Published in final edited form as:

J Opt Soc Am A Opt Image Sci Vis. 2003 March ; 20(3): 430–438.

Ideal-observer computation in medical imaging with use of Markov-chain Monte Carlo techniques

Matthew A. Kupinski

Optical Sciences Center and Department of Radiology, University of Arizona, Tucson, Arizona 85721

John W. Hoppin

Program in Applied Mathematics, University of Arizona, Tucson, Arizona 85721

Eric Clarkson and Harrison H. Barrett

Optical Sciences Center and Department of Radiology, University of Arizona, Tucson, Arizona 85721

Abstract

The ideal observer sets an upper limit on the performance of an observer on a detection or classification task. The performance of the ideal observer can be used to optimize hardware components of imaging systems and also to determine another observer's relative performance in comparison with the best possible observer. The ideal observer employs complete knowledge of the statistics of the imaging system, including the noise and object variability. Thus computing the ideal observer for images (large-dimensional vectors) is burdensome without severely restricting the randomness in the imaging system, e.g., assuming a flat object. We present a method for computing the ideal-observer test statistic and performance by using Markov-chain Monte Carlo techniques when we have a well-characterized imaging system, knowledge of the noise statistics, and a stochastic object model. We demonstrate the method by comparing three different parallel-hole collimator imaging systems in simulation.

1. INTRODUCTION

We advocate the position that to properly quantify image quality one must clearly define the task that is to be performed and the observer who will be performing this task.¹ Tasks in medical imaging are typically detection tasks, e.g., determining whether a tumor is present, or estimation tasks, e.g., determining the value of a clinically relevant parameter. Observers can be human, human-model, or ideal observers. An ideal observer is a mathematical construct that is optimal in a well-defined sense.

The class of ideal observers includes the Bayesian ideal observer,² which sets an absolute upper limit on task performance as measured by most common figures of merit. (For simplicity, any reference to ideal observers from now on implies Bayesian ideal observers.) Because no other observer can outperform the ideal observer by these standards, it is reasonable to use ideal-observer performance as a figure of merit in comparing imaging systems. That is, the ideal observer is a useful tool for comparing and optimizing imaging systems. Furthermore, it is often useful to compare the performance of other observers relative to the performance of the best observer.

Address correspondence to Matthew A. Kupinski, Optical Sciences Center, University of Arizona, 1630 East University Boulevard, Tucson, Arizona 85721; e-mail, mkupinski@optics.arizona.edu.

The ideal observer requires complete knowledge of the statistics of the images or image data. Thus computing the ideal observer is often considered either too difficult or impossible for image data, which tend to be large; e.g., images of size 512×512 are characterized by a probability density function of a 512^2 random vector. Researchers often resort to unrealistic assumptions that dramatically simplify the math, such as assuming fixed backgrounds.³ Clarkson and Barrett⁴ worked out the ideal-observer test statistic for a number of situations in which the statistics of the backgrounds followed certain distributions, including Gaussian and Laplacian distributions. A great deal of research has been done on the performance of nonoptimal observers such as human observers and model observers in different types of backgrounds. See studies by Abbey and Barrett,⁵ Burgess *et al.*,⁶ and Bochud *et al.*⁷ for examples. However, the performance of the ideal observer has not been computed for complicated backgrounds that are not characterized by simple density functions.

In this paper we present a method for computing the ideal observer in simulation for a class of stochastic object models. The computation of the ideal observer is accomplished by using Markov-chain Monte Carlo techniques^{8,9} (to be referred to hereafter as Markov-Chain Monte Carlo). This method can be used for hardware optimization or comparison. In this paper we present the results of a simulation study in which we compare the ideal-observer performance of three different parallel-hole collimator imaging systems.

2. BACKGROUND

The process of data acquisition in imaging may be represented mathematically by

$$\mathbf{g} = H\mathbf{f} + \mathbf{n}, \quad (1)$$

where $\mathbf{f}$ is the continuous object being imaged, H is a continuous-to-discrete operator representing the imaging system, $\mathbf{n}$ is the measurement noise, and $\mathbf{g}$ is an $M \times 1$ vector of image data. In radiological imaging, the object $\mathbf{f}$ may represent the three-dimensional activity distribution of radio tracer within a body or the two-dimensional fluence of x-rays impinging on a detector. The imaging system H describes the mapping of continuous objects to discrete image data. The statistics of the random vector $\mathbf{n}$ may depend on the object $\mathbf{f}$, e.g., Poisson noise. The image data $\mathbf{g}$ may be an image as in the case of direct digital detectors or may require processing before being shown to a human observer as in the case of sinogram data. In either case the ideal observer uses the raw image data $\mathbf{g}$ and does not require any reconstruction or processing.

For a linear system, the continuous-to-discrete imaging operator H can be more clearly represented by

$$g_m = \int_S d\mathbf{r} h_m(\mathbf{r}) f(\mathbf{r}) + n_m, \quad (2)$$

where g_m and n_m are the m th elements of the vectors $\mathbf{g}$ and $\mathbf{n}$, respectively, $\mathbf{r}$ is a spatial coordinate (two or three dimensional), $f(\mathbf{r})$ is the object $\mathbf{f}$, $h_m(\cdot)$ is the m th detector sensitivity function, and S is the field of view. The entire operator H is composed of M functions $h_m(\cdot)$.

