeding in getting a good estimate of the likelihood ratio that we are interested in. The second is that we are somehow calculating a good estimate of some other likelihood ratio. Since we have a natural tendency to regard the latter possibility as unlikely, we are drawn toward the first conclusion.

We also have the constraint

$$\frac{d^2 M_0}{d\beta^2} = \langle (\log \Lambda)^2 \Lambda^\beta \rangle_0 > 0. \quad (18)$$

Therefore a plot of $M_0(\beta)$ versus β for $0 \leq \beta \leq 1$ should show a curve that is concave upward and passes through the points $(0, 1)$ and $(1, 1)$. Points on this curve can be estimated by using an ensemble of signal-absent images and the result examined as a test of our calculations.

In terms of the function $M_0(\beta)$ the AUC for the ideal observer can be calculated from the relation [5]

$$2(1 - AUC_\Lambda) = \frac{1}{\pi} \int_0^\infty \left| M_0 \left(\frac{1}{2} + i\alpha \right) \right|^2 \frac{d\alpha}{\alpha^2 + \frac{1}{4}}. \quad (19)$$

This equation can also be used to check our estimate for AUC_Λ , although the right-hand side is probably harder to estimate accurately than the corresponding number in (12). However, note that, if we write the moment generating function in terms of the log-likelihood $\lambda = \log \Lambda$, then

$$M_0 \left(\frac{1}{2} + i\alpha \right) = \int_{-\infty}^\infty \exp \left(\frac{1}{2} \lambda \right) pr(\lambda|H_0) \exp(i\alpha\lambda) d\lambda, \quad (20)$$

which is the Fourier transform of an integrable function, since $M_0(\frac{1}{2}) < 1$. Therefore we know that $|M_0(\frac{1}{2} + i\alpha)| \rightarrow 0$ for large α , and the high frequency components of the Fourier transform, which are the ones that are difficult to estimate, will not contribute significantly to the integral in (19). We will see a demonstration of this asymptotic behavior in the example below.

The fact that $M_0(\beta)$ satisfies $M_0(0) = M_0(1) = 1$ motivates the definition of the likelihood generating function $G(\beta)$ given by

$$M_0(\beta) = \exp \left[\beta(\beta - 1) G \left(\beta - \frac{1}{2} \right) \right]. \quad (21)$$

This function completely determines the statistics of Λ under both hypotheses [5]. The value of this function at the origin is given by

$$G(0) = -\frac{1}{4} \ln \langle \Lambda^{\frac{1}{2}} \rangle_0 = \frac{1}{4} d_B, \quad (22)$$

where d_B is the Bhattacharyya distance between the two densities $pr(\mathbf{g}|H_0)$ and $pr(\mathbf{g}|H_1)$:

$$d_B = -\ln \int [pr(\mathbf{g}|H_0) pr(\mathbf{g}|H_1)]^{\frac{1}{2}} d^M g \quad (23)$$

This number seems to have some significance for estimating AUC_Λ . It will appear in the bounds that we discuss below and can also be used in the approximation

$$2(1 - AUC_\Lambda) \approx 1 - \operatorname{erf} \left\{ \left[\frac{1}{2} G(0) \right]^{\frac{1}{2}} \right\} \quad (24)$$

which has proved to be accurate in many cases that we have examined where we can compute the exact value for AUC_Λ [4,6]. We call this the $SNR_{G(0)}$ approximation, in analogy with the more familiar approximation

$$2(1 - AUC_T) \approx 1 - \operatorname{erf} \left[\frac{1}{2} SNR_T \right] \quad (25)$$

which is often used to relate AUC_T to the ordinary signal-to-noise ratio for an arbitrary test statistic T . Following this analogy we use the symbol $SNR_{G(0)}$ for the quantity $[2G(0)]^{\frac{1}{2}}$.

4. BOUNDS ON THE AUC OF THE IDEAL OBSERVER

It has been shown elsewhere [7] that

$$\frac{1}{2}G(0) \leq -\ln[2(1 - AUC_\Lambda)] \leq \frac{1}{2}G(0) + \left[G(0) - \frac{1}{8}G''(0)\right]^{\frac{1}{2}} \quad (26)$$

