



Published in final edited form as:

J Opt Soc Am A Opt Image Sci Vis. 2007 December ; 24(12): B99–B109.

Analysis of observer performance in unknown-location tasks for tomographic image reconstruction

Anastasia Yendiki^{1,*} and Jeffrey A. Fessler²

¹ HMS/MGH/MIT Martinos Center for Biomedical Imaging, 149 13th Street, Charlestown, Massachusetts 02129, USA

² Department of Electrical Engineering and Computer Science, The University of Michigan, 4240 EECS Building, 1301 Beal Avenue, Ann Arbor, Michigan 48109-2122, USA

Abstract

Our goal is to optimize regularized image reconstruction for emission tomography with respect to lesion detectability in the reconstructed images. We consider model observers whose decision variable is the maximum value of a local test statistic within a search area. Previous approaches have used simulations to evaluate the performance of such observers. We propose an alternative approach, where approximations of tail probabilities for the maximum of correlated Gaussian random fields facilitate analytical evaluation of detection performance. We illustrate how these approximations, which are reasonably accurate at low probability of false alarm operating points, can be used to optimize regularization with respect to lesion detectability.

1. INTRODUCTION

Several applications of emission tomography involve the detection of a spatially localized target signal in an image reconstructed from noisy data. When choosing among different reconstruction methods or tuning the parameters of individual methods, a reasonable approach is to seek the choice that leads to optimal performance in such detection tasks. In this work we focus on penalized-likelihood reconstruction methods for emission tomography. These methods involve one or more regularization parameters that control the noise-resolution trade-off in the reconstructed images. Instead of choosing the regularization parameters by numerical criteria, such as cross validation or *L*-curves [1,2], we would like to optimize these parameters analytically with respect to signal detectability.

In clinical practice, detection tasks are usually performed by human observers. The performance of humans can be evaluated empirically by tracing their *receiver operating characteristic* (ROC) in simple detection tasks [3–5], or their *localization ROC* (LROC) in tasks that involve both detection and localization [6]. However, the experiments required for such an evaluation are too time-consuming to perform for many values of a reconstruction parameter or to repeat every time that some aspect of the imaging process changes. Thus we turn to the mathematical observers that have been proposed in the literature to model human performance [7] and that allow analytical treatment.

Detection tasks where the observer knows the possible location of the target signal *a priori* have been analyzed with respect to optimal regularization methods [8–12]. Known-location tasks are also the ones for which model and human observer correlation has been

*Corresponding author: ayendiki@nmr.mgh.harvard.edu.

investigated most extensively. In these tasks human observers appear to be effective in compensating for second-order image statistics. As a result, they are modeled well by mathematical observers that perform prewhitening, such as the *channelized Hotelling observer* (CHO) [13–18]. However, optimizing regularization parameters has less of an effect on the detection performance of such observer models, as has been concluded both experimentally [17] and analytically [19]. Thus it is of interest to optimize such parameters with respect to detection performance in tasks where the prewhitening capabilities of humans degrade. Such a degradation seems to occur in tasks where the location of the target signal is not known *a priori* [20,21].

Location uncertainty complicates the analysis of detectability. Observer models that have been compared to humans in unknown-location tasks usually base their decisions on the maximum value of a local test statistic over all possible signal locations. The exact distribution of the maximum of a correlated random field has the form of a multiple integral that is intensive to compute. A “brute-force” approach to evaluating the performance of such an observer model would be to perform a large number of time-consuming tomographic reconstructions of Monte Carlo–simulated projection data and produce realizations of the maximum test statistic from the reconstructed images. When optimizing some reconstruction parameter with respect to detection performance, this simulation would have to be repeated for every value of the parameter of interest.

To avoid performing multiple reconstructions, investigators have proposed analytical approximations of the moments of the local test statistics and used them to produce Monte Carlo–simulated realizations of the maximum test statistic directly [22,23]. This approach results in very significant time savings when compared to the brute-force method, but it still requires performing multiple simulations when the evaluation of detection performance needs to be repeated for various values of one or more parameters of interest.

We propose an alternative approach that does not require simulations. Although the exact expression for the distribution of the maximum is complicated, simple approximations of this distribution at high thresholds for correlated Gaussian random fields have been proposed by Adler [24]. These approximations assume (local) stationarity, so they are more accurate under the signal-absent than the signal-present hypothesis. The latter requires a different treatment, and we adopt the approach of Siegmund and Worsley to analyze it [25]. We illustrate how these analytical results, which are accurate at low probability of false alarm operating points, can be combined to approximate the performance of the observer models of interest as a function of regularization parameters.

This paper is structured as follows: Section 2 establishes the models that we will consider for the object, imaging system, reconstruction method, and observer. Section 3 reviews the tail probability approximations that we apply in Section 4 to analyze the detection performance of the observer models of interest in unknown-location tasks. Section 5 illustrates how these analyses can be used to optimize regularized reconstruction methods with respect to detection performance, and Section 6 discusses these results.

2. DETECTION TASK

A. Object Model

Let f denote the true object being imaged (or its approximation in \mathbb{R}^{n_p} , where n_p is the number of coefficients in a discrete representation of f). The object f consists of a *background* f_b and it may or may not also contain a spatially localized *signal* of interest f_s . We assume that, when the target signal is present within the object, it is centered at one of a finite set of locations $\ell = 1, \dots, n_L$. We denote the target signal centered at location ℓ by $f_{s,\ell}$.

In emission tomography, where the object f is a radioactivity distribution, the background could represent the distribution of radioactivity in the absence of disease and the signal could represent the additional radioactivity absorbed in the area of a lesion. Thus an additive model for the background and signal is reasonable. The detection/localization task at hand is then a decision among the following $n_L + 1$ hypotheses:

$$\begin{aligned} H_0: f &= f_b \quad (\text{signal absent}) \\ H_\ell: f &= f_b + f_{s,\ell} \quad (\text{signal present at location } \ell, \ell = 1, \dots, n_L). \end{aligned} \quad (1)$$

The background f_b and the signal $f_{s,\ell}$ are random due to patient variability. We assume that they are statistically independent under any of the signal-present hypotheses H_ℓ , $\ell = 1, \dots, n_L$. We denote the expectations of f_b and $f_{s,\ell}$ by $\bar{f}_b \triangleq E[f_b]$ and $\bar{f}_{s,\ell} \triangleq E[f_{s,\ell}] = E[f_s | H_\ell]$, respectively. We denote their covariances by $\kappa_b = \text{Cov}\{f_b\}$ and $\kappa_{s,\ell} \triangleq \text{Cov}\{f_{s,\ell}\} = \text{Cov}\{f_s | H_\ell\}$, respectively. The covariances $\kappa_{s,\ell}$, $\ell = 1, \dots, n_L$ express variability in the shape of the target signal and not in its location.

B. Image Reconstruction

To reflect medical practice, we assume that the decision is made by an observer applied on a *reconstructed image* \hat{f} . The image \hat{f} is reconstructed from a measurement y that is acquired by a tomographic imaging system. For a given object f , y is random due to imaging noise. Specifically, the entries in y are statistically independent and Poisson-distributed conditionally on f with moments

$$E\{y|f\} = \mathcal{A}f + r, \quad (2)$$

$$\text{Cov}\{y|f\} = \text{diag}\{\mathcal{A}f + r\}, \quad (3)$$

where the linear operator \mathcal{A} models the tomographic imaging system and the vector r models scatter and/or random coincidences. We assume that the reconstructed image is given by

$$\hat{f}(y) = \mathcal{Z}y \quad (4)$$

for some linear reconstructor \mathcal{Z} . The linearity assumption holds either exactly or approximately for several common tomographic reconstruction techniques.