As discussed above, classification tasks in medicine consist of determining the presence of a signal, e.g., tumor. In this paper we consider signal-known-exactly tasks; we consider the tumor to be some known signal $\mathbf{f}_s$ in a random background $\mathbf{f}_b$. As a result there exist two possible hypotheses for any given image: signal absent (H_0) or signal present (H_1). We mathematically represent imaging under these two hypotheses as

$$H_0: \mathbf{g} = H\mathbf{f}_b + \mathbf{n}, \quad (3)$$

$$H_1: \mathbf{g} = \mathcal{H}(\mathbf{f}_b + \mathbf{f}_s) + \mathbf{n}. \quad (4)$$

For convenience we define the background image and signal image to be

$$\mathbf{b} = \mathcal{H}\mathbf{f}_b, \quad (5)$$

$$\mathbf{s} = \mathcal{H}\mathbf{f}_s. \quad (6)$$

A. Ideal Observer

The ideal observer employs the likelihood ratio, or any monotonic transformation thereof, to make decisions.¹⁰ The likelihood ratio is defined as

$$\Lambda(\mathbf{g}) = \frac{\text{pr}(\mathbf{g}|H_1)}{\text{pr}(\mathbf{g}|H_0)}, \quad (7)$$

where $\text{pr}(\mathbf{g}|H_i)$ is the probability density of the image data under hypothesis H_i . An observer compares the test statistic, e.g., $\Lambda(\mathbf{g})$, to a predetermined critical value (or threshold) and makes a decision accordingly. A fixed critical value corresponds to a single point on a receiver operating characteristic (ROC) curve.^{11,12} An ROC curve is generated by varying this critical value over the real line. The area under the ROC curve (AUC)¹² is a commonly used performance metric that is maximized by the ideal observer.

Traditionally, to compute the ideal observer, researchers were forced to make strict and unrealistic assumptions regarding the imaging system. For example, if we were to assume that the background $\mathbf{b}$ and the signal $\mathbf{s}$ were nonrandom quantities, then the randomness in $\mathbf{g}$ would be due only to noise. Hence the densities $\text{pr}(\mathbf{g}|H_i)$ would be known, and the likelihood ratio could be easily calculated for various noise models. However, the stochastic nature of normal patient anatomy has a dramatic effect on the performance of an observer.^{5,13} For example, a lesion in a mammogram may be very easy or very difficult to detect depending upon whether there is dense tissue surrounding the tumor. Ultimately, if the task is to detect a lesion in a mammogram, then the background structure of a mammogram must be accounted for.⁶

B. Object Models

Characterizing the full probability density of $\mathbf{b}$ is a very difficult task, which many researchers have addressed.^{14,15} However, the probability density $\text{pr}(\mathbf{b})$ is particular to the imaging system operator $\mathcal{H}$ since $\mathbf{b} = \mathcal{H}\mathbf{f}_b$, and is not especially useful if we wish to perform hardware optimizations. We require a characterization of the randomness in the background *objects* $\mathbf{f}_b$ and not the background *images* $\mathbf{b}$. However, for the sake of notational simplicity we will often refer to $\text{pr}(\mathbf{f}_b|\mathcal{H})$ as $\text{pr}(\mathbf{b})$, which we can know in simulation given knowledge of our object model and a complete description of the imaging system $\mathcal{H}$.

We will employ a stochastic object model known as the lumpy object model, developed by Rolland and Barrett,¹³ although this method can easily be extended to other parameterized object models. The lumpy object model may be represented by the equation

$$\mathbf{f}_b = f_b(\mathbf{r}) = \sum_{n=1}^N L(\mathbf{r} - \mathbf{c}_n|a, s), \quad (8)$$

where $\mathbf{r}$ is a two- or three-dimensional spatial coordinate, N is the random number of lumps in the object, $L(\cdot)$ is the lump function, $\mathbf{c}_n$ is the center of the n th lump, a is the magnitude of the lump function, and s is the width. Typically $L(\cdot)$ is chosen to be a Gaussian function, but other functions can be chosen as well. The variable N is a Poisson-distributed random variable with

mean $\bar{N}$, and the centers $\mathbf{c}_n$ are uniformly distributed within the field of view. Note that the randomness in this object model is completely characterized by the random variables N and $\{\mathbf{c}_n\}$. Examples of two-dimensional lumpy objects with different parameters are given in Fig. 1.

It follows from Eqs. (2) and (5) that the m th background element is given by

$$b_m = \sum_{n=1}^N \int_s d\mathbf{r} h_m(\mathbf{r}) L(\mathbf{r} - \mathbf{c}_n | a, s). \quad (9)$$

Thus, to generate a single pixel of a noise-free and signal-free background image, we need to compute the N inner products of a sensitivity function with the lumps. To generate an entire image, we compute these inner products over all M pixels, i.e., all sensitivity functions $h_m(\cdot)$.