The expressions on the left and right in this inequality can be estimated from an ensemble of signal-absent images. If our estimate for AUC_Λ does not satisfy these inequalities, then we take that as an indication that we need to increase our sample sizes or modify the Markov chain. To calculate the square root in the upper bound we can start by noting that

$$pr(\mathbf{g}|H_\beta) = \frac{\Lambda^\beta}{M_0(\beta)} pr(\mathbf{g}|H_0) = \frac{pr^{1-\beta}(\mathbf{g}|H_0) pr^\beta(\mathbf{g}|H_1)}{M_0(\beta)}$$

is a probability density on \mathbf{g} for $0 \leq \beta \leq 1$. This defines an exponential family of probability densities which contains $pr(\mathbf{g}|H_0)$ and $pr(\mathbf{g}|H_1)$. This family has been studied extensively in information theory [8]. We can define expectations under the H_β hypothesis via

$$\langle t(\Lambda) \rangle_\beta = \frac{1}{M_0(\beta)} \langle t(\Lambda) \Lambda^\beta \rangle_0 \quad (27)$$

Then, setting $\lambda = \log \Lambda$ as before, we have

$$2 \left[G(0) - \frac{1}{8}G''(0) \right] = \langle \lambda^2 \rangle_{\frac{1}{2}} - \langle \lambda \rangle_{\frac{1}{2}}^2 = var_{\frac{1}{2}}(\lambda). \quad (28)$$

This equation allows us to estimate the upper bound in (26) from an ensemble of signal-absent images. It is interesting to note that, if $\hat{\beta}_{ML}(\mathbf{g})$ is the maximum likelihood estimate of the parameter β from the data vector \mathbf{g} , then this function could also be used as a test statistic in the detection problem. As such it corresponds to an observer that is equivalent to the ideal observer. We do not know whether this curious fact has any practical implications.

We may get another inequality for AUC_Λ by following Shapiro [9]. We use the minimum total error rate to define the quantity P_e via

$$\min_{\Lambda_0} \{FPF(\Lambda_0) + FNF(\Lambda_0)\} = 1 + FPF(1) - TPF(1) = 2P_e. \quad (29)$$

We then have the inequalities

$$(2P_e)^2 \leq 2(1 - AUC_\Lambda) \leq 2P_e. \quad (30)$$

To estimate P_e we need both signal-present and a signal-absent ensembles of images, in contrast to the bounds in (26). The inequality in (30) can be a more stringent test, however, since the lower bound in this inequality has been shown to be higher than that in (26) in at least one example [8].

A special case of some bounds derived by Burnashev [10] for the minimum total error rate may be written as

$$\frac{1}{4}G(0) \leq -\ln(2P_e) \leq \frac{1}{4}G(0) + \frac{1}{2} \left[\langle \lambda^2 \rangle_{\frac{1}{2}} \right]^{\frac{1}{2}} \quad (31)$$

These inequalities can also be used to check likelihood ratio calculations although they do not involve AUC_Λ directly. One can derive further inequalities for AUC_Λ by combining (30) with (31), but these are weaker than those given in (26) and are therefore not needed.




Figure 1: The ROC curve for the simulated system.




Figure 2: A plot of the integrand of equation 19

5. AN EXAMPLE

As an example we consider an idealized SPECT imaging system consisting of a collimated detector array viewing a two-dimensional distribution of activity. The random backgrounds are generated by the lumpy background model and Poisson noise is added at the detector outputs. For this idealized system we assume that the sensitivity functions for the detectors are Gaussians, that the lumps in the background are Gaussians, and that the signal to be detected is a low-amplitude Gaussian of known width and location. The use of Gaussians throughout makes the calculations more analytically tractable, but it does not imply that the detector outputs are normally distributed. The Markov chains for the likelihood calculations had length 300,000 with the first 2000 iterations discarded as a burn-in period. Expectations and estimates of AUC_{Λ} that required only a signal absent ensemble were estimated from a set of 631 images, while those that required both ensembles were estimated from a set of 480 image pairs. The large sample size used here is dictated by the highly non-Gaussian statistics of the likelihood ratio. In particular the long tail on the distribution of Λ makes the estimation of moments of this distribution difficult.

The estimate of AUC derived from the empirical ROC curve is 0.86. The empirical ROC curve is shown in Figure 1. From equation (12) we get an estimate for AUC_{Λ} of 0.87. The integrand in (19) is plotted in Figure 2, where we can see that truncating the integral at $\alpha = 5$ is quite reasonable. The estimate for AUC_{Λ} from this equation is also 0.87. For curiosity we can check the approximation in (24) to see if it holds up in this case. The quantity on the right in this expressions is estimated as 0.87.




Figure 3. A plot of $M_0(\beta)$. The curve has the appropriate shape and comes very close to the point (1, 1) at the right endpoint.