Here we are interested in regularized image reconstructors. A simple example that satisfies the linearity assumption is the *quadratically penalized weighted least squares* (QPWLS) family of reconstructors:

$$\mathcal{Z} = (\mathcal{F} + \mathcal{R})^{-1} \mathcal{A}' \Pi^{-1}, \quad (5)$$

where “ $'$ ” denotes the adjoint of an operator (or the transpose of a matrix), \mathcal{R} is a (linear) regularization operator, $\mathcal{F} \triangleq \mathcal{A}' \Pi^{-1} \mathcal{A}$, and $\Pi \triangleq \text{diag}\{\mathcal{A}E\{f\} + r\} \approx \text{diag}\{\mathcal{A}\bar{f}_b + r\}$. (Since we are interested in detecting small perturbations on the object background, we can assume that the signal intensity is weak with respect to the background intensity.) The QPWLS

reconstruction strategy [Eq. (5)] yields an estimated image that satisfies $\hat{\mathbf{f}}(\mathbf{y}) = \arg \max_{\mathbf{f}} [- (\mathbf{y} - \mathbf{A}\mathbf{f})' \mathbf{\Pi}^{-1} (\mathbf{y} - \mathbf{A}\mathbf{f}) - \mathbf{f}' \mathcal{R} \mathbf{f}]$. Note that, when the measurement likelihood is Poisson, QPWLS does not correspond to maximum *a posteriori* (penalized-likelihood) reconstruction. Although in the following we consider the QP-WLS reconstructor for simplicity, it is straightforward to extend our analysis to more general penalized-likelihood reconstructors using linearization approximations [12].

A commonly used form for the regularizer \mathcal{R} is that of a quadratic *roughness penalty*, such that

$$\mathbf{f}' \mathcal{R} \mathbf{f} = \beta \sum_{j=1}^{n_p} \sum_{k \in \mathcal{N}_j} (f_j - f_k)^2, \quad (6)$$

where f_j is the intensity of the object \mathbf{f} at the j th pixel, \mathcal{N}_j is a *neighborhood* of pixels around the j th pixel, and $\beta \geq 0$ is a *regularization parameter*. Regularizers of the form of Eq. (6) penalize differences between neighboring image pixels, thus favoring smoother images. For simplicity, we consider a first-order neighborhood \mathcal{N}_j , consisting of the four closest (top, bottom, left, and right) neighboring pixels. Then β is the only free parameter in \mathcal{R} .

C. Observer Model

Following the literature, we consider observers whose decision rule relies on computing some scalar *local test statistic* $t_\ell = t_\ell(\hat{\mathbf{f}}(\mathbf{y}))$ for each of the candidate locations and then comparing the *maximum test statistic*,

$$t_{\max} \triangleq \max_{\ell=1, \dots, n_\ell} t_\ell, \quad (7)$$

to a data-independent threshold τ . If $t_{\max} > \tau$, it is decided that the signal is present; otherwise, it is decided that the signal is absent.

Local test statistics that have been used to model the suboptimality of human observers are linear, and they involve a set of M bandpass filters, attempting to mimic the visual system [13]. These test statistics sample the output of the M bandpass filters at a location of interest ℓ , to obtain a local feature vector $\hat{\mathbf{c}}_\ell \in \mathbb{R}^M$, to which a *local template* $\mathbf{w}_\ell \in \mathbb{R}^M$ is then applied:

$$t_\ell = \mathbf{w}'_\ell \hat{\mathbf{c}}_\ell, \quad \hat{\mathbf{c}}_\ell = \mathbf{C}'_\ell (\hat{\mathbf{f}} - \mathbb{E}[\hat{\mathbf{f}}_b]) + \boldsymbol{\varepsilon}_\ell, \quad (8)$$

where $\mathbf{c}_\ell = [c_{1,\ell}, \dots, c_{M,\ell}]$ consists of M operators. The m th of these operators applies the impulse response of the m th bandpass filter and samples the output at location ℓ . The *internal noise* vector $\boldsymbol{\varepsilon}_\ell \in \mathbb{R}^M$ models inherent uncertainty in the observer's decisions and is typically assumed to contain independent, zero-mean, Gaussian-distributed entries. The mean of the reconstructed background, $\hat{\mathbf{f}}_b \triangleq \mathbf{Z}(\mathbf{A}\mathbf{f}_b + \mathbf{r})$, is subtracted from the reconstructed image in Eq. (8) to signify that the observer determines the most suspicious location by comparing intensities relative to the background, rather than absolute intensities [22,23].

Two such models have been proposed in the literature to model the performance of human observers in detecting unknown-location lesions in tomographic images: The *maximum channelized Hotelling observer* (MaCHO), which includes both the first- and second-order

statistics of the channel outputs in its template, and the *maximum channelized nonprewhitening observer* (MaCNPW) observer, which includes the first-order statistics only. In particular, the MaCNPW observer has been found to be well correlated with humans in simple unknown-location tasks [20,21]. This indicates that humans may not be effective at compensating for correlations in the images when they have to search for a lesion in more than one location.

Here we consider a generalization of these models that allows for an intermediate degree of prewhitening accuracy. For lack of a better term, we call this generalized model a *maximum channelized partially prewhitening* (MaCPPW) observer and define its local template at some location ℓ as

$$\mathbf{w}_\ell \triangleq [(1 - \gamma)\mathbf{I} + \gamma(\frac{1}{2}\text{Cov}\{\widehat{\mathbf{c}}|H_\ell\} + \frac{1}{2}\text{Cov}\{\widehat{\mathbf{c}}|H_0\})]^\dagger \cdot (\mathbf{E}[\widehat{\mathbf{c}}|H_\ell] - \mathbf{E}[\widehat{\mathbf{c}}|H_0]), \quad (9)$$

where “ \dagger ” denotes the pseudoinverse and $\gamma \in [0,1]$. For $\gamma = 0$ and $\gamma = 1$, the above becomes equal to the well-known MaCNPW and MaCHO models respectively. An intermediate value of γ yields an observer with intermediate prewhitening accuracy. The expression in Eq. (9) is only one of several ways in which a prewhitening deficiency could be introduced in the channelized observer’s template. We use Eq. (9) here simply to evaluate our performance approximations, without making any claims about correlation with human observers.

D. Detection Performance

Our goal is to optimize the reconstructor \mathbf{z} in general, or β in the particular case of QPWSL, with respect to the detection performance a specific observer. Since we are ultimately interested in optimizing reconstruction parameters locally, we focus on an ROC curve as a metric of the observer’s signal detection performance within a search area around a given location $\ell \in \{1, \dots, n_L\}$. Specifically, we consider the ROC curve that is obtained by plotting the *probability of detection* (deciding that the signal is present when it is actually present at ℓ),

$$P_{D_\ell}(\tau) \triangleq \mathbf{P}\{t_{\max} \geq \tau | H_\ell\}, \quad \text{for some } \ell \in \{1, \dots, n_L\}, \quad (10)$$

versus the *probability of false alarm* (deciding that the signal is present when it is actually absent),

$$P_{FA}(\tau) \triangleq \mathbf{P}\{t_{\max} \geq \tau | H_0\}. \quad (11)$$

The curve is traced by varying the threshold τ . To trace an ROC or LROC curve for the test statistic t_{\max} in Eq. (7), we need the cumulative distribution function (CDF) of t_{\max} , from which we can compute the threshold-exceeding probabilities (10) and (11).

We are interested in an observer model that searches for the signal within an area consisting of n_L voxels around some voxel of interest ℓ . As seen in Eq. (10), we use the probability of detection at the voxel ℓ to quantify detection performance. Our motivation for this analysis is optimizing the regularization parameter for some voxel ℓ to maximize detectability at that voxel, an optimization that is to be repeated for each voxel in the image. An alternative approach would be to consider the overall probability of detection within the search area,

which would require defining *a priori* probabilities of the signal presence at each location in the search area, but this is not the approach that we follow here.