The lumpy object model is not appropriate for all imaging problems. The methodology we are about to describe, however, can be applied to many parameterized object models. We show in this paper that we can compute the ideal observer for this lumpy object model. This is a clear step forward for ideal-observer computations, which have traditionally been limited to flat or Gaussian background models. In future papers we will discuss the extension of this work to different object models.

3. METHOD

Using our knowledge of the imaging system, its associated noise properties, and the lumpy object model, we would like to calculate likelihood ratios for a set of simulated image data $\{\mathbf{g}\}$. We begin by rewriting Eq. (7) as

$$\Lambda(\mathbf{g}) = \frac{\int d\mathbf{b} \text{pr}(\mathbf{b}) \text{pr}(\mathbf{g}|\mathbf{b}, H_1)}{\int d\mathbf{b} \text{pr}(\mathbf{b}) \text{pr}(\mathbf{g}|\mathbf{b}, H_0)}. \quad (10)$$

For a given background $\mathbf{b}$, the randomness in $\mathbf{g}$ is due only to the noise, and hence $\text{pr}(\mathbf{g}|\mathbf{b}, H_i)$ characterizes the noise distribution of the imaging system, which we know *a priori*. The numerator and denominator of Eq. (10) turn out to be exceptionally small numbers, because they are integrals over a product of high-dimensional probability density functions. Thus estimating the ratio accurately by computing the numerator and denominator separately is infeasible.

To make computation of the likelihood ratio possible,¹⁶ we rewrite Eq. (10) as

$$\Lambda(\mathbf{g}) = \int d\mathbf{b} \frac{\text{pr}(\mathbf{b}) \text{pr}(\mathbf{g}|\mathbf{b}, H_1)}{\int d\mathbf{b}' \text{pr}(\mathbf{b}') \text{pr}(\mathbf{g}|\mathbf{b}', H_0)}, \quad (11)$$

where we have taken the denominator into the integral. Next, we multiply and divide by $\text{pr}(\mathbf{g}|\mathbf{b}, H_0)$, which leads to

$$\Lambda(\mathbf{g}) = \int d\mathbf{b} \frac{\text{pr}(\mathbf{g}|\mathbf{b}, H_1)}{\text{pr}(\mathbf{g}|\mathbf{b}, H_0)} \frac{\text{pr}(\mathbf{b}) \text{pr}(\mathbf{g}|\mathbf{b}, H_0)}{\int d\mathbf{b}' \text{pr}(\mathbf{b}') \text{pr}(\mathbf{g}|\mathbf{b}', H_0)}. \quad (12)$$

Using Bayes's law we are able to rewrite Eq. (12) as

$$\Lambda(\mathbf{g}) = \int d\mathbf{b} \Lambda_{\text{BKE}}(\mathbf{g}|\mathbf{b}) \text{pr}(\mathbf{b}|\mathbf{g}, H_0), \quad (13)$$

where $\Lambda_{\text{BKE}}(\mathbf{g}|\mathbf{b})$ is the background-known-exactly (BKE) likelihood ratio defined as

$$\Lambda_{\text{BKE}}(\mathbf{g}|\mathbf{b}) = \frac{\text{pr}(\mathbf{g}|\mathbf{b}, H_1)}{\text{pr}(\mathbf{g}|\mathbf{b}, H_0)} \quad (14)$$

and

$$\text{pr}(\mathbf{b}|\mathbf{g}, H_0) = \frac{\text{pr}(\mathbf{g}|\mathbf{b}, H_0)\text{pr}(\mathbf{b})}{\int d\mathbf{b}' \text{pr}(\mathbf{g}|\mathbf{b}', H_0)\text{pr}(\mathbf{b}')} \quad (15)$$

The BKE likelihood ratio is the ideal-observer test statistic in the situation where the background is fixed and not random. Thus we have derived the random-background likelihood ratio in terms of an integral of the BKE likelihood ratio (*a posteriori* mean).

For an imaging system with Poisson noise

$$\text{pr}(\mathbf{g}|\mathbf{b}, H_1) = \prod_{m=1}^M \exp[-(b_m + s_m)] \frac{(b_m + s_m)^{g_m}}{g_m!}, \quad (16)$$

and

$$\text{pr}(\mathbf{g}|\mathbf{b}, H_0) = \prod_{m=1}^M \exp(-b_m) \frac{(b_m)^{g_m}}{g_m!}. \quad (17)$$

We are thus able to rewrite Eq. (14) as

$$\Lambda_{\text{BKE}}(\mathbf{g}|\mathbf{b}) = \prod_{m=1}^M \left(1 + \frac{s_m}{b_m}\right)^{g_m} \exp(-s_m). \quad (18)$$

Similarly, for an imaging system with Gaussian noise with covariance K_n , the BKE likelihood ratio becomes

$$\Lambda_{\text{BKE}}(\mathbf{g}|\mathbf{b}) = \exp[(\mathbf{g} - \mathbf{b} - \mathbf{s}/2)^\dagger K_n^{-1} \mathbf{s}], \quad (19)$$

where $\mathbf{x}^\dagger$ represents the adjoint of $\mathbf{x}$.