Figure 4. The AUC estimates and bounds as a function of the number of images. The dotted lines are the AUC bounds from reference 7. The solid lines with the plus signs are the Shapiro bounds. The solid line with the stars is the FPF-estimate of AUC. The solid line is the G_0 -estimate of AUC. The dashed-line is the moment-estimate of AUC. The dashed-dotted line is the estimate of AUC using the empirical ROC curve.

Next we show a plot of $M_0(\beta)$ for $0 \leq \beta \leq 1$ for this system in Figure 3. The curve has the appropriate shape and comes very close to the point (1, 1) at the right endpoint, $M_0(1) = 1.04$. As additional checks we examine the inequalities. We estimate all of the quantities in each inequality. The first one, in (26) gives us $.64 \leq 1.61 \leq 1.73$; the second one, in (30), reduces to $.18 \leq .28 \leq .42$; and finally the third one, in (31) comes out as $.32 \leq .86 \leq 1.1$. Our conclusion from all of these checks is that our calculations are consistent with known properties of the likelihood ratio and its ROC curve. While we can never prove that we are getting a good estimate of the AUC of the ideal observer for this system and task, we take our results as evidence that our calculations are usable.

To examine the dependence on sample size we show plots of the estimates of AUC_Λ , and the upper and lower bounds on this number, versus sample size in Figure 4. The four estimates of AUC_Λ are from the empirical ROC curve, which ends at a sample size of 480 pairs, and equations (12), (19) and (24) which continue out to a sample size of 631 signal-absent images. The pairs of bounds are from (30), which ends a 480 pairs, and (26), which continues to 631 single images. The purpose of this plot is to get some sense of how large a sample that we need to get stable estimates of the various quantities.

6. FUTURE WORK

There are many ways to extend the simulation described in the previous section and we will continue to check our calculations for these more complicated systems and tasks as we did for the simpler one given above. We are currently examining systems that view three-dimensional objects and use pinhole apertures instead of collimators. We are also adding signal uncertainty to the task in order to see how this changes the performance of the ideal observer. The lumpy backgrounds are not entirely satisfactory for modeling complicated background textures, so we plan to extend these calculations to the clustered lumpy background model as a step toward more realism. Eventually we will also have to take structures in the background into account. We believe that the basic methods for estimating the AUC of the ideal observer and checking these calculations that we have outlined above can be used in all of these more complicated situations with sufficient computer power. We also need to better quantify the degree of expected agreement among the various estimates of AUC_{Λ} and the upper and lower bounds for this number. For this purpose we will need to examine the statistics of all of our estimates in more detail, beyond the simple variance estimates that we do now. In this way we will be able to more fully justify our belief in these methods.

ACKNOWLEDGMENTS

This study was made possible by grants R37 EB000803, P41 RR14304 and K01 CA87017 from the National Institutes of Health. Its contents are solely the responsibility of the authors and do not necessarily represent the official views of the National Institutes of Health.

REFERENCES

1. E. Clarkson, M. A. Kupinski and H. H. Barrett, "Transformation of characteristic functionals through imaging systems," *Optics Express* 10 (2002) 536-539
2. M. A. Kupinski, E. Clarkson, J. Hoppin, and H. H. Barrett, "Experimental determination of object statistics from noisy images," *Journal of the Optical Society of America A* (accepted)
3. M. A. Kupinski, J. Hoppin, E. Clarkson, and H. H. Barrett, "Ideal observer computation using Markov-chain Monte Carlo," *Journal of the Optical Society of America A* (accepted)
4. E. Clarkson and H. H. Barrett, "Approximations to ideal-observer performance on signal-detection tasks," *Applied Optics* **39** (2000) 1783-1793
5. H. H. Barrett, C. K. Abbey and E. Clarkson, "Objective assessment of image quality. III. ROC metrics, ideal observers and likelihood-generating functions," *J. Opt. Soc. Am. A* **15** (1998) 1520-1535
6. E. Clarkson and H. H. Barrett, "Statistical decision theory and tumor detection," Chapter 4 in *Image Processing Techniques for Tumor Detection*, R. Strickland, ed. (New York, Dekker, 2002)
7. E. Clarkson, "Bounds on the area under the receiver operating characteristic curve for the ideal observer," *Journal of the Optical Society of America A* 19 (2002) 1963-1968
8. S. Kullback, *Information Theory and Statistics*, (Dover, New York, 1968)
9. J. Shapiro, "Bounds on the area under the ROC curve," *J. Opt. Soc. Am. A* **16** (1999) 53-57
10. M. V. Burnashev, "On one useful inequality in the testing of hypotheses," *IEEE Trans. Info. Th.* **44** (1998) 1668-1670