If the local test statistics t_ℓ , $\ell = 1, \dots, n_L$, were statistically independent, then the CDF of their maximum would be simply equal to the product of their individual CDFs. Thus, under the assumption of independent, Gaussian-distributed local test statistics, it is possible to derive the area under the LROC curve for observer models of the form of Eqs. (7) and (8) [26]. Furthermore, even with a somewhat weaker statistical independence assumption, one can show that maximizing the area under the ROC curve (AUC) of t_{\max} would be equivalent to maximizing the area under its LROC curve [6].

Typical image reconstruction methods, and regularized ones in particular, produce images where the intensities at neighboring locations are correlated. Thus local test statistics at neighboring locations are also correlated, with more regularization leading to wider autocorrelation functions. Analyses that use independence assumptions can be applied to tasks where the candidate locations are at a distance from each other that is greater than the autocorrelation width. For tasks where all the pixels within some search area are candidate locations, correlations between the local test statistics must be taken into account.

Threshold-exceeding probabilities such as (10) and (11) are difficult to obtain in closed form when the t_ℓ 's are correlated, even if their joint distribution is available, since the exact distribution of the maximum of correlated random variables has the form of a multiple integral. As a result, investigators have proposed to trace the LROC of t_{\max} for images reconstructed from tomographic data using regularized methods via simulations [22,23]. We propose an alternative approach to evaluating the performance of the t_{\max} observer in the presence of correlations. Although closed-form expressions for the threshold-exceeding probabilities of t_{\max} are generally not available for correlated Gaussian random fields, approximations of these probabilities for high values of the threshold τ have been developed. We use them here to trace a portion of the ROC curve.

3. THRESHOLD-EXCEEDING PROBABILITIES OF THE MAXIMUM TEST STATISTIC

By analyzing the Euler characteristic of excursion sets, Adler has derived approximations for the distribution tails of the maximum of a correlated random field [24,27]. In particular, if $t_{\max} = \max_{\mathbf{x} \in S} T(\mathbf{x})$ is the maximum value of a 2-D stationary random field $T(\mathbf{x}) = T(x_1, x_2)$ over a set S , then the probability of t_{\max} exceeding a high threshold τ is approximately

$$P\{t_{\max} \geq \tau\} \approx \sum_{d=0}^2 R_d(S) \rho_d(\tau), \quad (12)$$

where the factors $R_d(S)$, $d=0,1,2$, depend on the geometry of the search area S and the functions $\rho_d(\tau)$, $d=0,1,2$, depend on the distribution of $T(\mathbf{x})$.

Approximation (12) is most accurate for search areas that are convex with sufficiently smooth boundaries [24]. If the search area S is a disk of radius r , then $R_0(S) = 1$, $R_1(S) = \pi r$, $R_2(S) = \pi r^2$. If the stationary field $T(\mathbf{x})$ is Gaussian-distributed with *zero mean*, variance σ_T^2 , and autocovariance function $R_T(\mathbf{x}) = R_T(x_1, x_2)$, then we can denote the dependence of $\rho_d(\cdot)$ on the moments of the field by writing $\rho_d(\tau) = \rho_d(\tau; \sigma_T, \Lambda_T)$, where Λ_T is the 2×2 matrix with the ij th element equal to $\{\Lambda_T\}_{ij} = -\partial^2 R_T(0, 0) / \partial x_i \partial x_j$, $i, j = 1, 2$. In particular, we have [28]

$$\rho_0(\tau; \sigma_T, \Lambda_T) \triangleq 1 - \Phi(\tau/\sigma_T), \quad (13)$$

$$\rho_1(\tau; \sigma_T, \Lambda_T) \triangleq \frac{|\det \Lambda_T|^{1/4}}{2\pi\sigma_T} e^{-\tau^2/2\sigma_T^2}, \quad (14)$$

$$\rho_2(\tau; \sigma_T, \Lambda_T) \triangleq \frac{|\det \Lambda_T|^{1/2}}{(2\pi)^{3/2}\sigma_T^3} \tau e^{-\tau^2/2\sigma_T^2}, \quad (15)$$

where $\Phi(\cdot)$ is the standard normal CDF. Approximations of the form of (12) have been applied to the problem of detecting activation in functional neuroimaging [28,29]. According to Worsley *et al.*, they lead to satisfactory accuracy for tail probabilities as high as 0.2 [28].

The analysis leading to approximation (12) assumes a continuous random field. When the field is defined on a lattice, the results hold asymptotically as the lattice becomes finer [24]. Thus, expression (12) can be applied to approximate the tail distribution of the maximum test statistic in Eq. (7) when the discrete local test statistics t_ℓ , $\ell = 1, \dots, n_L$ can be considered stationary and their mean and autocovariance is known.

Under the signal-absent hypothesis, the t_ℓ 's may be considered stationary to within the accuracy of some local shift-invariance approximations discussed in the following section. The same can be said in terms of the second-order statistics of the t_ℓ 's under the signal-present hypothesis, assuming that the contribution of the signal profile variability to the covariance of the t_ℓ 's is insignificant. However, the mean of the t_ℓ 's cannot be considered constant throughout the search area in the presence of a spatially localized target signal such as the ones that we are interested in. Thus, expression (12) is not appropriate for the signal-present hypothesis.

An alternative approach for approximating threshold-exceeding probabilities in the signal-present case follows an argument by Siegmund and Worsley [25]. This approach decomposes the probability of detection at some location \mathbf{x}_ℓ as

$$\begin{aligned} \mathbf{P}\{t_{\max} \geq \tau | H_\ell\} &= \mathbf{P}\{T(\mathbf{x}_\ell) \geq \tau | H_\ell\} + \mathbf{P}\{t_{\max} \geq \tau, T(\mathbf{x}_\ell) < \tau | H_\ell\} = 1 - \Phi\left(\frac{\tau - \mu_T(\mathbf{x}_\ell)}{\sigma_T}\right) + \\ &\int_0^\infty \mathbf{P}\{t_{\max} \geq \tau | T(\mathbf{x}_\ell) = \tau - s, H_\ell\} \times \varphi\left(\frac{\tau - s - \mu_T(\mathbf{x}_\ell)}{\sigma_T}\right) ds, \end{aligned} \quad (16)$$

where $\Phi(\cdot)$ and $\varphi(\cdot)$ are the standard normal CDF and probability distribution function (PDF), respectively, and $\mu_T(\mathbf{x}) \triangleq \mathbf{E}[T(\mathbf{x}) | H_\ell]$. Assuming that the maximum is most likely to occur near \mathbf{x}_ℓ , i.e., near the center of the target signal, the integrand in Eq. (16) will be nonnegligible only for small values of s .

To derive an expression for the conditional probability inside the integral, Siegmund and Worsley also assume that the field $T(\mathbf{x})$, in the signal-present case and in the immediate neighborhood of the target signal, can be approximated as quadratic in \mathbf{x} . Under these

assumptions, they show that the conditional probability in the integrand of Eq. (16) is approximately

$$P\{t_{\max} \geq \tau | T(\mathbf{x}_\ell) = \tau - s, H_\ell\} \approx P\{-\dot{T}(\mathbf{x}_\ell)' E[\ddot{T}(\mathbf{x}_\ell) | T(\mathbf{x}_\ell) = \tau - s]^{-1} \dot{T}(\mathbf{x}_\ell) \geq 2s\}, \quad (17)$$

where $\dot{T}(\cdot)$ and $\ddot{T}(\cdot)$ denote the gradient and Hessian with respect to \mathbf{x} , respectively, evaluated at \mathbf{x}_ℓ . For Gaussian random fields, the conditional expectation in Eq. (17) can be obtained using

$$E[\dot{T}(\mathbf{x}) | T(\mathbf{x})] = E[\dot{T}(\mathbf{x})] + \text{Cov}\{\dot{T}(\mathbf{x}), T(\mathbf{x})\} \frac{T(\mathbf{x}) - E[T(\mathbf{x})]}{\text{Var}\{T(\mathbf{x})\}} \quad (18)$$

and the linearity of derivatives. Approximating the field around the center of the target signal as quadratic and circularly symmetric yields a conditional expectation in Eq. (18) that is proportional to the identity matrix, and thus Eq. (17) corresponds to a χ^2 CDF. Substituting this CDF in Eq. (16) results in the following approximation for the probability of detection:

$$P\{t_{\max} > \tau | H_\ell\} \approx 1 - \Phi\left(\frac{\tau - \mu_T(\mathbf{x}_\ell)}{\sigma_T}\right) + \varphi\left(\frac{\tau - \mu_T(\mathbf{x}_\ell)}{\sigma_T}\right) \frac{1}{\sigma_T} \left| \frac{\partial^2 R_T(\mathbf{0})}{\partial^2 x_i} \right| / \left| \frac{\partial^2 \mu_T(\mathbf{x}_\ell)}{\partial^2 x_i} \right|. \quad (19)$$

In the following, we will combine expressions (12) and (19) with locally shift-invariant approximations to compute, respectively, the probabilities of false alarm and detection for observers of the form of Eq. (8).