Ideally we would use Monte Carlo techniques to calculate the integral in Eq. (13) by sampling backgrounds from $\text{pr}(\mathbf{b}|\mathbf{g}, H_0)$, evaluating $\Lambda_{\text{BKE}}(\mathbf{g}|\mathbf{b})$, and averaging. Sampling from $\text{pr}(\mathbf{b}|\mathbf{g}, H_0)$ is difficult because the density function is generally not known. As a result, we consider alternate methods. The advantage that we have, and must use, is our knowledge of our model for $\mathbf{f}_b$. We are representing our background objects $\mathbf{f}_b$ as lumpy backgrounds that are completely characterized by the number of lumps N (following a Poisson distribution) and the centers $\{\mathbf{c}_n\}$ (uniformly distributed throughout the field of view) for fixed a , s , lump function $L(\cdot)$, and field of view S . We represent the stochastic parameters as an unordered set $\boldsymbol{\theta} = \{\mathbf{c}_1, \mathbf{c}_2, \dots, \mathbf{c}_N\}$.

We know that for any given $\boldsymbol{\theta}$ there exists a unique $\mathbf{b}$. This relationship allows us to rewrite Eq. (13) as

$$\Lambda(\mathbf{g}) = \int d\boldsymbol{\theta} \Lambda_{\text{BKE}}(\mathbf{g}|\mathbf{b}(\boldsymbol{\theta})) \text{pr}(\boldsymbol{\theta}|\mathbf{g}, H_0), \quad (20)$$

where

$$\text{pr}(\theta|\mathbf{g}, H_0) = \frac{\text{pr}(\mathbf{g}|\mathbf{b}(\theta), H_0)\text{pr}(\theta)}{\int d\theta' \text{pr}(\mathbf{g}|\mathbf{b}(\theta'), H_0)\text{pr}(\theta')}. \quad (21)$$

Unlike $\text{pr}(\mathbf{b})$, the function $\text{pr}(\theta)$ has a simple analytic form that we know from the definition of our object model. For a more detailed explanation of the replacement of $\mathbf{b}$ with θ in the above expressions we refer the reader to Appendix A.

A. Markov-Chain Monte Carlo

We want to estimate the likelihood ratio in Eq. (20) by using Monte Carlo integration, i.e.,

$$\widehat{\Lambda}(\mathbf{g}) = \frac{1}{J} \sum_{j=1}^J \Lambda_{\text{BKE}}(\mathbf{g}|\mathbf{b}(\theta^{(j)})) \quad (22)$$

where each $\theta^{(j)}$ is a sample from the density $\text{pr}(\theta|\mathbf{g}, H_0)$. To generate the $\theta^{(j)}$, we construct a Markov chain with $\text{pr}(\theta|\mathbf{g}, H_0)$ as the stationary density for the chain. We choose an initial parameter vector $\theta^{(0)}$ and a proposal density $q(\theta|\theta^{(i)})$ for the Markov chain. Given $\theta^{(i)}$ we draw a sample parameter vector $\tilde{\theta}$ from the proposal density and add it to our Markov chain with probability

$$\tilde{\theta} \text{ accepted} = \min \left[1, \frac{\text{pr}(\tilde{\theta}|\mathbf{g}, H_0)q(\theta^{(i)}|\tilde{\theta})}{\text{pr}(\theta^{(i)}|\mathbf{g}, H_0)q(\tilde{\theta}|\theta^{(i)})} \right]. \quad (23)$$

If $\tilde{\theta}$ is accepted, then $\theta^{(i+1)}$ is $\tilde{\theta}$; if $\tilde{\theta}$ is not accepted, then $\theta^{(i+1)}$ is $\theta^{(i)}$. The Markov chain is the sequence consisting of

$$\theta^{(0)}, \theta^{(1)}, \dots, \theta^{(K)}, \quad (24)$$

We adopt a Metropolis–Hastings approach⁸ and design our proposal density to be symmetric, i.e., $q(\theta|\theta^{(i)}) = q(\theta^{(i)}|\theta)$, resulting in a cancellation of the proposal densities in Eq. (23). Furthermore, we can rewrite $\text{pr}(\theta|\mathbf{g}, H_0)$ using Bayes's rule [Eq. (21)] and cancel the denominators, leaving an expression that we can compute exactly for a given noise model and object model. In the case of Poisson noise (which we will use for our simulation studies below), the expression in Eq. (23) becomes a ratio of

$$\begin{aligned} \text{pr}(\mathbf{g}|\mathbf{b}(\theta), H_0)\text{pr}(\theta) &= \text{pr}(\mathbf{g}|\mathbf{b}(\theta), H_0)\text{pr}(N)\text{pr}(\{\mathbf{c}_n\}) \\ &= \left\{ \prod_{m=1}^M \exp[-b_m(\theta)] \frac{b_m(\theta)^{g_m}}{g_m!} \right\} \\ &\times \left[\exp(-\bar{N}) \frac{\bar{N}^N}{N!} \right] \left(\frac{N!}{\text{FOV}^N} \right) \end{aligned} \quad (25)$$