4. APPROXIMATIONS OF DETECTION PERFORMANCE

A. Moments of the Local Test Statistics

Let $\mathbf{t} \in \mathbb{R}^{n_L}$ be a vector whose components are the local test statistics t_ℓ , $\ell = 1, \dots, n_L$. Let $\mathcal{V} = [\mathbf{v}_1, \dots, \mathbf{v}_{n_L}]$ be a matrix whose columns consist of the image-domain templates $\mathbf{v}_\ell \triangleq \mathbf{c}_\ell \mathbf{w}_\ell$. The latter are vectors in the same space as the reconstructed object $\hat{\mathbf{f}}$ and are formed by linear combinations of the channel responses.

To derive the moments of the test statistic vector \mathbf{t} under each of the hypotheses in Eq. (1), we combine Eq. (8) with the moments of the tomographic data \mathbf{y} from Eqs. (2) and (3), the linear reconstructor of Eq. (4), and the assumption that \mathbf{f}_b and $\mathbf{f}_{s,\ell}$ are independent. This yields the following expressions for the mean $\boldsymbol{\mu}_t$ and covariance $\boldsymbol{\Pi}_t$ of the test statistic vector under each of the signal-absent and signal-present hypotheses:

$$\boldsymbol{\mu}_{t|H_0} = \mathbf{0}, \quad (20)$$

$$\boldsymbol{\mu}_{t|H_\ell} = \mathcal{V}' \mathcal{Z} \mathcal{A} \bar{\mathbf{f}}_{s,\ell}, \quad \ell = 1, \dots, n_L, \quad (21)$$

$$\mathbf{\Pi}_{i|H_\ell} = \mathcal{V}' \mathbf{Z} \mathbf{\Pi}_\ell \mathbf{Z}' \mathcal{V} + \text{diag}\{\mathbf{w}'_\ell \mathbf{\Pi}_{\varepsilon_\ell} \mathbf{w}_\ell, \ell=1, \dots, n_L\}, \ell=0, \dots, n_L, \quad (22)$$

where $\mathbf{\Pi}_{\varepsilon_\ell}$ is the $M \times M$ covariance matrix of the internal noise vector ε_ℓ and $\mathbf{\Pi}_\ell \triangleq \text{Cov}\{\mathbf{y}|H_\ell\}$, i.e.,

$$\mathbf{\Pi}_0 = \text{diag}\{\mathcal{A} \bar{\mathbf{f}}_b + r\} + \mathcal{A} \mathcal{K}_b \mathcal{A}', \quad (23)$$

$$\mathbf{\Pi}_\ell = \mathbf{\Pi}_0 + \text{diag}\{\mathcal{A} \bar{\mathbf{f}}_{s,\ell}\} + \mathcal{A} \mathcal{K}_{s,\ell} \mathcal{A}', \quad \ell=1, \dots, n_L. \quad (24)$$

For typical problem sizes, applying the operators $\mathbf{Z} \mathcal{A}$ and $\mathbf{Z} \mathbf{\Pi}_\ell \mathbf{Z}'$ explicitly to compute the moments of \mathbf{t} in Eqs. (21) and (22) is time-consuming for common statistical image reconstruction methods. Therefore, this approach would be impractical when one needs these moments to compute measures of detectability for many values of some reconstruction parameter that is to be optimized.

We use locally shift-invariant analysis to derive Fourier-domain approximations to the moments in Eqs. (21) and (22). A similar approach to computing the moments of the local test statistics has been followed and evaluated by others [22,23]. In this work we will combine such moment approximations with expressions (12) and (19) to speed up the computation of the ROC-related probabilities (11) and (10), respectively.

B. Locally Shift-Invariant Moment Approximations

For the Fourier analysis that follows, we will assume that we have a discrete representation of the object $\mathbf{f} \in \mathbb{R}^{n_p}$. Let \mathcal{U} be a discrete Fourier operator, mapping an object in \mathbb{R}^{n_p} to some vector in \mathbb{C}^{n_p} known as the *spectrum* of the object. The Fourier operator \mathcal{U} is linear and, due to the shift property of the Fourier transform, it can be defined through its action on an object \mathbf{e}_0 that consists of an impulse centered at the origin:

$$\mathcal{U} \mathbf{e}_0 = \frac{1}{\sqrt{n_p}} \mathbf{1}, \quad (25)$$

where $\mathbf{1}$ is the vector of n_p ones. Without loss of generality, we choose the origin of the Fourier transform to coincide with the location where the signal \mathbf{f}_s is centered. (The $1/\sqrt{n_p}$ factor ensures orthonormality.)

We consider tasks where the mean target signal $\mathbf{f}_{s,\ell}$ has the same shape at all candidate locations $\ell = 1, \dots, n_L$. Then $\mathbf{f}_{s,\ell}$ is a copy of a common mean signal profile \mathbf{f}_s shifted to location ℓ . We also assume that the observer's channel responses at different image locations are shifted copies of each other. Thus we write

$$\bar{\mathbf{f}}_{s,\ell} = \mathcal{U}^{-1} \mathbf{E}_\ell \mathbf{X}_0, \quad (26)$$

$$C_\ell = \mathbf{U}^{-1} \mathbf{E}_\ell \mathbf{T}_0, \quad (27)$$

where $\mathbf{E}_\ell \triangleq \sqrt{n_p} \text{diag}\{\mathcal{U}e_\ell\}$ consists of the complex exponential that corresponds to a circular shift from the origin (0,0) to location ℓ , and e_ℓ is an impulse centered at location ℓ . Finally, \mathbf{X}_0 is the spectrum of \mathbf{f}_s , when the latter is assumed centered at (0,0), and the M columns of \mathbf{T}_0 contain the frequency responses of the observer's M channels, when their impulse responses are assumed centered at (0,0).

We now assume that the mean target signal profile \mathbf{f}_s is localized in space. In general, the operator \mathcal{F} is shift-variant for typical tomographic systems. However, if \mathcal{F} is applied to an object that is contained within a small spatial extent around the origin, we can approximate \mathcal{F} as shift-invariant within the spatial extent of this object. The closer the object comes to resembling an impulse, the greater the accuracy of this approximation. Thus we approximate the operator \mathcal{F} (and similarly \mathcal{R}) within a small spatial extent as

$$\mathcal{F} \approx \mathbf{U}^{-1} \mathbf{\Lambda} \mathbf{U}, \quad (28)$$

$$\mathcal{R} \approx \mathbf{U}^{-1} \mathbf{\Omega} \mathbf{U}, \quad (29)$$

where $\mathbf{\Lambda} \triangleq \text{diag}\{\lambda_k, k=1, \dots, n_p\}$ and $\mathbf{\Omega} \triangleq \text{diag}\{\omega_k, k=1, \dots, n_p\}$. The λ_k 's and ω_k 's contain the local frequency response of \mathcal{F} and \mathcal{R} , respectively, computed at the location of the signal. Such local diagonalization approximations have been used by other investigators to analyze observer performance with penalized-likelihood reconstruction [9,10,22,23].

Similarly, we assume that under any single hypothesis the object can be approximated as locally stationary around each of the locations $\ell = 1, \dots, n_L$. This is more likely to be accurate when the signal profile is deterministic under all signal present hypotheses (*i.e.*, $\mathbf{K}_{s,\ell} = \mathbf{0}$, $\ell = 1, \dots, n_L$) or when the signal profile variability is negligible when compared to the background variability. Thus we approximate the object covariance within some small spatial extent by

$$\mathcal{K}_b \approx \mathbf{U}^{-1} \mathbf{N}_0 \mathbf{U}, \quad (30)$$

$$\mathcal{K}_b + \mathcal{K}_{s,\ell} \approx \mathbf{U}^{-1} \mathbf{N}_\ell \mathbf{U}, \quad \ell=1, \dots, n_L, \quad (31)$$

where $\mathbf{N}_\ell \triangleq \text{diag}\{v_k^\ell, k=1, \dots, n_p\}$. The v_k^ℓ 's contain the local power spectrum of \mathbf{f} for hypothesis H_ℓ .

We take $\mathbf{\Lambda}$, $\mathbf{\Omega}$, and \mathbf{N}_ℓ to contain the local frequency responses of the respective operators at location ℓ_c , where $\ell_c \in \{1, \dots, n_L\}$ corresponds to the center of the observer's search area. Thus we assume that these frequency responses capture the approximate behavior of the respective operators throughout the search area, if the latter is not too big.

We present here only the final expressions and refer the reader elsewhere for step-by-step derivations [19, Section 4.C]. Substituting the Fourier representations (26)–(31) into the MaCPPW local template in Eq. (9) yields

$$\mathbf{w}_\ell \approx [(1 - \gamma)\mathbf{I} + \gamma(\mathbf{T}'_0[\mathbf{\Lambda} + \mathbf{\Omega}]^{-2}[\mathbf{\Lambda} + \mathbf{\Lambda}^2\check{\mathbf{N}}_\ell]\mathbf{T}_0 + \mathbf{\Pi}_{\varepsilon_\ell})]^\dagger \cdot \mathbf{T}'_0(\mathbf{\Lambda} + \mathbf{\Omega})^{-1}\mathbf{\Lambda}\mathbf{X}_0, \quad (32)$$

where $\check{\mathbf{N}}_\ell \triangleq \frac{1}{2}\mathbf{N}_\ell + \frac{1}{2}\mathbf{N}_0$. By using the same $\mathbf{\Lambda}$, $\mathbf{\Omega}$, and \mathbf{N}_ℓ (*i.e.*, the ones computed at the center ℓ_c) for every \mathbf{w}_ℓ , $\ell = 1, \dots, n_L$, we are approximating the templates as locally shift-invariant within the search area, *i.e.*, $\mathbf{w}_\ell \approx \mathbf{w}_{\ell_c}$, $\ell = 1, \dots, n_L$. Then we can write

$$\mathcal{V} \approx \sqrt{n_p}\mathcal{U}^{-1}\text{diag}\{\mathbf{V}_0\}\mathcal{U}\mathcal{I}_{SA}, \quad (33)$$

where $\mathbf{V}_0 \triangleq \mathbf{T}_0\mathbf{w}_\ell$ is a linear combination of the channel responses and $\mathcal{I}_{SA} \triangleq [e_1 | \dots | e_{n_L}]$ has as its columns the impulses e_ℓ , $\ell = 1, \dots, n_L$, corresponding to the n_L locations in the search area.

Finally, substituting the frequency-domain representations (26)–(33) into the moments of the test statistic vector from Eqs. (21) and (22) yields the following locally shift-invariant approximations of these moments:

$$\mu_{i|H_\ell} \approx n_p \mathcal{I}'_{SA} \mathcal{U}^{-1} \Phi \mathcal{U} e_\ell, \quad \ell = 1, \dots, n_L, \quad (34)$$

$$\mathbf{\Pi}_{i|H_\ell} \approx n_p \mathcal{I}'_{SA} \mathcal{U}^{-1} \Psi \mathcal{U} \mathcal{I}_{SA} + (\mathbf{w}'_{\ell_c} \mathbf{\Pi}_{\varepsilon_{\ell_c}} \mathbf{w}_{\ell_c}) \mathcal{I}, \quad \ell = 0, \dots, n_L, \quad (35)$$

where

$$\Phi \triangleq \text{diag} \left\{ \frac{V_k^* X_k \lambda_k}{\lambda_k + \omega_k}, \quad k = 1, \dots, n_p \right\},$$

$$\Psi \triangleq \text{diag} \left\{ \frac{|V_k|^2 \lambda_k (1 + \lambda_k V_k^2)}{(\lambda_k + \omega_k)^2}, \quad k = 1, \dots, n_p \right\},$$

and V_k, X_k are the k th elements of $\mathbf{V}_0, \mathbf{X}_0$, respectively.

C. Approximations of P_{FA} and P_D

Regardless of the exact distribution of the data \mathbf{y} , the term of t_ℓ in Eq. (8) that is linear in $\hat{\mathbf{f}}$ can be considered approximately Gaussian due to the central limit theorem. Thus the local test statistics can be considered approximately Gaussian. Assuming that the t_ℓ 's form a correlated Gaussian random field, we will use the approximations in (12) and (19) to evaluate the probabilities of false alarm and detection in Eqs. (11) and (10).

Approximation (12) assumes stationarity. For a typical shift-variant tomographic system, the t_ℓ 's are not globally stationary. However, as discussed above, we will consider the system to be locally shift-invariant and the object to be locally stationary over a small area around

each of the candidate locations $\ell = 1, \dots, n_L$. This implies that the t_ℓ 's are approximately stationary locally around each candidate location.

We use approximation (12) to calculate the probability of false alarm in Eq. (11) for channelized linear local test statistics and high detection thresholds τ . In particular, we write

$$P_{\text{FA}}(\tau) \approx \sum_{d=0}^2 R_d(S) \rho_d(\tau; \sigma_{T|0}, \Lambda_{T|0}) \quad (36)$$

where we use the notation $\sigma_{T|\ell}, \Lambda_{T|\ell}$ to refer to moments under hypothesis H_ℓ . These moments can be computed fast from the Fourier approximation (35). In particular,

$$\sigma_{T|0}^2 \approx \sum_{k=1}^{n_p} |V_k|^2 \lambda_k (1 + \lambda_k v_k^0) / (\lambda_k + \omega_k)^2 + \mathbf{w}'_{\ell_c} \mathbf{\Pi}_{\varepsilon_{\ell_c}} \mathbf{w}_{\ell_c}. \quad (37)$$

It would be possible to use an approximation like (36) for the probability of detection (10) as well, simply replacing $\rho_d(\tau; \sigma_{T|0}, \Lambda_{T|0})$ with $\rho_d(\tau - \mu_{T|\ell_c}; \sigma_{T|\ell_c}, \Lambda_{T|\ell_c})$ [30]. This would be satisfactory if the mean of the test statistic field, $\mu_T(\mathbf{x}) \triangleq E[T(\mathbf{x})|H_\ell]$, were constant throughout the search area under each of the signal-present hypotheses. However, the mean cannot be considered constant throughout the search area in the presence of a spatially localized target signal. In this case, using a constant $\mu_T(\mathbf{x}) = \mu_T(\mathbf{x}_{\ell_c})$ would lead to an overestimation of the probability of detection. Instead, we will use the signal-present approximation (19) for the probability of detection, combining it with Fourier approximations of the local test statistic moments:

$$P_{\text{D},\ell_c}(\tau) \approx 1 - \Phi\left(\frac{\tau - \mu_{T|\ell_c}}{\sigma_{T|\ell_c}}\right) + \varphi\left(\frac{\tau - \mu_{T|\ell_c}}{\sigma_{T|\ell_c}}\right) \cdot \frac{1}{\sigma_{T|\ell_c}} \left| \frac{\partial^2 R_{T|\ell_c}(\mathbf{0})}{\partial^2 x_i} \right| \left| \frac{\partial^2 \mu_{T|\ell_c}(\mathbf{x}_{\ell_c})}{\partial^2 x_i} \right|. \quad (38)$$