evaluated at $\tilde{\theta}$ and $\theta^{(i)}$, respectively. The bracketed terms in Eq. (25) are $\text{pr}(\mathbf{g}|\mathbf{b}(\theta), H_0)$, $\text{pr}(N)$, and $\text{pr}(\{\mathbf{c}_n\})$, respectively. The final $N!$ in Eq. (25) arises because $\text{pr}(\{\mathbf{c}_n\})$ is invariant under permutations of $\{\mathbf{c}_n\}$ (see Appendix A). We are exploiting our lumpy object model in which we know that the probability of a set of N centers $\{\mathbf{c}_n\}$ is given by the probability that we choose N lumps (Poisson with mean $\bar{N}$) times the probability that we observed the centers $\{\mathbf{c}_n\}$ (uniform within the FOV and invariant under permutations).

In practice it is difficult to design a proposal density $q(\tilde{\theta}|\theta^{(i)})$ that is symmetric because the number of elements in $\tilde{\theta}$ and $\theta^{(i)}$ can be different. Furthermore, we cannot forgo symmetry,

and we choose to compute the ratio of the proposal densities because it is difficult to design a practical, yet easily computable, density of an unordered set conditioned on another unordered set. We avoid both of these difficulties by designing a proposal density on a matrix Φ , i.e., $q(\Phi|\Phi^{(i)})$. The matrix Φ is designed such that we can write θ as a function of Φ , i.e., $\theta(\Phi)$. The matrix Φ has the advantage that the number and position of elements is fixed. It has the disadvantage that not every possible θ can be represented as $\theta(\Phi)$. However, the states $\mathbf{b}$ not achievable with Φ are those for which $\text{pr}(\mathbf{b}|\mathbf{g}, H_0)$ is zero for all practical purposes (i.e., beyond computer precision).

This matrix Φ is composed of a binary column vector α of dimension N' followed by a list of centers of dimension $N' \times 2$ for two-dimensional lumpy objects and $N' \times 3$ for three-dimensional lumpy objects, i.e.,

$$\Phi = \begin{bmatrix} \alpha_1 & \mathbf{c}_1 \\ \alpha_2 & \mathbf{c}_2 \\ \vdots & \vdots \\ \alpha_{N'} & \mathbf{c}_{N'} \end{bmatrix}, \quad (26)$$

where the centers $\mathbf{c}_n$ are row vectors. Here we are using an unusual notation by inserting multiple elements into the matrix Φ using the row vectors $\mathbf{c}_n$. However, this keeps our definition of Φ general for two- or three-dimensional lumpy objects. The mapping function $\theta(\Phi)$ is defined as

$$\theta(\Phi) = \{\mathbf{c}_n; \alpha_n = 1\}. \quad (27)$$

The proposal density $q(\Phi|\Phi^{(i)})$ is sampled by first flipping (i.e., a 1 to a 0 or a 0 to a 1) each of the binary numbers α_n with probability η . Then we randomly choose one center that is turned on (i.e., $\alpha_n = 1$) in both Φ and $\Phi^{(i)}$ and randomly shift the position of this center by sampling an isotropic Gaussian distribution with fixed width σ . In Appendix B we write down an expression of the density $q(\Phi|\Phi^{(i)})$ and show that it is symmetric. The flipping of the α_n allows us to add or remove lumps, and the shifting of one of the centers allows us to make a small change to a given lump. We are limited here to N' lumps, but we choose N' to be much larger than the number of lumps N that went into the generation of the image $\mathbf{g}$. The probability that we will ever accept N' lumps is beyond the floating-point range of computers. Finally, we have a symmetric proposal density and are left only to compute the ratio $\text{pr}(\theta(\Phi)|\mathbf{g}, H_0)/\text{pr}(\theta(\Phi^{(i)})|\mathbf{g}, H_0)$, the numerator and denominator of which are shown in Eq. (25). Note that there are typically only a couple of different lumps between $\mathbf{b}(\theta(\Phi^{(i+1)}))$ and $\mathbf{b}(\theta(\Phi^{(i)}))$. Thus to generate the new background $\mathbf{b}^{(i+1)}$ [Eq. 9] from the old background $\mathbf{b}^{(i)}$, we need only to replace the few lumps that have changed, which dramatically speeds up the simulation procedure.

4. SIMULATIONS

The ultimate goal of our work is to compare imaging systems based on ideal-observer performance. As an example, we simulate parallel-hole collimator imaging systems for nuclear medicine. We vary the parameters of the collimator, e.g., bore length and height, to achieve different noise levels and resolution in the generated images. More specifically, we model the imaging system transfer functions $h_m(\cdot)$ as Gaussian centered on the m th pixel. Each imaging system has different widths and heights for the $h_m(\cdot)$ functions, corresponding to different collimator parameters. By choosing Gaussian transfer functions and the lumpy object model with Gaussian lumps, we are able to compute a closed-form solution for the background, using Eq. (9).