We compute $\mu_{T|\ell_c}$ and $\sigma_{T|\ell_c}$ from the locally shift-invariant approximations (34) and (35), respectively:

$$\mu_{T|\ell_c} \approx \sum_{k=1}^{n_p} V_k^* X_k \lambda_k / (\lambda_k + \omega_k), \quad (39)$$

$$\sigma_{T|\ell_c}^2 \approx \sum_{k=1}^{n_p} |V_k|^2 \lambda_k (1 + \lambda_k v_k^{\ell_c}) / (\lambda_k + \omega_k)^2 + \mathbf{w}'_{\ell_c} \mathbf{\Pi}_{\varepsilon_{\ell_c}} \mathbf{w}_{\ell_c}. \quad (40)$$

In the special case where the target signal is variable in its location only and not in its shape, *i.e.*, $\kappa_{\ell_c} = \mathbf{0}$, we have $v_k^{\ell_c} = v_k^0$, which implies $\sigma_{T|\ell_c} = \sigma_{T|0}$ and $\Lambda_{T|\ell_c} = \Lambda_{T|0}$. Finally, we obtain the two derivatives in approximation (38) by applying finite differences to approximations (34) and (35).

In the following, we will use the approximations of the probability of false alarm and detection given in (36) and (38), respectively, to compute high-threshold points of the ROC curves corresponding to various observer models of interest.

5. SIMULATION RESULTS

We present an example of using the above approximations to evaluate the effect of regularization on the detection performance of channelized observers applied to QPWLS-reconstructed images. In this example, \mathcal{A} corresponds to a 2-D PET system model with the characteristics of a CTI ECAT 931 scanner (matrix size 128×128 , pixel size 4.7 mm, 192 radial samples with 3.1 mm spacing, 160 projection angles over 180°), as implemented in the ASPIRE software package [31]. We scale \mathcal{A} to produce measurements with a total of 10^6 counts. We simulate the background events \mathbf{r} by adding 10% of the total counts uniformly distributed over all measurement bins. The object background \mathbf{f}_b has a Gaussian autocorrelation function with a full width at half-maximum (FWHM) of 12 pixels and a standard deviation of 0.075. The mean background $\bar{\mathbf{f}}_b$ is the anthropomorphic phantom shown in Fig. 1, which corresponds to a slice of the Zubal phantom [32]. The target signal $\mathbf{f}_{s,\ell}$ has a known Gaussian profile with a FWHM of 4 pixels and an amplitude of 0.3, but unknown location ℓ . The search area is assumed to be a disk. Any pixel inside the disk is a candidate location for the target signal. Figure 1 shows the mean background, the signal profile, and the largest of the concentric disk search areas that we consider in this example.

We assume a channelized observer model that involves four circularly symmetric, square-profile channels, as defined by Abbey *et al.* [17]. We consider local templates of the form of Eq. (9), where the CHO or CNPW templates are limiting cases. We evaluate detection performance for QPWLS reconstruction with the uniform quadratic regularizer in Eq. (6) and various values of the regularization parameter β . We compare results obtained from the proposed analytical approximations to results obtained from simulations.

(i) Proposed analytical approach

Using the analytical approximations (38) and (36), respectively, we calculate the probabilities of detection and false alarm for different values of the threshold τ . We repeat this for each value of β . Then we apply interpolation to these results to find the probability of detection at a fixed value of the probability of false alarm (0.02 in this example) for every β .

(ii) Empirical approach

We obtain the moments of the local test statistics for a range of different β 's from the local shift-invariant approximations (34) and (35). We then produce 5×10^5 realizations of the local test statistics, t_ℓ , $\ell = 1, \dots, n_L$, under each of the signal-present and signal-absent hypotheses, drawing from a Gaussian distribution with the respective moments. This yields multiple realizations of t_{\max} under each hypothesis. For each β , we compare the realizations of t_{\max} to a range of different thresholds τ to estimate the probabilities of detection and false alarm at every threshold. We then compute, as with the analytical approach, the probability of detection at a fixed value of the probability of false alarm for every β . We also compute the AUC for each β by numerical integration of the empirical probability of detection versus the empirical probability of false alarm.

Both the analytical and empirical methods perform a fast computation of the moments of the local test statistics based on local shift-invariance approximations. Khurd and Gindi followed a similar approach to compute the area under the LROC curve, using local shift-invariant moment approximations for CHO and CNPW local templates [22]. They compared

the results to those produced by applying the observer templates to reconstructions of multiple noisy data sets. They reported very good agreement (for a range of signal contrasts) and a 10^3 -fold reduction in computation time when using local shift-invariant analysis. Thus, we do not repeat such an evaluation of local shift-invariance approximations. The comparison herein focuses on the approximation introduced by (38) and (36), rather than the approximation introduced by assuming a shift-variant system to be locally shift-invariant.

We compute the above analytical and empirical metrics for MaCPPW observers, including the MaCHO ($\gamma = 1$), MaCNPW ($\gamma = 0$), and three intermediate values of γ . Figure 2 shows results for these observers with a search area diameter of 9 pixels and QPWLS reconstruction with the roughness penalty in Eq. (6) and various values of β . The abscissa of the plots shows the resolution of the QPWLS reconstruction at the center of the search area, defined as the FWHM (in pixels) of the local impulse response [33] at that location. This local FWHM is a measure of the amount of smoothing imposed by QPWLS. It equals 1 pixel for $\beta = 0$, which corresponds to unregularized WLS, and it increases as β increases.

Figure 3(a) shows how the optimal QPWLS resolution varies with the size of the search area. Figure 3(b) shows the maximum increase in the AUC of the MaCNPW observer versus the diameter of the search area. The increase is defined as the ratio $(AUC^* - AUC^0)/AUC^0$, where AUC^* and AUC^0 denote the AUC achieved by the observer when regularized QPWLS with optimal β and unregularized WLS is used, respectively.

As seen in Fig. 3(a), the optimal QPWLS resolutions for different search area sizes are around 4 pixels in this example. For reference, Fig. 4 shows QPWLS-reconstructed images with resolutions of 3, 4, and 5 pixels, all from the same data set with Poisson noise and the signal present in the center of the search area.

Figure 5 compares the analytical and empirical probabilities of detection and false alarm for the MaCNPW observer with two different search area sizes. Figures 5(a) and 5(b) show the probabilities of false alarm and detection, respectively, for a search diameter of 7 pixels and a fixed QPWLS resolution of 3 pixels. Figures 5(d) and 5(e) show the same probabilities for a search diameter of 15 pixels and a fixed QPWLS resolution of 3 pixels. The interpolated probability of detection at a fixed probability of false alarm ($P_{FA}=0.02$), with varying QPWLS resolution, is shown in Figs. 5(c) and 5(f) for a search diameter of 7 and 15 pixels, respectively.

6. DISCUSSION

The plots of Fig. 2 show that for the MaCHO observer, the peak detection performance achieved with regularized QPWLS is very close to the performance achieved with unregularized WLS. However, as the prewhitening capabilities of the observer deteriorate, the improvement achieved at the optimal amount of regularization becomes more and more significant. The MaCNPW observer is the most affected by the amount of regularization. The same trend is captured by all metrics, *i.e.*, the analytical and empirical probability of detection at a fixed low probability of false alarm, as well as the empirical AUC. The optimal resolution was slightly lower for the AUC.

As indicated by these results, the observer's prewhitening capabilities are important in determining whether optimized regularization can improve detectability significantly over unregularized reconstruction. This is in agreement with some prior results from the analysis of known-location tasks. In particular, we have shown that the gains in known-location detectability between optimized regularization and no regularization increase as the prewhitening deficiency of the observer increases [19, Section 3.4]. This occurs because, for an intermediate amount of regularization, the observer's channel outputs are closest to

white; hence prewhitening is less crucial for achieving good detection performance. Results by Qi also show that optimizing regularization parameters with respect to known-location detection performance has a more dramatic effect for a nonprewhitening (“maximum-contrast”) observer than the CHO [10]. Thus it seems that the effect of regularization depends less on whether the signal location is known or unknown to the observer and more on whether the observer can perform prewhitening. However, these two aspects of the observer model are not entirely unrelated in practice, since experimental results indicate that the prewhitening capabilities of human observers may deteriorate in unknown-location tasks [20,21]. In the simulations presented here, the size of the search area also played a role, with regularization being somewhat less beneficial for observers with very small or very large search areas. Although we show plots only for observers with square-profile channels, we obtain similar results with the difference-of-Gaussians channel sets from Abbey *et al.* [17]

Considering the range of QPWS resolutions that achieve near-peak AUC for the MaCNPW observer, as seen in Fig. 2, the variation in optimal QPWS resolution versus search diameter, as seen in Fig. 3(a), is small. Thus, a moderate uncertainty about the observer’s search area would not result in a significantly different conclusion about the optimal amount of regularization.

As seen in Fig. 3(b), the improvement in the AUC of the MaCNPW observer afforded by optimal regularization varies somewhat with the size of the search area. The maximum improvement is achieved for an intermediate search diameter (in this case, 9 pixels).

Approximation (38) for the probability of detection assumes that the test statistic field is quadratic around the target signal location. This quadratic approximation is more accurate for a search diameter of 7 pixels, as shown in Fig. 5(b), since this search area size is comparable to the support of the target signal. For a larger search area, however, this approximation breaks down, as shown in Fig. 5(e).

Even if the analytical approximations are not always accurate at predicting the value of the probability of detection, they follow the trends of the empirical curves and are maximized at the same or nearly the same QPWS resolution. Thus, analytically computed plots like the ones in Figs. 2(a), 5(b), and 5(e) can be useful for choosing the regularization parameter β to optimize the probability of detection for a given low probability of false alarm within a small search area around some pixel.

For the small search diameter of 7 pixels, the analytical approach offers a tenfold reduction in computation, when compared to the empirical approach of simulating the local test statistics from the Fourier approximations of their moments. The computational savings increase with the size of the search area. They also increase with the number of realizations produced for the empirical approach. This tenfold reduction in computation augments the 10^3 -fold reduction, which is afforded by the use of locally shift-invariant approximations instead of image reconstructions and which was reported by Khurd and Ghindi [22]. The additional savings can be important when detectability needs to be evaluated at every voxel in the image to optimize regularization parameters locally.

The limitation of the proposed method is that analytical approximations (36) and (38) apply only to high thresholds τ . Thus these approximations are useful when a partial area under the curve or a single point on the curve (at low false-positive rates) is the desired figure of merit for observer performance. However, they cannot be used to trace the entire ROC and find the AUC. In the example of Section 5, the probability of detection at a low probability of false alarm behaved similarly to the AUC with respect to the regularization parameter (see Fig. 2). In general, however, the ranking of ROC curves at low probabilities of false alarm may not predict their ranking at high probabilities of false alarm, as ROC curves for

different regularization parameters may cross. If one is interested in optimizing regularization with respect to performance at high probabilities of false alarm, one should check for crossing ROC curves by applying the simulation method for at least one voxel in the image before applying the analytical method to optimize regularization for the rest of the image. The analytical approach could also complement the empirical approach by providing an initial estimate of the optimal resolution, around which simulations can then be performed if more accuracy is needed.

We have considered observers that search for the target signal within a local search area, rather than within the entire image. We have made this choice for two reasons. First, locally shift-invariant approximations to the moments of the local test statistics are accurate within a local region. If the system is shift-variant and/or the object background is not stationary, these approximations are not accurate within the entire image. Second, our ultimate goal is to use the analyses above to optimize the amount of regularization locally at each pixel in the image. It seems reasonable for the optimal amount of regularization at some pixel to be determined by the image statistics within a local area around that pixel.

For the example presented above, we did not include internal noise in the observer model. The reason is that internal noise would manifest itself as an impulse at the center $R_T(\mathbf{0})$ of the autocovariance function. In this case, the analytical approximations would break down, because they require the autocovariance function to be smooth so that its second derivative can be computed. In particular, the analytical approximations would deteriorate for large β , where the internal noise dominates. This issue arises because the observer models with internal noise that have been proposed in the literature assume that this noise is uncorrelated between different candidate locations. Whether this assumption is crucial, however, must ultimately be validated by human observer studies. In general, internal noise is intended to model the phenomenon that a human observer may make different decisions when presented with the same image on different occasions. An alternative way to express this uncertainty that would make the model more amenable to analysis would be to add noise to the maximum test statistic itself instead of adding noise to the local test statistics individually. Once again, the usefulness of such a model would have to be tested by studies with human observers.

The analytical approximations can also be applied to optimize anisotropic regularization, which has been reported to yield greater performance improvement than isotropic regularization in known-location tasks [12]. In addition to lesion detectability, the proposed method can be applied to the optimization of regularization with respect to the detectability of activations in functional neuroimaging. Improving the accuracy of the approximations for a wider range of search area sizes and for nonzero internal noise, as well as relaxing the assumption of a circularly symmetric signal shape, are topics for future investigation. Furthermore, since the threshold-exceeding probability approximations used here have been derived for continuous random fields, the analysis can be applied naturally to tomographic reconstruction methods that yield continuous-space images \hat{f} .

Finally, it would be of interest to extend the analysis of the probability of detection to the probability of correct localization, thus considering a LROC instead of a ROC curve. Under certain assumptions of statistical independence on the local test statistics, it has been shown that maximizing the AUC is equivalent to maximizing the area under its LROC curve [6]. However, further investigation is needed to analyze the probability of correct localization for the correlated local test statistics considered here.

Acknowledgments

This work was supported in part by grant NCI P01 CA87634. The authors thank the reviewers, who improved the paper with their careful reading and thoughtful comments.