There are numerous practical issues in Markov-chain Monte Carlo that need to be addressed. For example, if $\Phi^{(0)}$ is a completely random realization, then one must “burn in” the chain for a number of iterations to remove bias due to the initial samples. Furthermore, we must choose an N' , a flip probability η , and a shift width σ for the random shift. We are aided in this problem because we generate the image $\mathbf{g}$ from the lumpy model and therefore know the number of lumps N^* and centers $\{\mathbf{c}_n^*\}$ that went into making the image. Thus the first N^* entries into $\Phi^{(0)}$ are randomly shifted versions of $\{\mathbf{c}_n^*\}$, and they are all turned on, i.e., $\alpha_n = 1$ for $n = 1, \dots, N^*$. We chose N' to be 100 times bigger than N^* , with the remaining $100N^* - N^*$ centers in $\Phi^{(0)}$ being uniformly distributed within the field of view. These remaining centers are all initially turned off, i.e., $\alpha_n = 0$ for $n = (N^* + 1), \dots, N'$. The flip probability η was chosen to be $0.04/N'$, which makes flips a rare occurrence. The Gaussian shift density was chosen to have a width of $s/10$, where s is the width of the lump function $L(\cdot)$ [Eq. (8)]. Because our initial $\Phi^{(0)}$ is chosen with knowledge of the centers that went into the creating of $\mathbf{g}$, our burn-in times are relatively short (500 iterations). Finally, we iterated the chain for a total of 150,000 iterations, computing $\Lambda_{\text{BKE}}(\cdot)$ at each iteration.

The parameter N' was chosen because the probability that the chain would ever accept ten times more lumps than went into the creation of the image is effectively zero. The parameters η and σ were determined empirically by studying the rejection rate of the chain. For example, if 99 out of 100 proposals are rejected on average, then the chain is not very efficient. Conversely, if the chain accepts all the proposed solutions, then it is likely that insignificant changes are occurring. This insignificance might be caused by a small σ . The burn-in of 500 iterations was determined by visual examination of $\Lambda_{\text{BKE}}(\cdot)$ as the chain progressed. It was clear that after a couple of hundred iterations, the chain had converged to the stationary density. If we were to choose $\Phi^{(0)}$ at random we would need to lengthen the burn-in time for our chain because it would take longer to arrive at the peak of the stationary density. Finally, the total number of iterations (150,000) was chosen because it produced estimates of $\Lambda(\mathbf{g})$ with little variation.

We simulated three different imaging systems (A, B, and C) with different widths and heights for the sensitivity functions given in Table 1. The images acquired by all three imaging systems are 64×64 , i.e., $\mathbf{g}$ is a $64^2 \times 1$ vector. For each imaging system we generated 50 signal-absent images and 50 signal-present images. We used a lumpy object model with $N = 25$, $a = 1$, and $s = 7$ [see Eq. (8)]. Example images of the same lumpy object from our three different imaging systems are shown in Fig. 2. For the signal we used a Gaussian with width 3 and height 0.1 located in the center of the object. We found that these signals were difficult to detect visually, although a psychophysical experiment was not performed. They are, however, often detected by the ideal observer.

As described in Subsection 3.A, we generated a Markov chain in order to sample from $\text{pr}(\theta | \mathbf{g}, H_0)$ in Eq. (20) to approximate the ideal observer by using Eq. (22). For each of the 300 images generated (100 for each imaging system), we generated 150,000 iterations of the Markov chain. In Fig. 3 we show $\Lambda_{\text{BKE}}(\mathbf{g} | \mathbf{b}(\theta(\Phi^{(i)})))$ evaluated at the iterations of an example Markov chain. The image $\mathbf{g}$ used to generate this plot was without signal and was acquired from imaging system B. In all of our calculations we ignore the first 500 iterations of the chain to allow for burn-in, and we average the $\Lambda_{\text{BKE}}(\cdot)$ of the remaining 149,500 iterations.

5. RESULTS

We calculated the ideal observer for 50 signal-absent and 50 signal-present images for three different imaging systems. Histograms of the log-likelihood ratio are given in Fig. 4. Imaging systems A and B seem to have substantial separation between the signal-absent and the signal-present data [Figs. 4(a) and (b)]. There is a great deal of overlap between the two histograms for imaging system C [Fig. 4(c)]. Larger separation indicates that the ideal observer can more

easily detect the signal. We used ROC analysis to quantify the performance of the ideal observer for these three imaging systems. Figure 5 shows ROC curves for the three imaging systems. These ROC curves were fitted by using LABROC4 software.¹⁷ The AUCs for imaging systems A, B, and C are 0.92, 0.88, and 0.67, respectively. These values provide a means of ranking imaging systems A, B, and C for the aforementioned detection task.

Because our method returns estimates of likelihood ratios, we reran the above experiment five times using the same 100 images from imaging system A. The five different runs used different random seeds, and thus the chains progressed differently. This experiment checks the interobserver variability that arises from the stochastic nature of Markov-chain Monte Carlo. That is, we wish to isolate the component of the variability that is due to the technique. We found that the variability in AUC as measured by the standard deviation of the five different AUC values was 0.006. This number could be reduced by running the chain for more than 150,000 iterations. This number does not reveal any information about the variability that will arise if different images (or more images) from system A are used. Similar runs were performed for imaging systems B and C with similar results.

6. CONCLUSIONS

We have presented a method for comparing imaging hardware based on the performance of the ideal observer on a signal-known-exactly detection task. The random backgrounds used in our simulations were generated by using the lumpy object model and simulated imaging systems. We regard this work as a first step toward calculations that would involve randomly varying signals, more realistic imaging-system models, and more complicated object models. Therefore the results for the particular imaging systems investigated here should be viewed as proof of concept rather than the final word on the trade-off between noise and resolution in collimator design.

Acknowledgements

We thank Hongbin Zhang and Brandon Gallas for helpful discussions and extensive work on very similar topics and Thomas Kennedy for sharing his expertise of Markov-chain Monte Carlo with us. This work was supported by National Science Foundation grant 9977116 and National Institutes of Health grants P41 RR14304, KO1 CA87017, and RO1 CA52643.

References

1. Barrett HH. Objective assessment of image quality: effects of quantum noise and object variability. *J Opt Soc Am* 1990;A 7:1266–1278.
2. Van Trees, HL. *Detection, Estimation, and Modulation Theory, Part I*. Academic; New York: 1968.
3. Fukunaga, K. *Statistical Pattern Recognition*. Academic; San Diego, Calif: 1990.
4. Clarkson E, Barrett HH. Approximation to ideal-observer performance on signal-detection tasks. *Appl Opt* 2000;39:1783–1794. [PubMed: 18345075]
5. Abbey CK, Barrett HH. Human- and model-observer performance in ramp-spectrum noise: effects of regularization and object variability. *J Opt Soc Am* 2001;A 18:473–488.
6. Burgess AE, Jacobson FL, Judy PF. Human observer detection experiments with mammograms and power-law noise. *Med Phys* 2001;28:419–439. [PubMed: 11339738]
7. Bochud FO, Abbey CK, Eckstein MP. Visual signal detection in structured backgrounds III. Calculation of figures of merit for model observers in statistically nonstationary backgrounds. *J Opt Soc Am A* 2000;17:193–206.
8. Robert, CP.; Casella, G. *Monte Carlo Statistical Methods*. Springer-Verlag; New York: 1999.
9. Gilks, WR.; Richardson, S.; Spiegelhalter, DJ., editors. *Markov Chain Monte Carlo in Practice*. Chapman & Hall; Boca Raton, Fla: 1996.

10. Barrett HH, Abbey CK, Clarkson E. Objective assessment of image quality III. ROC metrics, ideal observers, and likelihood-generating functions. *J Opt Soc Am* 1998;A 15:1520–1535.
11. Egan, J. *Signal Detection Theory and ROC Analysis*. Academic; New York: 1975.
12. Metz CE. Basic principles of ROC Analysis. *Semin Nucl Med* 1978;VIII:283–298. [PubMed: 112681]
13. Rolland JP, Barrett HH. Effect of random background inhomogeneity on observer detection performance. *J Opt Soc Am* 1992;A 9:649–658.
14. Zhu SC, Wu Y, Mumford D. Filters, random fields and maximum entropy (FRAME). *Int J Comput Vision* 1998;27:1–20.
15. Simoncelli EP, Olshausen B. Natural image statistics and neural representation. *Annu Rev Neurosci* 2001;24:1193–1217. [PubMed: 11520932]
16. H. Zhang and B. Gallas, University of Arizona, Tucson, Ariz. 85721 (personal communication, 2001).
17. Metz CE, Herman BA, Shen JH. Maximum likelihood estimation of receiver operating characteristic (ROC) curves from continuously-distributed data. *Stat Med* 1998;17:1033–1053. [PubMed: 9612889]
18. Dugundji, J. *Topology*. Allyn and Bacon; Boston, Mass: 1965.

APPENDIX A: REPLACING $\mathbf{b}$ WITH θ

We would like to show the equivalence of Eqs. (13) and (20) for the backgrounds generated with the lumpy object model. We use a standard method of transforming density functions given by

$$\text{pr}(\mathbf{b}|\mathbf{g}, H_0) = \int d\theta \delta^M(\mathbf{b} - \mathbf{b}(\theta)) \text{pr}(\theta|\mathbf{g}, H_0), \quad (\text{A1})$$

where M is the dimension of $\mathbf{b}$. In general this integral is difficult because knowledge of the roots of $\mathbf{b} = \mathbf{b}(\theta)$ is required. However, in our problem, $\text{pr}(\mathbf{b}|\mathbf{g}, H_0)$ is within an integrand, allowing us to compute the following:

$$\Lambda(\mathbf{g}) = \int d\mathbf{b} \Lambda_{\text{BKE}}(\mathbf{g}|\mathbf{b}) \text{pr}(\mathbf{b}|\mathbf{g}, H_0) \quad (\text{A2})$$

$$= \int d\mathbf{b} \int d\theta \Lambda_{\text{BKE}}(\mathbf{g}|\mathbf{b}) \times \delta^M(\mathbf{b} - \mathbf{b}(\theta)) \text{pr}(\theta|\mathbf{g}, H_0) \quad (\text{A3})$$

$$= \int d\theta \Lambda_{\text{BKE}}(\mathbf{g}|\mathbf{b}(\theta)) \text{pr}(\theta|\mathbf{g}, H_0), \quad (\text{A4})$$