References

1. Hansen PC, O'Leary DP. The use of the L -curve in the regularization of discrete ill-posed problems. *SIAM (Soc Ind Appl Math) J Sci Stat Comput* 1993;14:1487–1506.
2. Vogel CR. Non-convergence of the L -curve regularization parameter selection method. *Inverse Probl* 1996;12:535–547.
3. Lusted LB. Signal detectability and medical decision-making. *Science* 1971;171:1217–1219. [PubMed: 5545199]
4. Metz CE. Basic principles of ROC analysis. *Semin Nucl Med* 1978;8:283–298. [PubMed: 112681]
5. Chesters MS. Human visual perception and ROC methodology in medical imaging. *Phys Med Biol* 1992;37:1433–1484. [PubMed: 1631192]
6. Swensson RG. Unified measurement of observer performance in detecting and localizing target objects on images. *Med Phys* 1996;23:1709–1725. [PubMed: 8946368]
7. Barrett, HH.; Myers, KJ. *Foundations of Image Science*. Wiley; 2003.
8. Qi J, Huesman RH. Theoretical study of lesion detectability of MAP reconstruction using computer observers. *IEEE Trans Med Imaging* 2001;20:815–822. [PubMed: 11513032]
9. Xing Y, Hsiao IT, Gindi G. Rapid calculation of detectability in Bayesian single photon emission computed tomography. *Phys Med Biol* 2003;48:3755–3774. [PubMed: 14680271]
10. Qi J. Analysis of lesion detectability in Bayesian emission reconstruction with nonstationary object variability. *IEEE Trans Med Imaging* 2004;23:321–329. [PubMed: 15027525]
11. Yendiki A, Fessler JA. Analysis of observer performance in known-location tasks for tomographic image reconstruction. *IEEE Trans Med Imaging* 2006;25:28–41. [PubMed: 16398412]
12. Qi J, Huesman RH. Penalized maximum-likelihood image reconstruction for lesion detection. *Phys Med Biol* 2006;51:4017–4030. [PubMed: 16885621]
13. Myers KJ, Barrett HH. Addition of a channel mechanism to the ideal-observer model. *J Opt Soc Am A* 1987;4:2447–2457. [PubMed: 3430229]
14. Rolland JP, Barrett HH. Effect of random background inhomogeneity on observer detection performance. *J Opt Soc Am A* 1992;9:649–658. [PubMed: 1588452]
15. Wollenweber SD, Tsui BMW, Lalush DS, Frey EC, LaCroix KJ, Gullberg GT. Comparison of Hotelling observer models and human observers in defect detection from myocardial SPECT imaging. *IEEE Trans Nucl Sci* 1999;46:2098–2103.
16. Gifford HC, King MA, de Vries DJ, Soares EJ. Channelized Hotelling and human observer correlation for lesion detection in hepatic SPECT imaging. *J Nucl Med* 2000;41:514–521. [PubMed: 10716327]
17. Abbey CK, Barrett HH. Human- and model-observer performance in ramp-spectrum noise: effects of regularization and object variability. *J Opt Soc Am A* 2001;18:473–488.
18. Sankaran S, Frey EC, Gilland KL, Tsui BM. Optimum compensation method and filter cutoff frequency in myocardial spect: a human observer study. *J Nucl Med* 2002;43:432–438. [PubMed: 11884505]
19. Yendiki, A. Ph D dissertation. University of Michigan; 2005. Analysis of signal detectability in statistically reconstructed tomographic images.
20. Gifford HC, Pretorius PH, King MA. Comparison of human- and model-observer LROC studies. *Proc SPIE* 2003;5034:112–122.
21. Gifford, HC.; Kinahan, PE.; Lartzien, C.; King, MA. Evaluation of multiclass model observers in PET LROC studies. *Proceedings of IEEE Nuclear Science Symposium and Medical Imaging Conference; IEEE; 2004.* p. 4068-4071.
22. Khurd PK, Gindi GR. LROC model observers for emission tomographic reconstruction. *Proc SPIE* 2004;5372:509–520.

23. Qi J, Huesman RH. Fast approach to evaluate MAP reconstruction for lesion detection and localization. *Proc SPIE* 2004;5372:273–282.
24. Adler RJ. On excursion sets, tube formulas and maxima of random fields. *Ann Appl Probab* 2000;10:1–74.
25. Siegmund DO, Worsley KJ. Testing for a signal with unknown location and scale in a stationary Gaussian random field. *Ann Stat* 1995;23:608–639.
26. Khurd P, Gindi G. Rapid computation of LROC figures of merit using numerical observers (for SPECT/PET reconstruction). *IEEE Trans Nucl Sci* 2005;52:618–626.
27. Adler, RJ. *The Geometry of Random Fields*. Wiley; 1981.
28. Worsley KJ, Marrett S, Neelin P, Vandal AC, Friston KJ, Evans AC. A unified statistical approach for determining significant signals in images of cerebral activation. *Hum Brain Mapp* 1996;4:58–73. [PubMed: 20408186]
29. Worsley KJ. Detecting activation in fMRI data. *Stat Methods Med Res* 2003;12:401–418. [PubMed: 14599003]
30. Yendiki, A.; Fessler, JA. Analysis of observer performance in detecting signals with location uncertainty for regularized tomographic image reconstruction. *Proceedings of IEEE Nuclear Science Symposium and Medical Imaging Conference; IEEE; 2004. p. 2620-2624.*
31. Fessler, JA. Tech Rep. Vol. 293. *Communications and Signal Processing Laboratory, Department of EECS, University of Michigan; Ann Arbor, MI, 48109-2122: 1995. ASPIRE 3.0 user's guide: a sparse iterative reconstruction library. Available at <http://www.eecs.umich.edu/~fessler>*
32. Zubal, G.; Gindi, G.; Lee, M.; Harrell, C.; Smith, E. High resolution anthropomorphic phantom for Monte Carlo analysis of internal radiation sources. *IEEE Symposium on Computer-Based Medical Systems; IEEE; 1990. p. 540-547.*
33. Fessler JA, Rogers WL. Spatial resolution properties of penalized-likelihood image reconstruction methods: space-invariant tomographs. *IEEE Trans Image Process* 1996;5:1346–1358. [PubMed: 18285223]




Fig. 1. PSfrag replacements. Mean background f_b with the target signal f_s superimposed (a) and a profile through them (b). The largest of the search areas considered (diameter of 23 pixels) is indicated as a black circle in (a).




Fig. 2. Detection performance of MaCPPW observers versus QPWLS reconstruction resolution: P_D obtained analytically (a), P_D obtained empirically (b), and AUC obtained empirically (c). Results are shown for five different degrees of prewhitening accuracy. The search area is a disk with a diameter of 9 pixels.

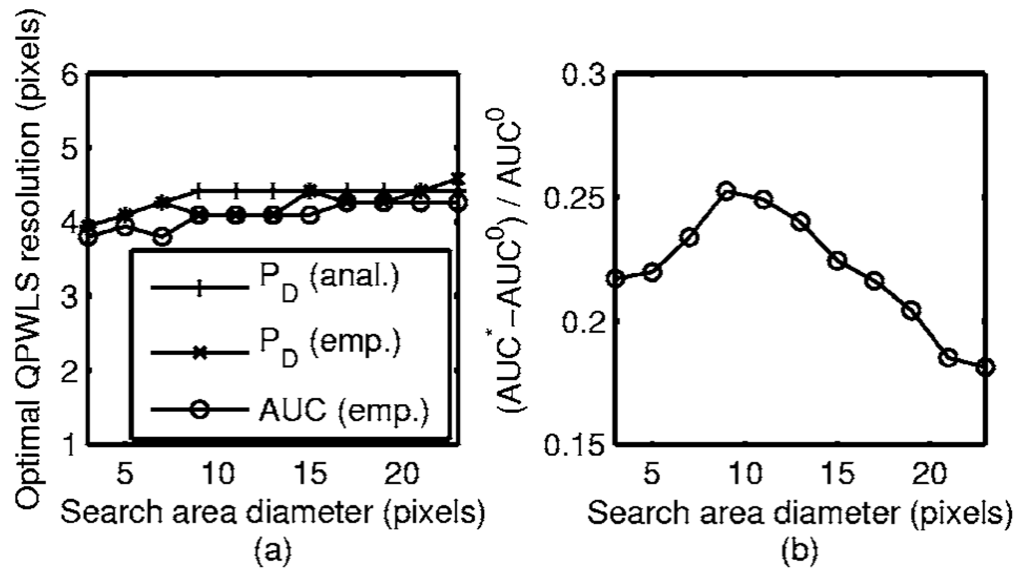


Fig. 3. QPWS resolution that maximizes the P_D (obtained analytically and empirically) or AUC (obtained empirically) versus search area diameter (a). AUC improvement with optimally regularized QPWS over unregularized WLS versus search area diameter (b). Results are shown for the MaCNPW observer.




Fig. 4. QPWLS reconstructions of a noisy Poisson data set with a resolution of 3 pixels (a), 4 pixels (b), or 5 pixels (c), and profiles through the three images (d). The signal is present in the center of the search area.




Fig. 5. Empirical and analytical probabilities of false alarm and detection for the MaCNPW observer with a search area diameter of 7 pixels [(a)–(c)] and 15 pixels [(d)–(f)]. The plots show the P_{FA} versus detection threshold [(a),(d)], the P_D versus detection threshold [(b),(e)], and the P_D (with error bars) versus QPWS resolution for a fixed $P_{FA}=0.02$ [(c),(f)]. (The error bars are small enough to fit in the plot markers.)