This integral expression is shorthand for the following:

$$\sum_N \text{pr}(N|\mathbf{g}, H_0) \int_{S_N} d\mathbf{c}_1 d\mathbf{c}_2 \dots d\mathbf{c}_N \times \Lambda_{\text{BKE}}(\mathbf{g}|\mathbf{b}(\{\mathbf{c}_n\})) \text{pr}(\{\mathbf{c}_n\}|\mathbf{g}, H_0). \quad (\text{A5})$$

The above integral is over a region S_N contained within R^{2N} for two-dimensional lumpy objects. This region is defined by the requirement that all centers be within the FOV and must satisfy the inequalities $\mathbf{c}_1 < \mathbf{c}_2 < \dots < \mathbf{c}_N$, where $<$ symbolizes a lexicographical ordering¹⁸ of the vectors $\mathbf{c}_n$. This ensures that there is only one point in S_N corresponding to the unordered set $\{\mathbf{c}_n\}$. Also note that the volume of this set S_N is given by $\text{FOV}^N/N!$, which explains the expression $\text{pr}(\{\mathbf{c}_n\})$ in Eq. (25).

APPENDIX B: SYMMETRY OF THE PROPOSAL DENSITY

The proposal density is given by

$$q(\tilde{\Phi}|\Phi) = \eta^f (1-\eta)^{N'-f} \sum_{j: \tilde{\alpha}_j = \alpha_j = 1} \frac{1}{C} \delta(\mathbf{c}_{-j} - \tilde{\mathbf{c}}_{-j}) \times G(\tilde{\mathbf{c}}_j - \mathbf{c}_j), \quad (\text{B1})$$

where f is the number of α 's that are flipped from Φ to $\tilde{\Phi}$ given by

$$f = \sum_{i=1}^{N'} |\alpha_i - \tilde{\alpha}_i| \quad (\text{B2})$$

The function $G(\cdot)$ is a properly normalized, symmetric Gaussian density with variance σ^2 . The variable C is the number of terms in the sum in Eq. (B1), given by

$$C = \alpha^\dagger \tilde{\alpha}. \quad (\text{B3})$$

In the case in which there are no centers in common between Φ and $\tilde{\Phi}$, the sum in Eq. (B1) is zero; i.e., its probability density is zero. For two-dimensional (respectively three-dimensional), lumpy objects, $\delta(\mathbf{c}_{-j} - \tilde{\mathbf{c}}_{-j})$ is a $2C - 2$ (respectively $3C - 3$)-dimensional Dirac delta function. The vector $\mathbf{c}_{-j}$ is a concatenation of all the center vectors $\mathbf{c}_i$ satisfying $\alpha_i = \alpha_j = 1$, except $\mathbf{c}_j$ itself. The vector $\tilde{\mathbf{c}}_{-j}$ is defined similarly.

The term to the left of the summation in Eq. (B1) gives the probability of making the flips in the α_i to get from Φ to $\tilde{\Phi}$. The summation term is a mixture of densities that represents the shift of one of the centers in common between Φ and $\tilde{\Phi}$. To sample from such a mixture, we choose one of the terms with probability $1/C$ and then sample from it.

The symmetry of $q(\tilde{\Phi}|\Phi)$ follows from the symmetry of the f [Eq. (B2)], C [Eq. (B3)], $\delta(\mathbf{c}_{-j} - \tilde{\mathbf{c}}_{-j})$, and $G(\tilde{\mathbf{c}}_j - \mathbf{c}_j)$ regarded as functions of Φ and $\tilde{\Phi}$.

(a)

(b)

Fig. 1. Examples of two-dimensional lumpy objects. (a) Parameters $N = 39$, $a = 1$, and $s = 10$; (b) parameters $N = 348$, $a = 1$, and $s = 4$. The field of view (FOV) for both objects is 128×128 .

(a)

(b)

(c)

Fig. 2. Image realizations of the same lumpy object for (a) imaging system A, (b) imaging system B, and (c) imaging system C. There is no signal in these images.

Fig. 3.

Λ_{BKE} versus iteration number. The iteration numbers here are the states of the Markov chain described in Subsection 3.A.

Fig. 4. Log-likelihood histograms for (a) imaging system A, (b) imaging system B, and (c) imaging system C.

Fig. 5. ROC curve comparison of the ideal-observer performance for the three imaging systems. The AUC values for imaging systems A, B, and C are 0.92, 0.88, and 0.67, respectively.

Table 1

Characteristics of the Three Imaging Systems

Imaging System	$h_m(\cdot)$ Width	$h_m(\cdot)$ Height	AUC
A	0.5	40	0.92
B	2.5	100	0.88
C	5	200	0.67